Second International Workshop on Agent Technology for Energy Systems

www.ates2011.org

Dr. Alex Rogers
School of Electronics and Computer Science
University of Southampton, UK
Email: acr@ecs.soton.ac.uk

Prof. Keith Decker
Department of Computer and Information Sciences
University of Delaware, Newark, USA
Email: decker@cis.udel.edu

Mr. Koen Kok
Energy Research Centre of The Netherlands (ECN)
The Netherlands
Email: j.kok@ecn.nl

Program Committee

Prof. Derek Bunn, London Business School, UK
Dr. Victoria Catterson, University of Strathclyde, UK
Prof. Paul Davidsson, Blekinge Institute of Technology, Sweden
Dr. Ying Guo, CSIRO ICT Centre, Australia
Dr. Alan Holland, University College Cork, Ireland
Dr. Christian Johansson, Blekinge Institute of Technology, Sweden
Dr. Sachin Kamboj, University of Delaware, USA
Dr. James Keirstead, Imperial College London, UK
Dr. Wolf Ketter, Rotterdam School of Management, The Netherlands
Mr. Koen Kok, Energy Research Centre, The Netherlands
Dr. Steffen Lamparter, Siemens Corporate Technology, Germany
Prof. Han La Poutre, Centrum Wiskunde & Informatica, The Netherlands
Dr. Jiaming Li, CSIRO ICT Centre, Australia
Dr. Ron Li, CSIRO ICT Centre, Australia
Mr. Andreas Ligtvoet, Delft University of Technology, The Netherlands
Dr. Robert Marks, University of New South Wales, Australia
Radu Casian Mihailescu, University Rey Juan Carlos, Spain
Prof. Julian Padget, University of Bath, UK
Dr. Steve Phelpsm, University of Essex, UK
Dr. Iyad Rahwan, Masdar Institute of Science and Technology, United Arab Emirates
Dr. Gopal Ramchurn, Unversity of Southampton, UK
Dr. Valentin Robu, Unversity of Southampton, UK
Dr. Andrew Roscoe, University of Strathclyde, UK
Dr. Koen Van Dam, University of Technology in Delft, The Netherlands
Mr. Stijn Vandael,Katholieke Universiteit Leuven, The Netherlands
Mr. Matteo Vasirani, University Rey Juan Carlos, Spain
Dr. Thomas Voice, Unversity of Southampton, UK
Terry Zimmerman, Smart Information Flow Technologies (SIFT), USA

Table of Contents

Full Papers

Proposal of Distributed Scheduling Heuristics using Mediation Agent

Takahiro Kawamura
Corporate Research & Development Center
Toshiba Corp.

Akihiko Ohsuga
Graduate School of Information Systems
University of Electro-Communications

ABSTRACT

This paper proposes an agent-based distributed scheduling system against the background of the deregulation of electric utility and the smart grid for the renewable energy, and then focuses on a maintenance scheduling in the context of real problems. A synchronous backtrack algorithm, a well-known method for distributed scheduling problems, has difficulties handling A)rapid schedule adjustments and B)impartial assignment. Thus, we propose two kinds of heuristics: 1)parallel assignment and 2)multiple priority strategies, and developed the distributed scheduling system which makes use of the heuristics.It consists of schedulers for each power station and mediation agents which have cloning and merging functions to support the implementation of the heuristics. Finally, the result of our experiment shows an improvement when handling the rapid adjustment and the impartiality issues with reasonable computational overhead.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multiagent systems

General Terms

Experimentation

Keywords

Multi-agent, Distributed Scheduling, Power Management, Smart Grid

1. INTRODUCTION

Configurations of electric power generation and its T&D (transmission and distribution) differ from one country to another. In Japan and some other countries, a single company by region integrates and controls all the power stations and their facilities in order to keep high reliability and consistency among them. However, in addition to the recent deregulation of electric utility where independent companies are allowed to supply the electric power, the latest trend of smart grid which aims at incorporation of as much as renewable energy and its efficient operation has been changing the configurations. The renewable source does not have much electric-generating capacity in comparison with fossil fuels, so that the small but many independent power generating companies will arise. Thus, a concentration of the controls over the companies has gotten more difficult in recent years. From the technical viewpoint, the concentration approach is also considered rather inconvenient because of excessive loading caused by data processing and transfer: each time additional partial modifications in the schedule arise (as a result of either a change in weather or insignificant production delays), the appropriate data have to be collected, reprocessed in the computer center, updated, and recirculated between the stations. Against this background, there has been several researches to propose a multi-agent system where an agent has a different social authority in the smart grid world. As one of these researches this paper proposes an agent-based distributed scheduling system[3] to make power outage maintenance schedules of the facilities belonging to different companies. This is assumed to be acceptable for the companies compared to the concentration approach, because they do not need to make all of the scheduling data and constraints available to another company.

Section 2 contains a brief description of the distributed maintenance scheduling method. Section 3 discusses A)rapid rescheduling adjustments and B)impartial assignment of initial power station schedules. This problem arises in the process of distributed maintenance scheduling using the synchronous backtrack algorithm, one of the typical methods used to solve a distributed constraint satisfaction problem. Therefore, the authors propose two kinds of heuristic: 1)parallel assignment and 2)multiple priority strategies. Furthermore, section 4 contains an explanation of the distributed maintenance scheduling system where these heuristics are implemented within a framework of mediation agents with cloning and merging functions. Section 5 describes the results of an experimental analysis of the effect of the proposed heuristics on A)rapid adjustments and B)impartial assignment in the context of real problems. Finally, section 6 provides analysis of the approaches adopted in this work.

2. DESCRIPTION OF MAINTENANCE SCHEDULING PROBLEM

Maintenance scheduling defines a date for the temporary interruption in electric power supply for the purpose of maintenance, performing repair work, replacement or installation of the additional power equipment (e.g., mains, transmission lines, generators, transformers). However, since the electric power system must supply electric power 24 hours a day, various constraints must be observed in relation to the specific equipment during the equipment shutdown. In addition, each piece of equipment is serviced by a local power station, therefore, local constraints of the power stations must also be satisfied. In the context of this problem, the shutdown of the equipment related to the power system is set as a parameter, and the value of such parameter is a shutdown date.

This problem can be considered as a form of constraint satisfaction problem, where both the parameters and the constraints are distributed between schedulers set up for each station. The schedulers $S_i (i = 1, ..., n)$ retain a number of jobs $J_{ij} (j = 1, ..., k)$, and their values are selected from the set $D_{ij} = \{1, ..., 365\}$. Between jobs J_{ij} and J_{nm}, **Same time constraint** P_{ijnm} and **Grouping constraint** Q_{ijnm} exist, where the former is simply to prevent unintended blackout, and the latter has two meanings: preventing overload for the rest in case of partial shutdown and efficiency of operations. The scheduler should assign maintenance operations, so that both constraints are satisfied. In Figure 1, $C_{J_{ij}}$ has a constant value (job duration), predetermined for each value J_{ij}, where J_{ij} indicates a job starting date, and $J_{ij} + C_{J_{ij}}$ correspondingly indicates a job completion date. Here we can define the constraints as follows:

J_{ij} and J_{nm} conflict with Same time constraint P_{ijnm},
if $J_{nm} \in Int(J_{ij}, C_{J_{ij}})$ or $J_{ij} \in Int(J_{nm}, C_{J_{nm}})$,
where $Int(J_{ij}, C_{J_{ij}}) \equiv (J_{ij}, J_{ij} + 1, ..., J_{ij} + C_{J_{ij}})$.

J_{ij} and J_{nm} conflict with Grouping constraint Q_{ijnm},
if $Int(J_{ij}, C_{J_{ij}}) \not\subset Int(J_{nm}, C_{J_{nm}})$, and
$Int(J_{nm}, C_{J_{nm}}) \not\subset Int(J_{ij}, C_{J_{ij}})$.

The criterion function to optimize each constraint is not determined. Often, one constraint may relate to all the jobs. Moreover, it is possible that all the jobs have constraints overlapping between the schedulers. Most of the time, scheduling is carried out by way of partial rescheduling rather than by drafting a one-time overall schedule. Scheduling is completed when all parameters satisfy appropriate constraints.

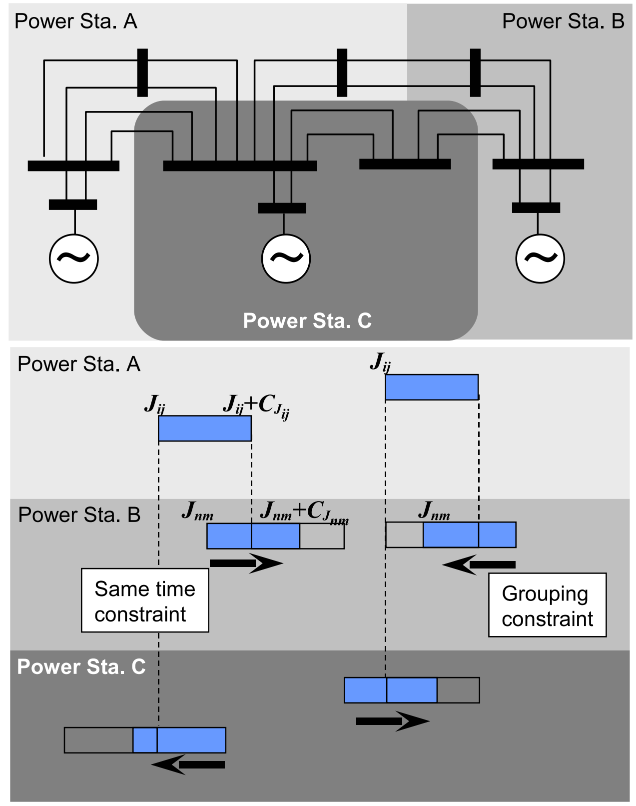

Figure 1: Same time constraint and Grouping constraint

3. PROBLEMS OF DISTRIBUTED MAINTENANCE SCHEDULING AND OUR APPROACHES

In the distributed maintenance scheduling using synchronous backtrack algorithm, which is a typical method to solve distributed constraint satisfaction problems, the following difficulties may arise.

A) Rapid adjustment

When scheduling is conducted based on the synchronous backtrack algorithm, a fixed strategy of parameter value assignment by individual schedulers, or a set of several alternately used strategies is implemented to allow resolving the constraints to the last one. Constraint resolution is strictly conducted sequentially. Therefore, rescheduling due to sudden change or modification of schedules takes some time. For example, it is impossible to find a solution as quickly as possible by using any of the assignment strategies.

B) Impartial assignment

Besides, although some parameters have equal high priority, the process of assigning values follows a certain preset order. As a result, despite the same priority, some parameters obtain their values last and appear to be in worse condition than those receiving their values first. For example, when it is desirable to minimize the deviation from initial schedules drawn up at each power station, those with a late adjustment queue will be affected by the other stations. schedules, which may result in rather considerable deviations from the initial schedule.

In order to resolve these problems, the authors propose to utilize the following two heuristics. Our contribution of this paper is a realization of those heuristics using a multi-agent system.

1) Parallel assignment

The normal assignment strategy of each individual scheduler is used to determine which value out of several possible ones should be assigned to the retained jobs. A variety of specific strategies may be utilized resulting from initial and final loading, priority of supply terms, minimization of surplus power, and so on. In this work, a method of parallel processing of the above assignment strategies is proposed. When several assignment strategies are stored in a scheduler, their search is carried out in parallel when a certain trigger occurs. Since it partially replaces breadth-first search with depth-first search, this actually constitutes parallel processing of the synchronous backtrack algorithm, which is sequential in nature. In this case, the assignment strategy is assumed to be retained in individual schedulers. Besides, the trigger is regarded as an extremity (unavailability) of the partial solution expansion (backtrack). More specific examples are given in the following section to explain system operation.

2) Multiple priority strategies

Priority determination between schedulers is to decide an order of jobs whose values have constraints related to other jobs and will be sequentially assigned. Typically, priority determination is conducted sequentially starting from either an earlier production process, or a process close to a final product, and the like. In this work, a method allowing duplication of the order of such determinations is introduced. When certain jobs have identical priority, the method of gradual adjustment is applied, according to which as soon as any of the jobs is determined as the first, another one is not added to

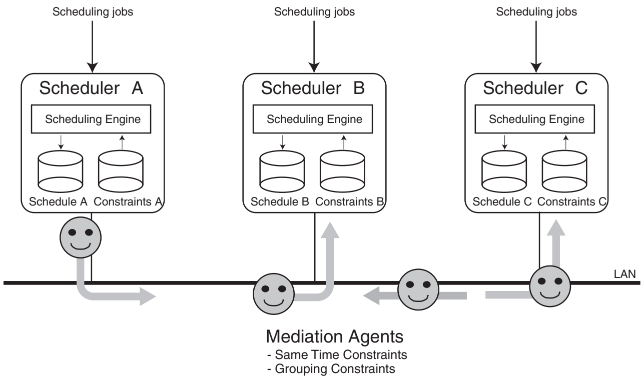

Figure 2: System architecture

it, but is gradually adjusted. According to the gradual adjustment method, the values being assigned to two parameters are lowered step by step based on the respective constraints to minimize the deviation from the initial values. However, the value lowering method depends on a type of constraint to be satisfied. It is assumed that the priority is preset for each job. More specific examples are given in the following section to explain the operation of the system.

4. AGENT-BASED MAINTENANCE SCHEDULING SYSTEM

In this work, the distributed maintenance scheduling is carried out based on our multi-agent framework Bee-gent (Bonding and Encapsulation Enhancement Agent), whose objective is to allow flexible interconnection between legacy systems, databases, and other existing systems by providing coordination and communication functions. Especially, interaction protocols describing interactions between system components are divided between global processing related to problem solving, and component-specific local processing. Coordination between system components in the global processing is ensured through consistent control by mobile mediation agents. Then, a well-known problem of distributed constraint satisfaction, which is vast amounts of direct communication between the system components (in this paper, schedulers) is replaced by local communication between the scheduler and the mediation agent. The actual implementation and evaluation for the framework are presented in [8].

In order to apply the heuristics of maintenance scheduling to a distributed scheduling system, the authors came up with a method in ease which a mediation agent repeatedly performs cloning and/or merging functions as necessary. Below, the system architecture is described, and the scheduling method using mediation agent is explained.

4.1 System architecture

This system consists of schedulers for each station and agents performing mediation functions (Figure 2).

The schedulers A, B, C control scheduling jobs entered by a user (a person responsible for drawing up schedules for each station) and arrange them in accordance with constraints. In case a constraint related to any of the entered jobs also affects a job retained by another scheduler, a mediation agent is formed to resolve the conflict. If necessary, the mediation agent operating in conjunction with an individual scheduler may be switched off. No direct exchange of messages between the schedulers is carried out.

Mediation agents ensure resolution of the overlapping constraints

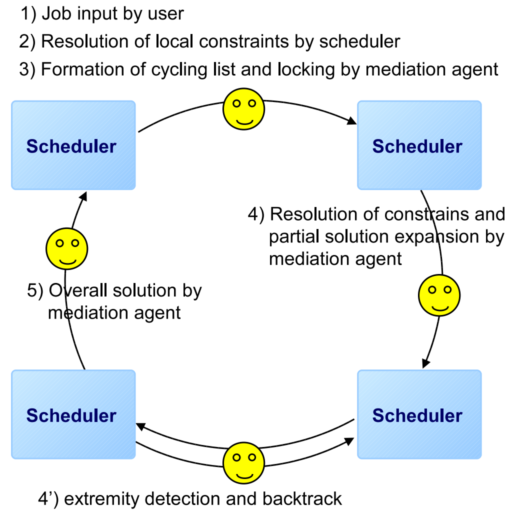

Figure 3: Synchronous backtrack using mediation agent

between the schedulers (**Same time constraint** and **Grouping constraint**). In the process of doing so, they migrate between the schedulers, as necessary, and assign the jobs. Due to the locking function described below, only one mediation agent may handle each job group related to specific constraints. Consequently, a limited number of mediation agents is formed in the entire system. Therefore, there is no need to arrange for centralized control over the mediation agents. On the other hand, several mediation agents may conduct simultaneous parallel scheduling with regard to the job groups not connected by constraints. Hence, it does not mean one mediation agent handles all the problems.

4.2 Implementation of synchronous backtrack algorithm using mediation agents

Described below is a process flow diagram of synchronous backtracking algorithm using mediation agents (the numbering in Figure 3 corresponds to the numbers utilized in the text below).

1. The user (responsible for scheduling the operation of power station) enters new jobs and constraints into a scheduler set up at each station. The user also updates information (e.g., modifies constraints) about already-scheduled jobs.

2. The scheduler assigns jobs related by constraints. The scheduler performs the initial assignment of jobs based on local constraints (e.g., number of workers, constraints associated with type of equipment). Then the scheduler checks the jobs for constraints (**Same time constraint** or **Grouping constraint**) overlapping with those of other schedulers. If they exist, the scheduler creates the mediation agent and hands over to it the relevant scheduling jobs and constraints.

3. Based on a list of schedulers related by constraints which is the minimum global information created by the user, the mediation agent creates an itinerary. It is a certain preset order (such as date order) of assigning values for the schedulers, and also a order of migration for the mediation agent. Then, it tries to lock groups of jobs related by constraints by the following locking function.

   Locking function is as follows:
   The mediation agent checks the itinerary and tries to lock

all other scheduling jobs related by constraints (to exclude any changes in job assignment). A message requesting that all corresponding jobs be locked is sent to other schedulers. Upon receipt of the message, the schedulers reply by sending a confirmation if locking is possible, and lock the scheduling job. As soon as the mediation agent receives messages confirming locking from each of the schedulers, it starts the scheduling procedure. Should a response arrive (even a single one) stating that locking is impossible, the message is sent out to all schedulers involved, canceling the locking. Then, after some time, the locking attempt is repeated. As a result, the possibility of simultaneous copying of the same job by several mediation agents or other mismatched processing is eliminated.

4. With locking successfully completed, the mediation agent begins resolving constraints according to the itinerary. However, constraint resolution is not carried out in the scheduler where the mediation agent was created, and the assignment of jobs that is obtained becomes the initial partial solution.

   Constraint resolution function is as follows:

   (a) When the mediating agent arises in the next scheduler, it requests the constraints extending over the other schedulers and the actually assigned value of the job related to the scheduling job.

   (b) The scheduler identifies the related jobs and together with constraints extending over the other schedulers transfers them to the mediation agent.

   (c) The mediation agent adjusts the assigned values with minimal deviation, so that they do not conflict with the actual partial solution and sends them back to the scheduler.

   (d) The scheduler checks whether they conflict with the local constraints. If they do, the scheduler adjusts them (except for the values equal to already-assigned values before) with minimal deviation so that they satisfy the local constraints, and transfers again to the mediation agent.

   (e) In case the agent fails to obtain the assigned value satisfying both constraints after multiple repetitions of (c),(d), it goes back to the previous scheduler (the scheduler the mediation agent visited before the current scheduler on the itinerary) and tries to obtain the assigned value different from that obtained earlier (backtracking).

   (e') Or, in case the assigned value satisfying both constraints is obtained, this new value is added to the partial solution and is stored by the mediation agent (termination). Then, the agent goes to the next scheduler.

5. Upon completion of the job assignment for all jobs connected by constraints, the mediation agent returns to the itinerary and communicates the new assigned value to each scheduler. Ultimately, the mediation agent informs the user of the completion of rescheduling and switches off.
   On the other hand, if further search is impossible, the mediation agent informs the user of an unsuccessful outcome of rescheduling and switches off. In regard to groups of scheduling jobs not relating by constraints, parallel processing can be carried out using several mediation agents.

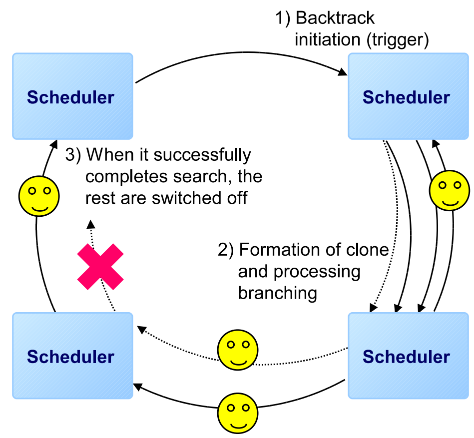

Figure 4: Parallel assignment strategy using mediation agent

4.3 Parallel assignment strategy using mediation agent cloning

Described below is a method of mediation agent cloning to realize 1)parallel assignment strategy for A)rapid adjustment when resolving **Same time constraint**. (the numbering in Figure 4 corresponds to the numbers utilized in the text below).

1. When the mediation agent finds the unavailability of the partial solution expansion (4(e) at section 4.2), it refers to the list of effective assignment strategies.

   In case that any parameter conflicts with **Same time constraint**, the scheduler can choose to utilize at least two assignment strategies by shifting its value by a minimum step back or forward. Both assignment strategies are not always applicable. For example, once the earliest start time is passed, no shifting forward is possible.

2. The mediation agent creates its own clones, each of them implementing only one assignment strategy (back or forward). Thus, parallel search for a solution is performed. A clone of the mediation agent is not allowed to perform backtracking from the scheduler where it was created to a previous one on the itinerary.

3. The very first mediation agent to find a solution satisfying all the constraints forces all other clones to switch off.

4.4 Multiple priority strategy based on merging of mediation agents

Described below is a method of merging of mediation agents to realize 2)multiple priority strategy for B)impartial assignment when resolving **Grouping constraint**. (the numbering in Figure 5 corresponds to the numbers used in the text below).

1. When several highest-priority jobs exist, the mediation agent creates clones for the schedulers containing these jobs and after setting a merging point begins resolving the constraints.

   If there are **Grouping constraint** between certain jobs, the mediation agent combines the assigned value of one job with that of another and makes modifications (grouping). However, when determining the solution order, these jobs may appear to be of equal importance. Usually, after successful completion of locking, the mediation agent interrogates the schedulers having jobs with the highest priority and assigns

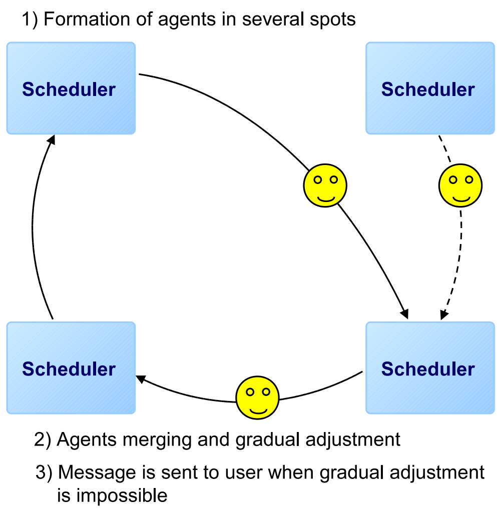

Figure 5: Multiple priority using mediation agent

Table 1: Example of priority and merging

	Priority	
Scheduler A	1	
Scheduler B	1	
Scheduler C	2	← a merging point
Scheduler D	3	

jobs in accordance with the preset priority. Since in the process of search the mediation agent gives equal treatment to jobs of equal priority, if there are several jobs with the highest priority, the mediation agent resolves the constraints by putting clones for the schedulers containing jobs after selecting the scheduler to act as a merging point. The scheduler selected to serve as a merging point is the one containing a job which has a unique priority. In the example given below (Table 1), mediation agent 1 which begins scheduling from scheduler A, and mediation agent 2 which begins scheduling from scheduler B, merge in scheduler C.

2. The mediation agent and its clone merge together in a merging point.

   The process of merging of mediation agents is explained below. The mediation agent with the earlier assigned internal ID (mediation agent 1) transfers the partial solution and constraints stored in it to the second agent (mediation agent 2). However, mediation agent 1 migrates to scheduler C, carrying constraints extending over to the other schedulers, to which it was referring when determining parameter A, and local constraints for parameter A obtained exceptionally from scheduler A. Mediation agent 2 does the same. Then, mediation agent 1 transfers the partial solution and constraints retained therein over to mediation agent 2 and switches off. Mediation agent 2 performs gradual adjustment of parameter A from scheduler A and parameter B from scheduler B in accordance with their respective constraints so that deviations in both cases would be minimal.

   The actual method of gradual adjustment varies depending on the type of constraint to be satisfied. However, the adjustment for **Grouping constraint** is simply to change the assigned parameter values until they coincide with each other at their intermediary point. For example,

   if $J_{ij} = t1$ and $J_{nm} = t2$,
   then $J_{ij} = t1 + \frac{\Delta t}{2}$ and $J_{nm} = t2 - \frac{\Delta t}{2}$,
   where $t2 > t1, \Delta t = t2 - t1$.

3. If the gradual adjustment is impossible (for example, when it is impossible to obtain assigned values satisfying both constraints or when there is a substantial change in the assigned value of only one parameter), the mediation agent informs the user that the priority must be changed and switches off.

5. EXPERIMENT AND EVALUATION

This section provides a description of experiments conducted to evaluate ways of possible improvement of the handling of A)rapid schedule adjustments and B)impartial assignment of schedules.

5.1 Experimental Setting

The number of power stations was assumed to be 10 while the total number of connected equipment was 200 units. Schedules for equipment downtime were created based on a “once a month” assumption, and the duration of works to be done was selected from the range of 1 day to 1 week depending on their complexity. The scheduling period was chosen to be 1 year. The above settings and other specified conditions are based on the configuration and scale of existing electric power systems operated by a certain power company. However, it is not the same as the real data, but our simulation. In addition, the following system constraints were taken into account:

- Grouping constraints
- Same time constraints
  - Same time constraints for the equipment in the loop
  - Same time constraints for the equipment supporting system’s reliability. Work on adjacent electrical lines should not be carried out on the same day because of serious consequences which may result from failures at the workplaces.
- Constraints on the number of personnel
- Other constraints related to characteristics of the equipment

By checking the entire system, the possibility of overloading associated with switching was carefully calculated for different processes. If overloading was detected, rescheduling of the system was performed. However, by formulating conditions leading to overloading and entering them as constraints, returns for rescheduling could be minimized.

This system was developed in Java for Windows (Figure 6). User interface of the schedulers includes a Gantt chart window indicating equipment shutdown, dialog windows for scheduling data input and modifications, and also a window displaying personnel allocation. Interface for the mediation agent is pre-defined and the scheduling data format is shared in this system.

5.1.1 Parallel assignment strategy experiments

During the series of experiments, the initial values utilized in each of the power station schedulers were set at random, after which an attempt was made to resolve the constraints resulting from shutting down of approximately 10 units among all the equipment shared between the stations. The authors tried to minimize the number of modifications in order to ensure the priority of initial schedules developed at individual stations. During the next step, a comparison was made between the actual measured time required to resolve the

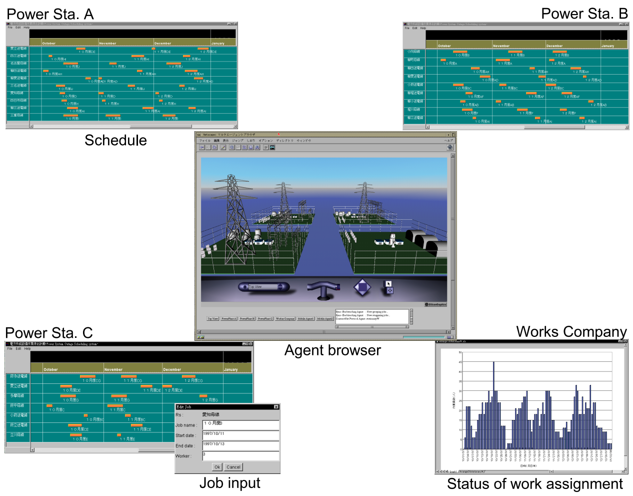

Figure 6: System Snapshots

constraints in this case and that when a single mediation agent performs solution search without cloning (in other words, synchronous backtracking without use of 1)parallel assignment strategy).

When performing search without cloning, the mediation agent proceeds on the basis of the predefined assignment strategy, without reference to another assignment strategy available from the scheduler of destination. In the course of experiments, both minimum step back and minimum step forward shifting were tested in an alternate mode. For example, if the mediation agent is set to perform a minimum step forward shifting strategy when in scheduler A, and the expansion of the partial solution in scheduler B cannot be realized, it returns to scheduler A and performs a minimum step back shifting strategy.

5.1.2 *Multiple priority strategy experiments*

In these experiments, the standard deviation was determined for the number of days modified from the initial configuration, and was compared to the actual spread in the modified number of days between power stations in the case of a single mediation agent performing a solution search without merging (in other words, synchronous backtracking without use of 2)multiple priority strategy). When performing a search without merging, the mediation agent determines the sequence for the parameters having identical priority based on a criterion and assigns them values in accordance with this sequence.

5.2 Results of experiments

5.2.1 *Experimental results of parallel assignment strategy application*

Figure 7 represents a comparison between the search time for 1)parallel assignment strategy using clones, and that for simple synchronous backtracking performed by a single agent. The figure summarizes the results obtained in the course of a series of experiments. The vertical axis is the period of time required to resolve the constraints, and the horizontal axis shows the number of clones. The point on the horizontal axis where the clone number is equal to 0 was determined by using synchronous backtrack without 1)parallel assignment strategy. This corresponds to the value of 100% on the vertical axis, which characterizes the search time.

Experiments confirm that by implementing 1)parallel assignment strategy with cloning function it is possible to reduce the search

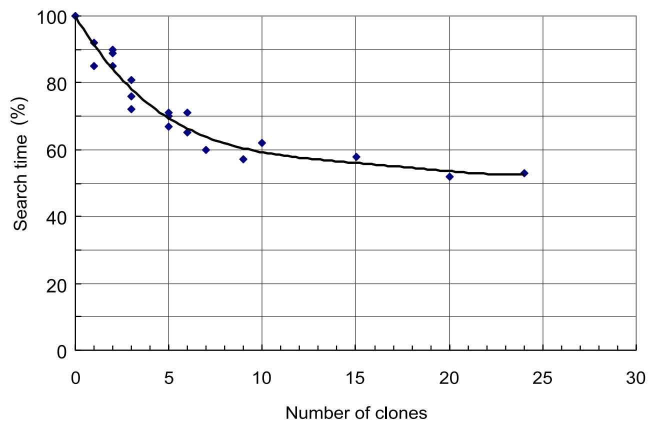

Figure 7: Efficiency of parallel assignment strategy

time by a maximum of 40 to 50% compared to the usual synchronous backtrack method. At this point, the number of clones was about 10 to 20. Since the number of clones corresponds to the volume of calculations required for search, it is very important to limit it in order to minimize the search time. From this point of view, the mediation agents in this system serve as a thread, and if their number is maintained within the limits, no significant consumption of calculation resources will take place. Hence, the results obtained with regard to the actual maintenance scheduling problems can be considered sufficient from the perspective of A)rapid schedule adjustments in view of the fact that they do not cause an increase in calculation volume and that they help reducing the search time nearly by half. However, these results are attributed to their comparative remoteness from the maintenance scheduling issues, as well as to the absence of optimization function. Therefore, when applied to other problems with more rigid constraints, the number of mediation agents (volume of calculations) will increase and so will the search time.

5.2.2 *Experimental results of multiple priority strategy application*

Figure 8 represents a comparison between the spread in the case of 2)multiple priority strategy supported by the merging function and that in the case of a single agent performing priority assignment (the usual synchronous backtracking). The figure presents a plot summarizing the results obtained in the course of a series of experiments. The vertical axis is the spread of the modified number of days (standard deviation) with respect to initial values assigned to each parameter, and the horizontal axis is the number of merges (number of gradual processing operations). A point on the horizontal axis where the number of merges is equal to 0 was determined by using synchronous backtrack without 2)multiple priority strategy. This corresponds to the value of 100% on the vertical axis, which characterizes the spread. The narrower the spread, the better.

The experiments show that implementation of 2)multiple priority assignment strategy supported by the merging function allows spread reduction of the time deviation for each parameter by more than 50% compared to the simple synchronous backtrack method. The number of gradual processing operations is from 2 to 8. Gradual processing requires a greater volume of calculations than the usual variable assignment process, and therefore the smaller their number, the better. Hence, the results obtained with regard to the actual maintenance scheduling problems can be considered suffi-

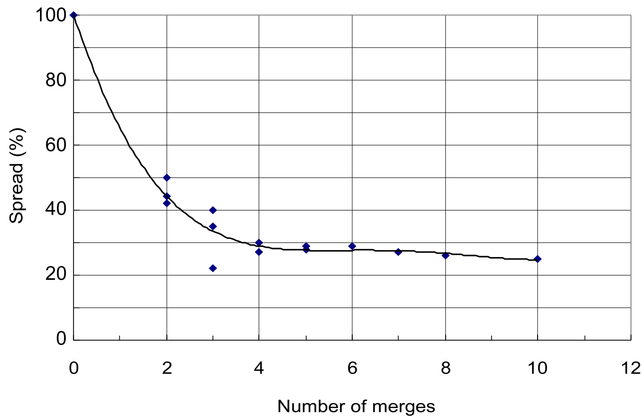

Figure 8: Efficiency of multiple priority

cient from the perspective of impartiality provided that they reduce the spread by more than half despite requiring only a slight increase in calculation volume. This means that B)impartial schedule assignment is possible when there are several power stations with identical job priority. However, these results are also dependent on the nature of the maintenance scheduling task, the presence of a small number of jobs, and the fact that the minor deviations from the schedule could be resolved by small-scale adjustments. Therefore, if gradual adjustment is applied to, for example, job shop scheduling problems, a decrease in solution efficiency may result. Apparently, in that case, the calculation volume will have to be restrained by setting an upper limit for the number of gradual processing operations.

6. DISCUSSION

The method of 1)parallel assignment strategy based on cloning of mediation agents represents partial parallel processing. However, there is the issue of minimizing the calculation volume and increasing the efficiency based on process branching, where backtracking serves as a trigger and the assignment strategy as a unit. The volume of calculations largely depends on such issues as the time during the search at which a set of jobs arises which are most difficult to adjust, or how the trigger is modified during process branching. For example, if there were two jobs at the beginning and at the end of search, which were the hardest to assign, then it would be necessary to correct the initial job at the final stage of search. That is why a single clone must be created at each backtracking, eventually leading to an extremely large number of them (large volume of calculations).

Consequently, it is necessary to develop a process branching method suited to the nature of the task. For example, assuming that a job set which is difficult to adjust can be determined by the domain analysis, it is possible to develop a method of process branching directly to any job set causing adjustment problems, rather than backtracking to the previous job.

It is important to analyze the quantitative relationship between the nature of the task and the number of clones (volume of calculations). In particular, it is necessary to define quantitatively the distance of the search order between several jobs that are difficult to adjust, or the number of hard-to-adjust jobs constituting a job set, or the total number of such sets, as well as the relationship between the number of clones and the actual search time. This will allow more precise establishment of the scope of application of this method.

The main feature of 2)multiple priority strategy based on merging of mediation agents when required is the actual method of gradual adjustment. The specific method of gradual adjustment depends on the type of constraint to which it is applied. In the course of experiments, we examined the effect of using a simple method according to which gradual adjustment was performed by changing the assigned values of parameters, so that they could coincide in their midpoint. The same is applicable to cases in which the number of parameters exceeds 3. Furthermore, as far as the job shop scheduling problem is concerned, if it is necessary to minimize stocks in an intermediate store when changing lines, the adjustment can be made to minimize the difference between the end time on a line and the start time on the next one.

In more general, some readers would question whether mobile agent technology fits to the electric power systems. However, the electric power system has less focus on interoperability than one for the internet, and even the proprietary system has a chance to be adopted if it has an obvious advantage. In fact, our Bee-gent based system for distributed power feeding record collection has been introduced in an information network of the electric system for several years. Moreover, the recent trend of the smart grid will further enhance and reform the information network in the near future. Therefore, it is more of a time to be fit to the electric power systems for the (mobile) agent systems.

7. RELATED WORKS

In the field of the distributed constraint satisfaction, [1] well surveyed this research area. It includes a description of the synchronous backtrack algorithm which is a basis of our heuristics, and some other improvement algorithm and techniques. One example of the improvement which is comparable with our approach is [2], where agents assign variables sequentially, but perform forward checking asynchronously. So that it's also trying to make the synchronous backtracking partially asynchronous as well as our 1)parallel assignment heuristics. However, this paper has focused on problems when applying the heuristics to improve the efficiency of a real schedule, rather than a constraint satisfaction problem in general.

Thus, we will show below some related approaches to the practical distributed scheduling problems. [3] surveyed applications of agents to the scheduling problems in manufacturing. It lists many different approaches which have been proposed in the past, and we are almost impossible to compare with them all in this section. We should note that it also assert importance of non-sequential scheduling based on the agents. One of the leading approaches is multi-scheduler-agents negotiation, such as [4]. After this research, several distributed scheduling systems based on the multi-agents have been proposed. However, this would need replacement of all the scheduling systems in the different power generating companies to agent-based ones in practise. On the other hand, our approach is adding on the mediation agents to the schedulers, and needs to connect to the schedulers with data file and apis. [5] is aimed at the maintenance schedule, but it also using the multi-agents negotiation.

The latest topics are naturally related to the smart grid. [6] and [7] are application of agents to the scheduling problems in the smart grid. [6] uses a bottom-up electronic market mechanism, and [7] determines schedules of various distributed energy resources by an auction mechanism. Both are economical approaches unlike ours, and not handling the maintenance scheduling in specific.

8. CONCLUSIONS

In this paper, an agent-based distributed scheduling system has been proposed against the background of the deregulation of electric utility and the smart grid for the renewable energy. Then, A)rapid schedule adjustments and B)impartial assignment are considered as problems arising in the process of distributed maintenance scheduling. In order to solve these problems, the authors proposed two kinds of heuristics: 1)parallel assignment and 2)multiple priority strategies, then realized them by mobile mediation agent's cloning and merging functions in a scheduling system. We believe this point is our novel contribution to the agent-based distributed scheduling research. Based on experiments oriented to real problems and measurements of system efficiency, it was found that the rapidity of schedule adjustment and the impartiality are improved without massive calculation. The authors plan to apply this system to solve other kinds of scheduling problems associated with the smart grid in the future.

9. REFERENCES

[1] M. Yokoo, K. Hirayama: "Algorithms for Distributed Constraint Satisfaction: A Review", Autonomous Agents and Multi-Agent Systems, Vol.3, Issue 2, pp.185-207, 2000.

[2] A. Meisels, R. Zivan: "Asynchronous Forward-checking for DisCSPs", CONSTRAINTS, Vo.12, No.1, pp.131-150, 2007.

[3] W. Shen, L. Wang, Q. Hao: "Agent-based distributed manufacturing process planning and scheduling: a state-of-the-art survey", IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews, Vol.36, No.4, pp.563-577, 2006.

[4] R. Sikora, M. J. Shaw: "Coordination mechanisms for multiagent manufacturing systems: Applications to integrated manufacturing scheduling", IEEE Transactions on Engineering Management, Vol.44, No.2, pp.175-187, 1997.

[5] M.A. da Rosa, V. Miranda, M. Matos, G. Sheble, A.M.L. da Silva: "Intelligent Agent-based Environment to Coordinate Maintenance Schedule Discussions", 15th International Conference on Intelligent System Applications to Power Systems, 2009.

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

[7] R. Bhuvaneswari, S.K. Srivastava, C.S. Edrington, D.A. Cartes, S. Subramanian: "Intelligent agent based auction by economic generation scheduling for microgrid operation", Innovative Smart Grid Technologies, 2010.

[8] T. Kawamura, T. Hasegawa, A. Ohsuga, S. Honiden: "Bee-gent: Bonding and Encapsulation Enhancement Agent Framework for Development of Distributed Systems", Systems and Computers in Japan, John Wiley & Sons, Inc., Vol.31, No.13, pp.42-56, 2000.

A methodology for agent-based modelling using institutional analysis – applied to consumer lighting

Andreas Ligtvoet
Delft University of Technology
a.ligtvoet@tudelft.nl

Amineh Ghorbani
Delft University of Technology
a.ghorbani@tudelft.nl

Emile Chappin
Delft University of Technology
e.j.l.chappin@tudelft.nl

ABSTRACT

In many (agent-based) simulations the programmers purport to model elements of human society. These models are often based on the researchers' understanding of and focus on specific elements of this society. We suggest that using a framework that is used in social and political science (the Institutional Analysis and Development framework) will greatly improve the design, comparison and validation of artificial societies. We apply this approach dubbed MAIA to a pre-existing model of consumer lighting.

Categories and Subject Descriptors

I.6.5 [**Simulation and Modelling**]: Model development

General Terms

Design, Standardization, Documentation

Keywords

Institutional analysis and development, socio-technical systems, simulation, consumer lighting, MAIA

1. INTRODUCTION

Agent-based simulation is a rapidly developing method both in the social sciences as well as in the engineering sciences, e.g. energy systems [14, 7]. As is the case with developing fields of research, there are many approaches with regard to conceptual design, software implementation and outcome analysis that are hard to compare. Especially when the models move away from general and abstract notions towards more specific and 'predictive' modes (although prediction in complex systems is arguably impossible), more information is required on the assumptions about the agents and the world they inhabit. Partially, this is a matter of clear and consistent reporting and open access to code and data [22]. The other part is building on a shared understanding of concepts, required elements, and framing of the problem at hand, for which we propose using a framework of institutional analysis.

Some colleagues have suggested the use of ontologies to facilitate shared use of code and concepts [25, 22]. This has certainly led to a basis for cooperation among colleagues and a close circle of friendly institutes, all from technical universities. However, as the aspects of energy systems that we are interested in are on the edge of social and technical sciences – emergent market behaviour, policy response, cooperative efforts, societal transitions – we need to go further and embrace frameworks of social and political science that allow us to design, test and communicate our models with other scientists. We have studied the Institutional Analysis and Design framework (IAD) of Elinor Ostrom and colleagues [11], and developed a methodology for agent-based modelling using institutional analysis (MAIA) introduced in [12].

In this paper we test MAIA by applying it to a model of consumer lighting that was developed without institutional analysis in mind. We first describe the MAIA methodology (section 2), then describe the consumer lighting background and the conceptualisation of the model based on MAIA (sections 3 and 4). We discuss the advantages observed in using MAIA for this particular case (section 5) and conclude on the usefulness of the proposed methodology (section 6).

2. INSTITUTIONAL ANALYSIS

Agent-based models are built with the idea that collective behaviour is a result of the aggregate of individual behaviours. In doing so, the social or institutional structures that are the basis for this behaviour are implicitly assumed in the modelling process. For example, two (simulated) agents entangled in a negotiation process will need to be understood in a human (social) context to make sense of the model outcome. Thus, to let a model resting on artificial agents reflect reality, social structures need to be (explicitly) specified and implemented. This can ideally be done independently of the individual agents that constitute the model.

To artificially implement social structures, many researchers have developed methodologies and frameworks which are mainly used for implementing multi-agent systems but not for simulation purposes [5, 9, 26]. Even though some of these are primarily based on social theories [5], their focus is mainly on roles rather than the more complex concepts existing in social environments such as laws, policies, regulations, culture and practices.

2.1 The Institutional Analysis and Development Framework

Institutional analysis aims to find and describe *unifying building blocks* in diverse, regularised social interactions, that can help to explain human behaviour in these situations. It is this unifying aspect that has led to the growth of the institution concept in several disciplines, including economics, philosophy, sociology, and politics [15].

So, what are these institutions? Ostrom defines an institution as "the set of rules actually used by a set of individuals to organise repetitive activities that produce outcomes affecting those individuals and potentially affecting others" [23]. While rules are one of the main elements, the Institutional Analysis and Development (IAD) framework clearly specifies different elements of the system description (see figure 1).

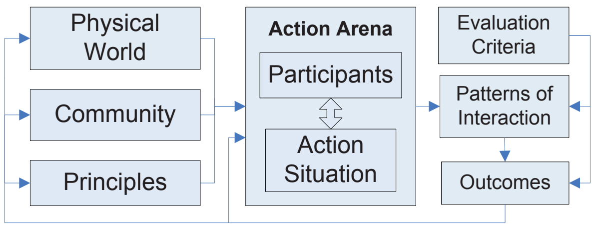

Figure 1: The IAD framework

2.2 The MAIA methodology for agent-based modelling

To apply the IAD framework for modelling agent-based artificial systems, many of the concepts need to be further specified. The MAIA framework has been developed for this purpose. The core concepts of this framework are inspired by the IAD while the details are added based on agent-based software methodologies including OperA [9]. The MAIA methodology covers the whole development cycle of agent-based modelling starting from conceptualisation and design to the coding and implementation of such models. This methodology consists of five structures presented in figure 2.

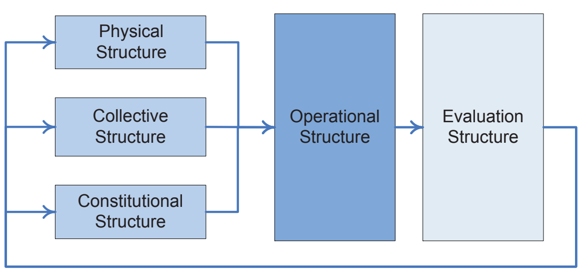

Figure 2: The five structures of MAIA

In this section we briefly describe the meta-model concepts that are used for system decomposition and would later be used for detailed design and implementations. The definitions of this section are used to decompose the consumer lighting case study in the next section. Some explanations of how the conceptualisation can be used in detailed design are given but not covered by the case due to space limitation. The order in which the concepts are defined is the sequence used for the decomposition of a system with MAIA.

2.2.1 *Physical Structure*

The physical structure is made of three distinct concepts namely resource, location and link. Resources have names and properties. They have a type which can either be private (e.g. money) or common pool (e.g water, electricity) [23]. A *location* (e.g. house) also has name, can have properties such as size, coordinate or capacity. A *link* is a physical connection that can be between resources or locations (e.g. electricity grid linking houses). Therefore, beside the name, a link has a begin node and an end node.

2.2.2 *Constitutional Structure*

Generally speaking, the roles the agents take, the groups they form, the dependencies they rely on and the institutions they follow in the social context are defined in the second component called the constitutional structure (see figure 3). The first step is to define the institutions of the system. Institutions can be *formal (i.e. law, regulation etc)* or *informal (i.e. culture)*. Fashion is an example of the informal constitutional environment in the society. Institutions can prohibit, oblige or permit roles to do certain tasks.

A *role* is defined independent of the underlying agent. Roles have predefined *objectives* which can be in contradiction with the agent goals. A role can be *external* to the system which means that agents cannot take that role, or *internal*. The *boundary rule* of a role defines when an agent can be considered as having the role. The *rights* of a role however define what the agent can do when taking a particular role.

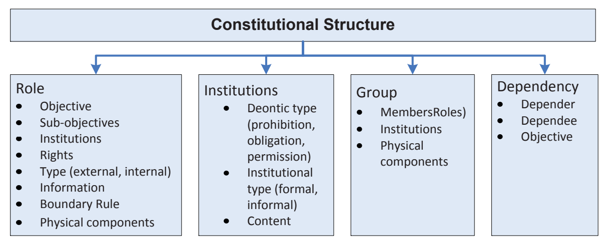

Figure 3: Meta-Model Concept: Constitutional Structure

To reach their objectives, roles depends on other roles which is represented in a *dependency graph* (figure 7). The nodes of the graph represent the roles and the edges are the objectives. A link is drawn between two nodes if a role depends on some other role to achieve his objective. Sometimes, an institution may belong to more than one role. In such case, a *group* is defined. Also, when a resource belongs to more than one role, it is given to a group.

2.2.3 *Collective Structure*

This structure specifies the *agent*, the *social network* and the *domain ontology*. Agents can take different roles. Each agent has some *preferences* which may also affect how he would follow the institutions of the role he is taking. The preferences of the agent together with his *properties* such as age or eduction define the personality of an agent which affects the decision making process. The *information processing capability* of an agent is the amount of information one may have access to. Similar to roles and groups, agents can also have *resources*. The *number* of agents taking a specific role is also defined in the agent table.The link between agents and roles is one of the most important aspects of the MAIA framework which will later be discussed.

The *social network* is a graph where the sender and receiver are agent-role pairs and the link between two nodes gives a general description of the interaction between two pairs. Finally, a domain ontology defines the common terms and concepts of the system. Therefore, every concept defined in system conceptualisation are put in the ontology which can later be used for the implementation of the model.

2.2.4 *Operational Structure*

The general focus of the operational structure is on the continuous activities of the system such as decision making of the participating agents and the functions that link decisions to outcomes.

This structure is divided into three main components, the action situation, the role enactment agreement and the action sequence. In the action situation, the actual processes of the system take place. Roles take *actions* which have some *costs* and may result in *benefits*. Each action situation may have its own institutions and resources. In a *role enactment agreement table*, the link between agents and roles is defined. Each agent can take one or more roles and vice versa. The specifications of the role and the agent are all present in this table and each possible combination occupies one row of the table. In the *behaviour identification* column, the decision making process of the agent is defined. This can be done by giving weights to different components or reading external data. The agents can take actions that are forbidden as a part of the role institutions. Since this is done explicitly in the model it is thus traceable.

The *action sequence* which is a component of the detailed design phase, defines the actual interaction between agent-roles. The interaction is the message exchanges and the action is an atomic process taken by only one agent. These set of diagrams are drawn based on the action situations defined in this part and the social network diagram in the collective structure.

2.2.5 Evaluation Structure

Model evaluation is a step that is often addressed after the software has been implemented and not taken into consideration at the conceptualisation or even design phase. The evaluation indicators of an agent-based model can be divided into two different categories:

1. *Reality closeness parameters*: When modelling a real world system, we first have to make sure that the artificial society is functioning as close to the real system as possible. Choosing the relevant parameters to evaluate the closeness of the model to reality is a function of this category. These parameters are suggested by the domain experts. In the consumer lighting example, light bulb consumption is chosen as one of the closeness parameters. It is important to realise that the value of these parameters are outcomes of the system and will be calculated after (or during) the model run. The values are not given during the implementation of the model.

2. *Problem domain parameters*: The purpose of the modelling exercise is to answer a set of questions defined by the problem owner. To better understand these questions and make efforts to answer them, certain parameters have to be assigned to each question. These parameter values are again the outcome of the system and only realised during runs or even after.

   The purpose of a scope matrix (see e.g. table 2) is to document the relation between the operational events and the global outcomes. Through this matrix one would have a general idea of which local behaviours may lead to certain emergent outcomes.

The feedback received from these two categories of parameters can be reflected in the physical, collective and constitutional structures of the MAIA framework to further affect agent actions and interactions at the operational structure.

3. CONSUMER LIGHTING MODEL

3.1 Background

Lighting is essential for modern living – it enables humans to do many things otherwise impossible. Edison's first carbon filament glow bulb had a lifetime of 45 hours and an efficiency of 2 lm/W.[1] Many gradual improvements in electric lighting technologies increased the lifetime of the bulbs and the electric efficiency [10]. By 1912, the glow bulb's efficiency had improved to reach a light output of 12 lm/W of electricity. Technological progress in incandescent bulbs halted at that point. Presently, almost 100 years later, the incandescent lamps are hardly more efficient: even now, over 98% of the electricity used is converted into heat and not into light.

More energy efficient alternatives have been developed, for example the compact fluorescent lamp (CFL) [2]. The CFL was first introduced by Philips in 1980, and offered four times energy savings and a much longer lifetime, with some disadvantages (size, weight). Subsequently, the CFL was much improved in the decades afterwards, and was known as the 'saving lamp'. The CFL enables a dramatic increase in the energy efficiency of lighting while, partly being a screw-in/plug-in replacement, it retains an amount of compatibility with existing luminaires. CFL's offer clear benefits for many applications, and many governments tried to stimulate their use (e.g. [20, 18]), but these stimulus policies have only seen limited success. Presently, CFL saving bulbs are present only in 55% of European households [4]. Another exciting development is solid-state lighting: the Light-Emitting Diode (LED). Proponents consider the LED as the ultimate lamp of the future, because it is very suitable to a wide range of applications and because it continues to achieve significant gains in electric efficiency [8, 24, 16, 2].

Consumers have adopted CFL and LED technology only partially because of a number of obstacles [19]. CFL and modern LED saving lamps are characterized by high up-front cost for consumers and poor light quality, which serve as a barrier for adoption. Consumers implicitly use high discount rates when purchasing energy efficient durable goods [13, 17]. Halogen lamps proved more attractive because they fitted in popular designs and do not have the disadvantages that CFLs have (cost, unfavourable colour, size, no dimming option, . . .).

In consumer lighting, changes are forthcoming. The European Union's phase-out of incandescent lighting is a clear strategy that will change the sector, it involves regulation designed to remove from stores the cheapest forms of inefficient household lighting [6]. Although implied, it is uncertain whether the lighting sector will become efficient overnight; consumers may switch to forms of inefficient lighting that are exempt from the phase-out; or consumer behaviour will change. The precise dynamics induced by the phase-out are unknown.

3.2 Model implementation

To get a better insight in these dynamics, Afman et al. developed an agent-based simulation using Java with elements of RePast [1]. This model encompasses agents –

[1] Light output is measured in lumen (lm). An ordinary incandescent 75 W bulb (which is now banned in the EU) emits more or less 900 lumen at 12 lm/W. With a theoretical maximum of 683 lm/W, the bulb is <2% efficient [2].

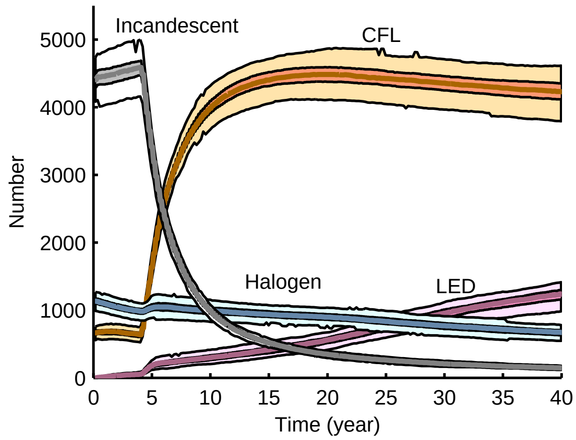

Figure 4: Model run with a ban on bulbs

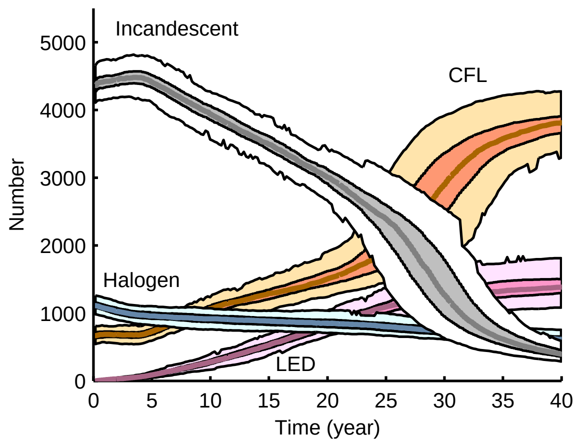

Figure 5: Model run with a tax on bulbs

consumers – that buy lamps, based on the available luminaires in their houses, their personal preferences and the preferences of their acquaintances. By introducing policies that affect consumers' behaviour (banning light bulbs, taxing light bulbs, or subsidising energy efficient alternatives), the system outcome is affected over a simulated period of 40 years.

3.3 Simulation results

The myriad of individual decisions of agents to purchase lamps drives the system as a whole: system level adoption of the technology types, average electricity consumption levels and average expenditure on lamp purchases are important indicators of a transition towards more efficient energy use.

The adoption of the four different lamp categories is plotted over simulated time (40 years), see figures 4, 5, and 6 per policy initiative. The results show that both the ban on bulbs and the taxation scheme on incandescent bulbs are effective to phase out the incandescent bulb in the long run. The ban on bulbs has a faster and more prominent effect. With the LED subsidy the incandescent bulbs remain dominant.

The simulation model confirms that, in the long run, the 'ban on bulbs' is the most *effective* way of achieving a lower electricity usage for lighting. A tax on bulbs of euro 2 is also effective. In contrast, a subsidy on LEDs is not effective. An important disadvantage of a ban is the burden on

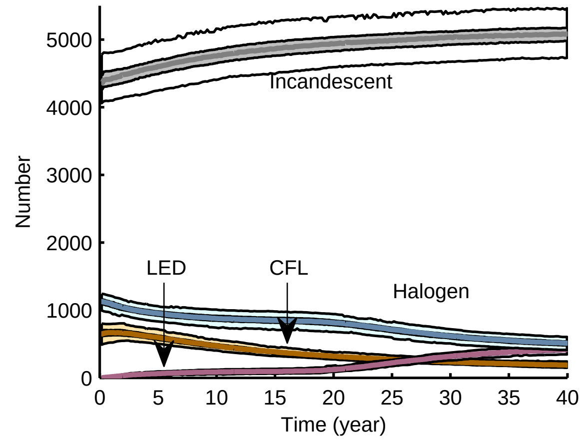

Figure 6: Model run with a subsidy for LEDs

consumers: expenditures spike during the start of the ban. This might be considered unacceptable. In contrast to the ban, a tax could be made income neutral. Whether it is a ban or a tax, it is crucial to attack all unwanted products. In this case halogen is *not* more electricity extensive then incandescents, but is not banned or taxed. If the penetration of halogen proceeds, it hampers the transition to lower electricity consumption in the sector.

4. RECONCEPTUALISATION OF THE CONSUMER LIGHTING MODEL

In this section we revisit the consumer lighting model using the MAIA framework. The description of the model is made in an effort to be compact yet complete. We will only use the components that are necessary for the consumer lighting case study. Due to space limitation, we will only illustrate some of the tables and diagrams used for this particular case.

4.1 Physical Structure

The physical structure (resources, locations and links) in the consumer lighting model is fairly small. There are two *resources*, namely luminaires and lamps. Luminaires have the properties: socket type (E27, GU10, E14), wattage and status (operational, failed). Lamps have more properties: the lamp technology (empirical data on 70 lamp technologies were collected in a variety of lighting stores), the expected lifetime, uncertainty of the lifetime, light output, electricity consumption, colour rendering index, colour temperature, voltage, shape, socket, and purchase price. There is also the ageing process that is a part of the lamp resource. Both lamp and luminaires are private resources.

Usage for each lamp is related to the location in the house: some lamps are used more often than others. Therefore, we define luminaireLocation as a *location* in this structure. We define a distribution function for the place of lamp in the house based on survey data on the number and usage of lamps in consumers' homes [1].

4.2 Constitutional Structure

The four different components of the constitutional structure are covered in this section.

4.2.1 Institutions

There are four institutions implemented in the model. First, word of mouth is considered as a culture which is an informal institution. People give their opinion about lamps to other people. When a consumer communicates its opinion, the opinions of a neighbour are averaged between its old value and the other consumer's opinion.

There are three possible policy interventions beside the base case (no governmental policy) in this model i.e. (1) a ban on bulbs, (2) an incandescent taxation scheme and (3) a subsidy scheme on LED lamps. Each of these are considered a formal institution that the agents are obliged to follow.

The first policy entails a complete ban on the standard incandescent light bulb, implemented between years 2 and 5 (relative to the start of the simulation). This policy is comparable to the EU ban on household light bulbs: first the incandescent bulbs with the highest wattages are removed from the stores, after which progressively the lower wattages are removed.

The second policy scenario introduces a taxation on the sale of incandescent light bulbs that increases progressively during the first five years of the simulation to a maximum of euro 2.00 per lamp (which is relatively large compared to a purchase price of euro 0.35 – euro 1.50).

The third policy is a subsidy on the purchase price of LED lamps. The subsidy is a discount of 33% of the purchase price. After five years, relative to the start of the simulation, the subsidy is slowly removed until year ten, where it is zero.

4.2.2 Roles

The *Consumer* is considered as a role in this case study. The objective of the consumer role is to have pleasant light in the house. The agent taking the consumer role must follow the word of mouth, tax and subsidy institutions. This role is internal, which means that agents in the model can take this role. The boundary rule is for the agent taking this role to own luminaires. In other words, when checking to see if an agent has the role of a consumer, we need to see if he owns luminaires. The information available to this role is the information about lamps and luminaires. The physical components belonging to this role are luminaires and lamps. The other roles of the system are presented in table 1. As can be seen in the table, except the consumer the rest of the roles are considered external. Agents cannot take external roles and thus there is no decision making in the behaviour of these external entities. Therefore, there is no need to define conditions (boundary rules) for an agent to take any of these roles.

4.2.3 Other Components

The group component was not defined for this case since there were no institutions that had meaning for more than one role, just as there were no physical structures for a group of roles. To give more details to this model, one option is to have co-inhabitants of an apartment building as a group where there are lamps in hallways that are shared by everyone in the group and the type of lamp to be chosen is a group decision. The last component of the constitutional structure is the group dependency graph which is shown in figure 7.

4.3 Collective Structure

The collective structure in the consumer lighting deals with the agents who are taking part in the society and the social network. All the defined components throughout the framework are put in a domain ontology for later use.

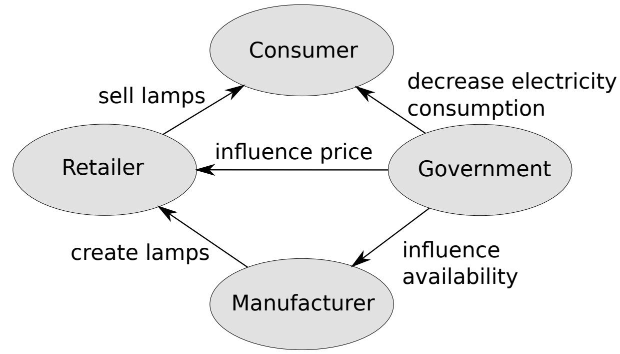

Figure 7: Dependency Graph for Consumer Lighting

4.3.1 Agents

There is only one type of agent, the consumerAgent, in the consumer lighting case study since the other entities in the system such as the retailer or manufacturer are considered external. The consumerAgent has properties including the number of lamps he owns, usage of lamps, total number of luminaires, socket types, the specific lamps installed initially and his opinion about the lamps. ConsumerAgents have preferences such as light colour, colour rendering and light output. We assume consumers as heterogeneous in their initial portfolio of lamps (total number of luminaires, socket types, and the specific lamps installed initially) and in their preferences for light colour, colour rendering and light output. Consumers start with neutral opinions (to become negative or positive during the simulation). Consumers form opinions about individual lamp types, technologies and brands. Since the luminaires and lamps belong to every agent taking the consumer role, there is no need to define these resources here. However, each agent has luminaireLocation which is different for every agent and thus given to the agent and not the role.

4.3.2 Social Network

There are 250 consumerAgents, arranged in a scale-free network, which is considered sufficient for modelling consumer networks [3, 21]. In the network, agents are connected to at least 15 others. The main interaction taking place in this model is the exchange of opinion between two consumers. The *buy* interaction takes place between the consumer and the retailers and there is also an *order* interaction between the retailer and the manufacturer. These interactions are the general interactions in the model which are represented in a graph similar to the dependency graph shown previously. These interactions will be further specified in sequence diagrams in the detailed design phase.

4.4 Operational Structure

4.4.1 Action Situation

The system to be modelled is divided into distinct action situations. Each action situation may consist of several actions and interactions. However, sometimes it is necessary to define a situation for an event that will be happening somewhere during the runs. In other words, the trigger for anything to happen through the model run should be defined

Table 1: The Role Table for Consumer Lighting

Name	Objective/Sub-objective	Institution	Type	Boundary Rule	Information	Physical Component
Consumer	Light in his house	fashion, tax, subsidy, ban	internal	agent owning luminaire	lamp and luminaire data	luminaire, lamp
Retailer	Sell lamps	ban	external	none	lamp and luminaire data	luminaire, lamp
Manufacturer	Produce lamps	ban	external	none	lamp and luminaire data	luminaire, lamp
Government	Decrease electricity consumption	none	external	none	lamp and luminaire data	none

here.

The first action situation defined for consumer lighting is the lamp replacement situation which involves the consumer role. The actions happening here are: lamp defect, check socket type, update opinion (based on the history of this defect) and replace lamp. In this version of the model, households do not change their luminaires. Therefore, the options to change lamps are limited to socket-compatible lamps. The information that becomes available in this situation is the socket type and the status of the lamp which could be operational or failed.

Similar to this action situation, there is the buy lamp action situation where the actual choosing and buying of the lamp takes place. Exchange of opinion and lamp supply are yet other action situations. When a consumer communicates its opinions in the exchange of opinion situation, the opinions of a neighbour are averaged between its old value and the other consumer's opinions.

One situation that is somewhat distinct compared to the others is the technological advancement action situation. Even though there are no roles involved in this situation, the consideration of it is especially important for the implementation of change in prices according to technology later on in the model. The action in this situation is 'change price'.

Although the prices of all lamp models differ, the lamp technology type determines the decline in price over time. Newer technologies – LED and CFL – are modelled to decline faster than proven technologies (halogen and incandescent). Due to technological progress, the efficiency of each technology increases over time, where the efficiency of more 'modern' technologies CFL and LED are assumed to increase faster than the old technologies.

4.4.2 Role Enactment Agreement

The only internal role in this model is the consumer. Therefore, the consumerAgent requires decision making and character specifications. The behaviour patterns of the consumerAgent are defined according to the properties of the consumerAgent and the specifications of the role. The consumer's lamps purchase decision is modelled using multi-criteria analysis incorporating criteria related to aspects of the lamps (purchase price, efficiency, and lifetime), preferences for subjective lamp qualities (colour, colour rendering index (CRI), light output), and opinions (perceptions) on the lamp's aspects (lamp model, brand, and technology type). A final criterion relates to what other consumers do (word of mouth).

A number of important behavioural assumptions underlie the criteria weight factors that determine the relative importance of the normalized scores. As the purchase price needs to be the most important criterion [19], it is assigned a high weight factor of 4. Then lamp efficiency, colour rendering, light colour, the household's opinion of lamp technology type and normative adaptation (word of mouth: imitating neighbours) are assigned a weight factor of 2: important, but not as strong as the purchase price. Last, the lamp's light output, lifetime, and the consumer's opinions on brand and lamp model are even less important, these get a weight factor of 1. Between household agents, weight factors differ by +/− 50% to make them heterogeneous.

A consumer's opinions (a [-1..1] scale) are changed as a result of own experiences with bought lamps and through information it receives from neighbours in the social network (word of mouth). Parameters and increment values used for autonomous opinion change are:

- if positive experience: + 0.1
- if negative experience: − 0.3
- if positive experience, contrary to existing opinion: + 0.2

4.5 Evaluation Structure

Three main evaluative parameters were defined regarding the main questions: system level adoption of the technology types, average electricity consumption level and average money expenditure on lamp purchases. The scope matrix relates the main parameters to the actions situation while helping us define new parameters which may also be as crucial in analysing the system. This parameter identification happens while trying to further specify the indirect relationship between the main parameters and the action situation. The scope matrix is represented in table 2. There is a direct relationship between system level adoption of technology types and the buy lamp action situation as this is where the actual buying of different technologies happens. There is an indirect relationship with all the other action situations.

The next step is to define a parameter that would have a direct relationship (d) with the opinion exchange action situation while explaining the first parameter in more detail. The general public opinion about different technology types seems to be the parameter we are looking for. This process can continue for all the other indirect relationships (i) until we have sufficient number of indicators specified. The letter (n) indicates that there is no relationship present between the two entities.

The reality closeness parameter matrix is a similar matrix where the indicators do not necessarily answer domain problem questions. The number of lamp defects or the price of lamps are two such parameters.

Table 2: Scope Matrix linking actions and outcomes (PI = policy initiative; d = direct i = indirect n = no relationship)

Action Situation / Outcome Parameters	buy lamp	word of mouth	lamp defect	ban	tax	subsidy
adoption of lamp type	d	i	i	d	i	i
perceptions on lamp types	n	d	i	i	i	i
electricity consumption level	d	i	d	i	i	i
money expenditure on purchases	d	i	i	i	i	i

5. DISCUSSION

The proposed method attempts to combine insights from computer science, social science (institutional economics), and multidisciplinary approaches focused on socio-technical systems. The reason we believe that this approach makes sense is because (1) a common framework can ensure comparable approaches and facilitate communication, (2) introducing the IAD will open up communication with social scientists who have focused on institutions for several decades and (3) tying it all in the MAIA approach brings computer and social sciences together. In other words: social science concepts are translated into formal concepts that can be used to program models.

5.1 Common frameworks

It is a basic value of science that our work should be clearly described and repeatable for falsification of the findings. In practise, researchers forget to mention on what platform the models are developed or to verify their outcomes [14]. Applying a common framework does not remedy such omissions. It does, however, facilitate comparison of different models using the same framework. For those knowledgeable of the framework, it helps to qualify the outcomes and to place them in a certain school of thought. It helps the modeller to structure the model and to give attention to the different elements of the framework. It also makes it easier to detect omissions or simplifications, for example regarding the decision making of agents. Except for very abstract models or toy models adhering to a modelling framework makes good sense. The question is: what framework?

5.2 Using the IAD framework

As indicated above, many researchers have developed frameworks that are used for implementing agent-based software systems [5, 9, 26]. The focus of these approaches is on individual roles (the 'micro' level) rather than the more aggregate phenomena such as laws, policies, regulations and culture (the 'macro' level). These frameworks are more thoroughly grounded in the computer sciences, and as such are less accessible for social and political scientists or non-scientific experts.

The use of the IAD framework comes with the important advantage that the concepts employed have been used, tested and validated by a wide array of scientists. This ensures that the building blocks are comprehensive and relieves the researcher somewhat from explaining the choices that she made. Furthermore, explicit attention is given to the evaluation structure, which again helps in structuring the research beforehand. The nature of the IAD is also closer to agent-based modelling than other tools such as system dynamics. This makes thinking in terms of agents easier when applying this framework.

5.3 Using MAIA

The last step is linking IAD back to computer science. With MAIA we formalise the concepts in the IAD and the links between the concepts. We believe it facilitates the writing of code and one of the authors (Ghorbani) is working on providing tools for this purpose.

Looking back at the original consumer lighting model, we cannot prove that using MAIA would have improved the model – this is also a function of the experience of the modeller. What we do find is that it can help in the conceptualisation phase and function as a guideline for including necessary elements of socio-technical systems. In the case of the consumer lighting, it has helped in checking if all assumptions made by the modelling team were well thought through and consistent. In particular, the distinction between agents and the roles they take up could be an interesting extension of the model. In addition, some entities in the model such as the government were considered as the environment in the implemented model while with the MAIA methodology, such elements are considered external roles. This is more logical as it gives more flexibility to the definition of the government as an active component. Instead of hard coding policy measures in the program, linking the (formally defined) institutions or policies to governmental enactment (action situation) would facilitate the implementation of different policies in the model and the analysis of their influence on other entities, including agents.

The use of role tables such as table 1 points at the fact that there is only one agent that is responsive to changes in the system. This then leads to ideas of adding of various retailerAgents and manufacturerAgents could have added insight into the interaction between consumers and retailers' strategies (e.g. marketing).

The MAIA framework also contains elements that are not necessary for the purpose of the model. For instance, the links between physical structures (physical connections of houses) are suggested by the MAIA framework, but were not modelled. Despite the fact that such links are unnecessary, the framework explicates that we could have modelled them. The framework is only intended to help in selecting the right elements of the model without forcing us to take any.

5.4 Limitations of using MAIA

Like any other software methodology, the MAIA framework has a learning curve which may be time consuming. After three case studies we have developed through MAIA, we realised that understanding the distinction between agents and roles may not be straight forward for users. However, once this is understood, it is a practical and useful concept of the framework. One other issue that can be considered as a drawback for MAIA is that there is no specific architecture for making agents. All MAIA gives to the user is the set of properties that should be considered when defining an agent, these include the properties of an agent, its preferences and its decision making process, among others. However, MAIA considers these information in a set of tables without structuring them into an architecture such as

BDI, reactive or psychological framework. This can also be considered a positive point as it gives more flexibility to the user.

6. CONCLUSION

Although there are many different ways of describing the social (or the technical) context of simulations, communication is served if a common framework is used. For the modeller, some of the assumptions she makes on the structure of society may seem trivial, but they may not be in an international context and thus will need to be explained, preferably in a standard way. Furthermore, the framework encourages the researcher to be explicit about the elements she includes and the ones she excludes. This may be too much effort for abstract models that focus on one or two elements, but for simulations of complex socio-technical systems we deem it invaluable to be as precise as possible about the building blocks of the model. The MAIA method, built on cross-disciplinary insights, may help in this respect.

We have already developed a few case studies with regard to energy and industrial networks to test the applicability of MAIA. We invite other researchers in the field of agent-based simulation to further test and apply the method with us.

7. ACKNOWLEDGMENTS

This research was made possible by the Next Generation Infrastructures Foundation.

8. REFERENCES

[1] M. R. Afman, E. J. L. Chappin, W. Jager, and G. P. J. Dijkema. Agent-based model of transitions in consumer lighting. In *3rd World Congress on Social Simulation: Scientific Advances in Understanding Societal Processes and Dynamics*, 2010.

[2] I. L. Azevedo, G. Morgan, and F. Morgan. The transition to solid-state lighting. *Proceedings of the IEEE*, 97(3):481–510, 2009.

[3] A. L. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.

[4] P. Bertoldi and B. Atanasiu. Electricity consumption and efficiency trends in the enlarged european union – status report 2006. 2007.

[5] P. Bresciani, A. Perini, P. Giorgini, F. Giunchiglia, and J. Mylopoulos. Tropos: An agent-oriented software development methodology. *Autonomous Agents and Multi-Agent Systems*, 8(4):203–236, 2004.

[6] CEC. Commission regulation (ec) no 244/2009 of 18 march 2009 implementing directive 2005/32/ec of the european parliament and of the council with regard to ecodesign requirements for non-directional household lamps. *Official Journal of the European Union*, L076:3–16, 2009.

[7] E. J. L. Chappin. *Simulations of Energy Transitions*. PhD thesis, Delft University of Technology, 2011. ISBN: 978-90-79787-30-2.

[8] S. Curtis. Efficiency gains boost high-power LED performance. *Compound Semiconductor*, 11(11):27–30, 2005.

[9] V. Dignum. *A model for organizational interaction: based on agents, founded in logic*. Siks dissertation series, Utrecht University, 2004.

[10] M. F. Gendre. Two centuries of electric light source innovations, 2003.

[11] A. Ghorbani, A. Ligtvoet, I. Nikolić, and G. Dijkema. Using institutional frameworks to conceptualize agent-based models of socio-technical systems. In S. Stepney, P. H. Welch, P. S. Andrews, and A. T. Sampson, editors, *Proceedings of the 2010 Workshop on Complex Systems Modelling and Simulation*, Odense, Denmark, August 2010.

[12] V. D. G. Ghorbani, A. Dignum. An analysis and design framework for agent-based social simulation. *to be published in proceeding of "International Workshop on Agent-based Modeling for Policy Engineering (AMPLE11)"*, 2011.

[13] J. A. Hausman. Individual discount rates and the purchase and utilization of energy-using durables. *The Bell Journal of Economics*, 10(1):33–54, 1979.

[14] B. Heath, R. Hill, and F. Ciarallo. A survey of agent-based modeling practices (january 1998 to july 2008). *Journal of Artificial Societies and Social Simulation*, 12:9, 2009.

[15] G. Hodgson and J. Calatrava. What are institutions. *Journal of Economic Issues*, 40:1, 2006.

[16] J. N. Holonyak. From transistors to lasers and light-emitting diodes. *Materials Research Society Bulletin*, 30:509–515, 2005.

[17] P. Kooreman. Individual discounting, energy conservation, and household demand for lighting. *Resource and Energy Economics*, 18(1):103–114, 1996.

[18] E. Martinot and N. Borg. Energy-efficient lighting programs. experience and lessons from eight countries. *Energy Policy*, 26(14):1071—1081, 1998.

[19] P. Menanteau and H. Lefebvre. Competing technologies and the diffusion of innovations: The emergence of energy-efficient lamps in the residential sector. *Research Policy*, 29(3):375–389, 2000.

[20] E. Mills. Efficient lighting programs in europe: Cost effectiveness, consumer response, and market dynamics. *Energy*, 18(2):IN1,131–144, 1993.

[21] M. Nekovee, Y. Moreno, G. Bianconi, and M. Marsili. Theory of rumour spreading in complex social networks. *Physica A: Statistical Mechanics and its Applications*, 374(1):457–470, 2007.

[22] I. Nikolić. *Co-Evolutionary Method For Modelling Large Scale Socio-Technical Systems Evolution*. PhD thesis, Delft University of Technology, 2009.

[23] E. Ostrom, R. Gardner, and J. Walker. *Rules, games, and common-pool resources*. Univ of Michigan Pr, 1994.

[24] U.S. Department of Energy. Multi-year program plan fy'09-fy'15 – solid-state lighting research and development. Technical report, U.S. Department of Energy, 2009.

[25] K. van Dam. *Capturing socio-technical systems with agent-based modelling*. PhD thesis, Delft University of Technology, 2009.

[26] M. Wooldridge, N. Jennings, and D. Kinny. The gaia methodology for agent-oriented analysis and design. *Autonomous Agents and Multi-Agent Systems*, 3(3):285–312, 2000.

A Machine Learning Approach for Fault Detection in Multi-variable Systems

Ying Guo
CSIRO ICT Centre
Marsfield, NSW 2122
Australia
+ 61 2 9372 4708
Ying.Guo@csiro.au

Josh Wall
Division of Energy
Technology, CSIRO
Australia
+ 61 2 4960 6108
Josh.Wall@csiro.au

Jiaming Li
CSIRO ICT Centre
Marsfield, NSW 2122
Australia
+ 61 2 9372 4707
Jiaming.Li@csiro.au

Sam West
Division of Energy
Technology, CSIRO
Australia
+ 61 2 4960 6129
Sam.West@csiro.au

ABSTRACT

The present paper relates to systems and methods for detecting faults in multi-variable systems/agents. The approach is particularly suitable for detecting faults in heating, ventilation and air-conditioning (HVAC) systems and will be described in relation to that exemplary but non-limiting embodiment.

The fault detection approach is based on statistical machine learning technology. This is achieved by learning the consistent nature of normal HVAC operation, and then using the statistical relationships between groups of measurements to identify anomalous deviations from the norm and identify faults in all subsystems for which sensor information is available, regardless of the specifics of the installation.

The approach models the dynamical sub-systems (agents) and sequence data in HVAC system. These models (agents) are constructed via a learning process from some training data of normal running HVAC systems. The trained models (agents) can then be used for automatic fault detection. Our algorithm can capture the fact that time flows forward by using directed graphical models (agents), it is adaptive to environment changes and is reliable in HVAC systems. The approach has been tested based on real data from commercial HVAC systems. We can successfully detect a number of typical faults. The experimental results are all very positive.

Using the adaptive learning approach requiring only incumbent sensors' data, this technique offers several advantages over rule-based techniques. In particular, the amount of customization, configuration and expert knowledge required to implement such techniques on different HVAC systems is greatly reduced. The results obtained by using the methods prove that this algorithm has the property of robustness and adaptability. Current intention can achieve better fault detection results when the building situation and/or environment properties changes, while traditional approaches are normally not adaptable to these changes automatically.

Categories and Subject Descriptors

J.2 [**Physical Sciences and Engineering**]: Engineering.

General Terms

Algorithms, Reliability, Experimentation, Detection, Fault.

Keywords

Machine learning, Dynamic Bayesian Network, multi-variable system, fault detection, HVAC system.

1. INTRODUCTION

A wide variety of systems can be described as multi-variable systems/agents. For example systems such as building **heating, ventilation and air-conditioning** (HVAC) systems are multi-variable systems/agents. Multi-variable systems/agents such as these commonly fail to satisfy performance expectations envisioned at design time. Such failings can be due to a number of factors, such as problems caused by improper installation, inadequate maintenance, or equipment failure. These problems, or "faults," can include mechanical failures, control problems, design errors, and inappropriate operator intervention [4]. Due to the complexity of such multi-variable systems/agents, it is relatively common for faults to go unnoticed for extended periods of time, and only identified when they become significant enough to cause complete equipment failure, excessive power consumption, and/or sufficient performance degradation to trigger user complaints.

By way of one example, faults in HVAC systems may be caused by a variety of factors. Mechanical faults can include stuck, broken, or leaking valves, dampers, or actuators, fouled heat exchangers, or other damaged/inoperative components. Control problems ran relate to failed or drifting sensors, poor feedback loop tuning or incorrect sequencing logic [2].

A variety of **fault detection and diagnostic** (FDD) techniques for multi-variable systems/agents are known, and their use provides for a number of benefits [3][21]. By detecting and acting on faults in multi-variable systems/agents significant energy savings can be realized. Additionally, if minor faults are detected before becoming major problems, the useful service life of equipment can be extended, maintenance costs can be reduced, and repairs can be scheduled when convenient (avoiding downtime and overtime work) [5][6].

Further, and again using a HVAC system by way of example, detecting faults allows for better control of temperature, humidity, and ventilation of occupied spaces. This, in turn, can improve employee productivity, guest/customer comfort, and/or product quality control.

Many current fault detection techniques for multi-variable systems/agents are rule-based [7][8][9][15]. The fault detection system integrates and interprets incoming data in accordance with a pre-determined set of rules, produces a risk profile, and autonomously initiates a response to a breach of these rules. Rule-based systems are, however, limited insofar as they are very specifically derived for/tailored to a particular system and are very difficult to update, change, or adapt to a different system. Additionally, rule-based systems typically fail miserably if conditions beyond the boundaries of the knowledge incorporated in them are encountered.

Although less common, another class of fault detection techniques used for multi-variable systems/agents are model-based systems [10][11][12]. Model-based systems use analytical mathematical models to identify faults. As with rule-based systems, however, model-based systems have a number of limitations. For example, model-based systems are generally complex and computationally intensive, and a large amount of skilled work is required to develop a model for a particular system. Also, in order to create a usable model many inputs are required to describe the system being modeled and the values of some of the required inputs may not be readily available.

In addition to the above limitations, most multi-variable systems/agents are installed in different buildings/environments. This generally means that rules or analytical models developed for a particular system cannot be easily applied to an alternative system. As such, the difficult process of determining and setting rules or generating analytical mathematical models must be tailored to each individual building/environment. In addition, the task of setting the thresholds used by such systems to raise alarms is involved, and prone to producing false alarms. Also, building conditions such as structure of the internal architecture design and even external factors (such as shading and the growth of plant life) often change after the system installation/initialization of a fault detection system, which can require rules/models that were originally appropriate to be revisited and updated.

It would be desirable to provide a fault detection system and/or method which overcomes or ameliorates one or more of the limitations of existing fault detection systems/methods. In the alternative, it would be desirable to provide a useful alternative to existing fault detection systems and methods.

Broadly speaking, the present approach relates to systems (agents) and methods for detecting faults in a multi-variable system /agent. In order to illustrate the features of the machine learning approach, examples of the proposed fault detection system and method will be described in relation to HVAC systems. It is to be understood, however, that systems and processes according to the approach may be used to detect faults in other multi-variable systems/agents.

Two most recent papers by Najafi et. al. [13][14] present an approach for fault diagnostics in HVAC systems based on Bayesian Networks . Their papers apply a Bayesian updating approach that is systematic in managing and accounting for most forms of model and data errors. The proposed method uses both knowledge of first principle modelling and empirical results to analyse the system performance within the boundaries defined by practical constraints. The focus of their application is to cope with sensor failure which does not necessary mean that the system has faults. It is mentioned in the paper that "The goal of this paper is to develop a systematic solution to improve the observability of the system when faced with limited data." The goal of our approach is to detect the real faults in the system with a model with strong adaptability. Our approach has good generalisation ability. That is, it can be easily applied to different HVAC systems without long training time. The adaptability and generalibility is not mentioned in their paper.

Using the presented machine learning approach requiring only incumbent sensors' data, this technique offers several advantages over existing techniques. In particular, the amount of customization, configuration and expert knowledge required to implement such techniques on different HVAC systems is greatly reduced. The results obtained by using the methods prove that this algorithm has the property of robustness and adaptability. Current intention can achieve better fault detection results when the building situation and/or environment properties changes, while traditional approaches are normally not adaptable to these changes automatically.

2. HVAC SYSTEM DESCRIPTION

Many HVAC systems use a central plant to provide hot water (e.g., with temperature of 32°C) for heating purposes or cold water (e.g., with temperature of 6°C) for cooling purposes. Figure 1 is a schematic diagram of a typical HVAC system with multiple zones for cooling purpose.

The water path is presented by black arrows, which shows that cold water with temperature flows from the central plant along the pipework and into the cooling coils. The chilled water valves (CHWV) are used to control the volume of water flowing through the coils. A cooling coil exchanges heat energy between air from the mixing box and water from the central plant, with the output water flowing back to the central plant with a slightly higher temperature.

The air path is represented by grey arrows. The outside air can be drawn in by an outside air fan (OAF) and outside air damper into the mixing box. A mix of outside air and return air then flows through the air filter and passes over the cooling coil. A supply air fan (SAF) then forces the air through an insulated supply duct into the zone area as supply air (SA) for the zone. The supply air gains part of its heat from the zone to replace the heat that is leaking through the walls, roof, etc. The supply air then passes through the room and into the return air inlets. Some return air is exhausted through an exhaust air damper. The remaining return air passes through the re-circulation air damper into the mixing box and mixes with the outside air. The outside air dampers and exhaust air dampers control the amount of air flow amount from the outside air (OA) and return air (RA) streams respectively.

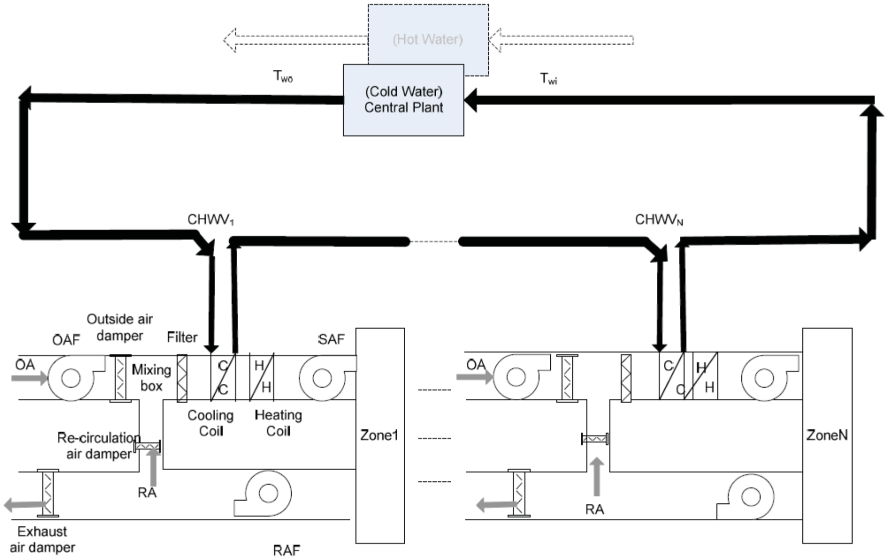

Figure 1: Schematic diagram of a typical HVAC system for cooling purpose.

3. FAULT DETECTION PROCESS

3.1 Overview of the approach

By way of overview, Figure 2 provides a flowchart illustrating the high-level processes involved in a fault detection process. The fault detection process will typically be performed by a computer system (agent). In the HVAC system example this agent may be (or form part of) the HVAC control system, or may be a standalone system which has access to the relevant data/information. The instructions may be stored on a memory of a computing device, a portable computer readable medium (e.g. a disc or a flash drive or suchlike), and/or conveyed to the computing system by means of a data signal in a transmission channel. Examples of such transmission channels include wired or wireless network connections enabled by a communications interface and various communications protocols.

As shown in Figure 2, the data from the components of the multi-variable system/agent is received. This data may, for example, include sensed data from various sensors within the agent and feedback data from various components of the agent. Additional data from external data sources is also received. Such additional data may include data from additional sensors (external agents) and/or other sources such as the Internet.

It will be appreciated that while receiving HVAC data is depicted as occurring before receiving other data, the order of these steps may be changed, or, in fact, these steps may occur concurrently. Further, and depending on the multi-variable system/agent in question, it may be the case that only data from the components of the agent is necessary to perform fault detection.

The received data is then processed in accordance with a plurality of fault detection techniques. This will be referred to as preliminary fault detection processing. Many models are generated accordingly. (Each model can be seen as an agent in the fault detection system). One specific dynamic machine-learning based fault detection technique based respectively on **Dynamic Bayesian Networks** (DBNs) is discussed below. DBN involves "learning" regular patterns of normal operation of the multi-variable system/agent, and then identifying deviations from these patterns as faults. Applying machine learning techniques to fault detection is advantageous in that they do not rely on fixed rules or models to determine a fault.

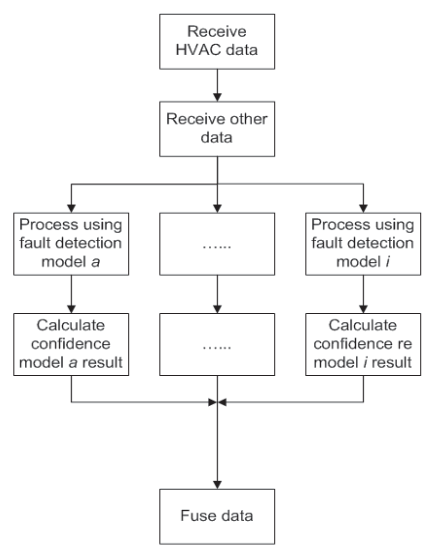

Figure 2: The high-level processes involved in a fault detection process.

The fault detection techniques may be techniques intended simply to detect abnormal operation of the multi-variable system/agent (as opposed to normal operation), techniques designed to detect specific faults in the system, or a combination of both. As discussed below, using a fault detection technique to detect abnormal operation in general versus using the same technique to detect a specific fault depends, in many instances, on how the fault detection technique is trained. As such, and for example, the fault detection techniques may include:

1. multiple instances of the same general type of fault detection technique (e.g. DBN based techniques) with each specific technique being trained to detect different faults;
2. different types of fault detection techniques (e.g. DBN based techniques, Hidden Markov Model based techniques, support vector machine based techniques, Kalman filter based techniques) trained to detect the same type of faults;
3. different types of fault detection techniques trained to detect different faults;
4. or a combination thereof.

The fault detection process may be performed concurrently (on multiple processors), or sequentially in any order desired. At next step, confidence levels of the outputs of each trained agent (model) are calculated. These calculations can be performed using data fusion technologies. The faults detected along with their corresponding levels of confidence, are fused together to provide a combined output (being either the present of a fault or otherwise) and a level of confidence associated with that output.

While the fault detection process has been depicted as a single linear process it will be appreciated that the process is continuous, in that system and other data may be continuously (or, at least, periodically) received and processed.

3.2 Preliminary fault detection

The data available (and relevant) is processed according to a plurality of fault detection techniques. The techniques may be adapted to detect generally whether the system appears to be operating normally or not, or may be trained for the detection of specific faults.

The plurality of techniques may include multiple instances of the same general fault detection technique trained on different datasets (i.e. trained to detect different types of faults), different types of fault detection techniques trained to detect the same type of fault, or different types of techniques trained to detect different faults. Employing multiple fault detection techniques and fusing the consequent data provides a valuable tool for the evaluation of risk and reliability in engineering applications where it is not possible to obtain precise measurement from experiments, or where knowledge is obtained from expert elicitation.

One specific fault detection technique based on DBN is implemented in the multi-variable systems/agents, and will be described in detail below. The DBN technique is based on statistical machine learning algorithms and serves to avoid some of the limitations of the rule and model based fault detection techniques commonly used in detecting faults in multi-variable systems/agents.

3.3 Dynamic Bayesian Network (DBN) fault detection process

A Bayesian network can be thought of as a graphical model that encodes probabilistic relationships among variables of interest. When used in conjunction with statistical techniques, the graphical model of a Bayesian network has several advantages for data analysis.

Because a Bayesian network encodes dependencies among all variables, it readily handles situations where some data entries are missing. This is a relatively common occurrence in multi-variable systems/agents. Additionally, a Bayesian network can be used to learn causal relationships between variables, and hence can be used to gain an understanding about a problem domain and predict the consequences of intervention. Further, because the Bayesian network model includes both causal and probabilistic semantics, it is a suitable representation for combining prior knowledge (which often comes in causal form) and data. Further still, Bayesian statistical methods in conjunction with Bayesian networks offer an efficient and principled approach for avoiding the over fitting of data.

Each variable in a Bayesian network has a probability distribution which describes the probability of the variable having a particular value (out of the possible values for the variable). For example, in a HVAC system one variable may be the speed of a particular fan, and the probability distribution for that fan may be produced by statistical measurement.

The Bayesian network also includes probabilistic links (arcs) between variables. These links are conditional probability distributions which describe the probability distribution of the target variable (i.e. the variable at which the link terminates) when the variable from which the link originates has a particular value. For example, and as discussed further below, in a HVAC system there may be a relationship (and hence a link) between a fan power variable and the fan speed variable. The conditional probability distribution associated with this link would be

$$p_{Fan\ speed}(fan\ speed\ value |\ fan\ power\ value).$$

A standard Bayesian network cannot represent the direct mechanism of temporal dependencies among variables. To cater for temporally dependent complex systems (which many multi-variable systems/agents, including HVAC systems, are) an additional dimension of time is added to a standard Bayesian network to generate a DBN. A DBN is an extension to a Bayesian network that represents a probability model that changes with time. For the purposes of fault detection in multi-variable systems/agents, DBNs offer a number of further improvements over standard Bayesian networks – such as relaxing some of the feedback restrictions typical of the standard directed acyclic graphs used for Bayesian networks. Accordingly, for modeling time-series data the use of directed graphical models, which can capture the fact that time flows forward, is advantageous as they enable the multi-variable system/agent to be monitored and updated as time proceeds, and to use both current and previous system behavior to predict future behavior of the system.

In a DBN, a probabilistic network is used to represent the joint probability distribution of a set of random variables. The variables of the DBN can be denoted as the states of a DBN because they include a temporal dimension. The states satisfy the Markov condition that the state of a DBN at time t depends only on its immediate past (i.e. its state at time $t-1$). As with a standard Bayesian network, a DBN uses nodes to represent variables in the system and links between the nodes to represent probabilistic relationships/dependencies between nodes (variables). In a DBN, however, links can be both between nodes in the same time-slice (intra-slice links) or between nodes in different time-slices (inter-slice links). Intra-slice links are condition probabilities between variables within the same time-slice. Inter-slice connections represent condition probabilities between variables from different time-slices. As such, each state of the dynamic model at a given time-slice t may depend both on the state of one or more nodes of the model at a previous time-slice (e.g. time t – 1), and the state of one or more nodes in the given time-slice t.

Defining a DBN requires three sets of parameters to be defined:

- The initial probability $P(x_o)$ which is the *priori* probability distribution at the initial state of the process;
- The state transition probability $P(x_t|x_{t-1})$ which specifies time dependency between the states; and
- The observation probability $P(y_t|x_t)$ which specifies the dependency of observation states regarding to the other variables at time-slice t.

Here *x and y* are the states of variables within the DBN.

Accordingly, in order to generate a DBN for a multi-variable system/agent the following need to be defined: the set of variables in the network; the initial states of each of the variables; the *a priori* probabilities for each variable in the network and their **conditional probability dependencies** (CPDs) (determined based on training datasets); and the dependency relationships between the variables. Determining the dependency relationships involves creating both intra-slice links (each intra-slice link being between two variables within the same time-slice), and inter-slice links (each inter-slice link being between a variable in a previous time-slice and a variable in the current time-slice).

Turning to the specific HVAC system example, Figure 3 shows a DBN designed for fault detection in a HVAC system. The DBN illustrates two time-slices (time t – 1) and (time t). The nodes (variables) of DBN are the various sensors and feedback mechanisms in the HVAC system. In this instance nine nodes are incorporated in the DBN: Fan Power, Fan Speed Set point, Fan Speed, Supply Air Temperature, Hot Water Valve, Chilled Water Valve, Supply Air Pressure, Return Air Damper Position, and Outside Air Damper Position. Additional variables from the data/information relevant to the HVAC system could, of course, also be included.

As can be seen, in each time-slice, solid lines are used to show the intra-slice links between nodes. For example, fan speed is related to both the fan power and fan speed set point. Broken lines are used to show inter-slice links between nodes. For example, at time t the fan speed is related to the supply air temperature at time t – 1 and the supply air pressure at time t – 1 (in addition to being related to the fan power at time t and the fan speed at time t).These dependency links/relationships are initially designed based on the understanding of the HVAC system and the experience of the building managers.

In experiments, all the nodes are continuous and all the CPDs are linear-Gaussian. As noted above, the DBN can be defined as $\{P(x_o),\ P(x_t|x_{t-1}),\ P(y_t|x_t)\}$. It is a two-slice temporal Bayesian network which defines $P(x_t \mid x_{t-1})$ by means of a directed acyclic graph (DAG) as follows:

$$P(x_t|x_{t-1}) = \prod_{i=1}^{N} P\left(x_t^i \middle| \pi(x_t^i)\right) \tag{1}$$

Here x^i stands for each node, and x^i is a node at time-slice t. x^i denotes the set of parents of x^i in the network, which can be at either time-slice t or t-1. The links joining the nodes in the first time-slice of the DBN do not have parameters (i.e. conditional probability dependencies) associated with them, but each link joining nodes in the second time-slice does have an associated CPD.

In the network, the first-order Markov assumption is satisfied. That is, the parents of a node are only in the same time-slice or the previous time-slice – i.e. all inter-slice links are from time t -1 to time t. The direction of intra-slice links can be arbitrary. Once trained, the parameters of the CPDs associated with the variables do not change over time (unless the DBN is re-trained). The semantics of DBN can be defined by "unrolling" the network to T time-slices. Hence the resulting joint probability distribution is defined by:

$$P(X_{1:T}) = \Pi_{t=1}^{T} \Pi_{i=1}^{N} P\left(x_t^i \middle| \pi(x_t^i)\right) \tag{2}$$

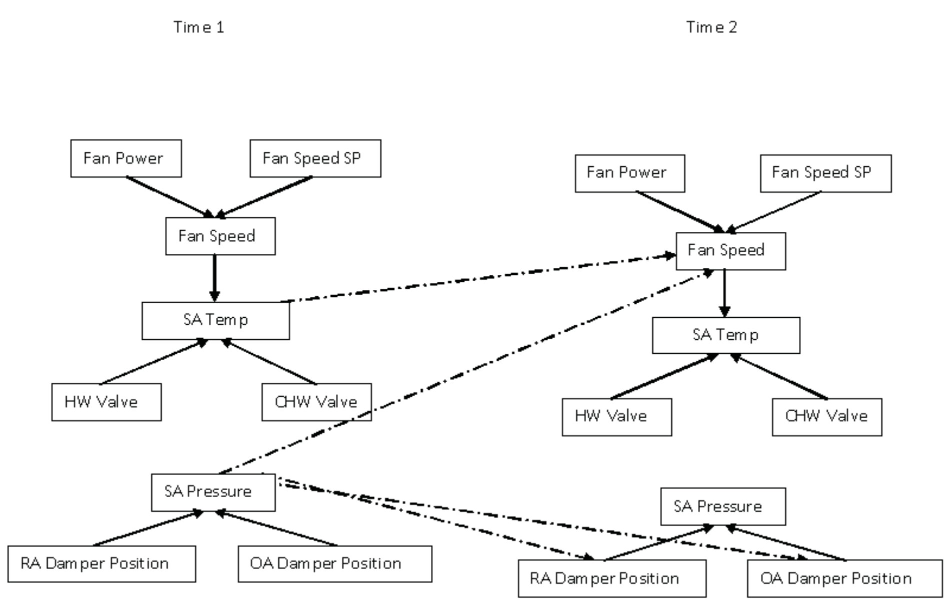

Figure 3: a DBN representative of a HVAC system.

As the structure of the DBN is pre-defined, the parameters that need to be trained are the set of conditional probabilities in equations (1) and (2) that quantify the nodes of the network. These parameters are estimated using the Maximum Likelihood (ML) algorithm, as described in [1]. The ML estimator chooses values for the conditional probabilities that maximize the probability of the training dataset.

In the present approach, the DBN is trained using measurements/data taken from the normal operation of the multi-variable system/agent (i.e. the system operating without errors or faults). The network learns the CPD parameters of the system when it runs smoothly. This trained network is then used to make inferences from testing datasets (as described below). Depending on the testing dataset, a DBN can be trained to generally detect abnormal behavior, or can be trained to detect specific faults.

For fault detection, testing datasets provide observed values of the nodes in the network and allow the probability distributions to be calculated based on these observed values. The *posterior* probabilities can then be calculated to obtain the likelihood value of the dataset, which measures how well the testing dataset fits the DBN. When a fault occurs, the likelihood value should be lower than the value with normal datasets.

The output of a DBN fault detection process, therefore, is a likelihood that provides information as to how well a new measurement (e.g. current sensor/data readings) fits to the trained model.

3.4 Fusion of preliminary fault detection data

In order to use the results of the preliminary fault detection processes, the data from the individual fault detection processes are aggregated or fused. The purpose of aggregation of information is to meaningfully summarize and simplify a corpus of data whether the data is coming from a single source or multiple sources. Familiar examples of aggregation techniques include arithmetic averages, geometric averages, harmonic averages, maximum values, and minimum values. Combination rules are special types of aggregation methods for data obtained from multiple sources, as is the case in a multi-variable system/agent.

4. TESTS OF HVAC FAULT DETECTION USING DBN

In this section, the results of two real-world tests will be discussed.

4.1 HVAC system fault detection using Dynamic Bayesian Network – experiment one

In the first test, a fault was generated by forcing a chilled water valve in a HVAC system to remain at 30% open (simulating a stuck valve). The DBN in respect of the system being tested was trained over a period of eight weeks in which the HVAC system was running normally. Consequently, in use the DBN should be able to detect abnormal operations relevant to the nodes (variables) (as in Figure 3). The fault detection procedure then used datasets from the HVAC system for five days after the system was trained. The chilled water valve in the HVAC system was set at 30% open from 10:00am to 10:45am on day two of the testing, and reset to normal operation after 10:45am.

The results of this test are shown in Figure 4. The x-axis of graph shows time and the y-axis is the log likelihood of the combined probability over the whole network – i.e. the probability distribution over U={X1, X2, …, Xk} where X is the variable expressed by the node of the network – (with the lower the likelihood, the higher the likelihood that a fault is occurring). As can be seen, the lowest points of the log likelihood are recorded during late morning on the second day of the test, matching the artificial fault introduced into the system.

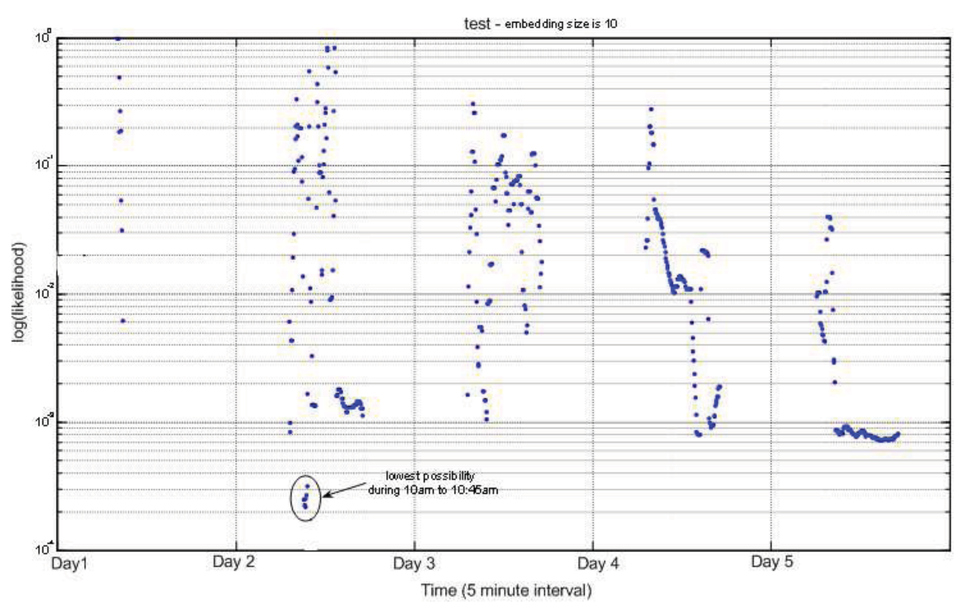

Figure 4: Test results using a DBN fault detection process in experiment one

4.2 HVAC system fault detection using Dynamic Bayesian Network – experiment two

In the second test, a fault was generated by forcing the output air damper to be stuck. The details of the experiment are described as follows:

14.54 – air handling unit (AHU) in cooling mode, high humidity day, output air damper initially closed, output air damper locked at 50% open.

14.56 – noticeable rise in chilled water valve position until 14.57. Flattens out at about 35%.

15.03 – output air damper opened to 100%.

15.04 – sharp rise in chilled water valve position. Flattens out at about 58%.

15.20 – output air damper reset to auto control (drops to 0% - closed). Chilled water valve position drops back to normal levels by 15.22.

In this experiment, nine sensors were used to build up a DBN, which was trained over the same eight weeks of normal system operation in the previous experiment. The output air damper sensor was not used in the training and testing. The results of this test are shown in Figure 5, with the x-axis being time. The green line in Figure 5 shows the opening percentage of the chilled water valve (as shown on the left y-axis), and blue line shows the log likelihood of the combined probability (as shown on the right y-axis). As can be seen, the testing results show a period marked as dots in the afternoon (2:55pm -3:25pm) which match to the fault that was set in the experiment.

5. DISCUSSIONS

The present paper relates to systems and methods for detecting faults in multi-variable systems/agents. The approach is particularly suitable for detecting faults in heating, ventilation and air-conditioning (HVAC) systems and will be described in relation to that exemplary but non-limiting embodiment.

Many current fault detection techniques for multi-variable systems/agents are rule-based. The fault detection system integrates and interprets incoming data in accordance with a pre-determined set of rules, produces a risk profile, and autonomously initiates a response to a breach of these rules. Rule-based systems are very specifically derived for/tailored to a particular system and are very difficult to update, change, or adapt to a different system. Additionally, rule-based systems typically fail miserably if conditions beyond the boundaries of the knowledge incorporated in them are encountered.

The fault detection approach in this paper is based on statistical machine learning technology. This is achieved by learning the consistent nature of normal HVAC operation, and then using the statistical relationships between groups of measurements to identify anomalous deviations from the norm and identify faults in all subsystems for which sensor information is available, regardless of the specifics of the installation.

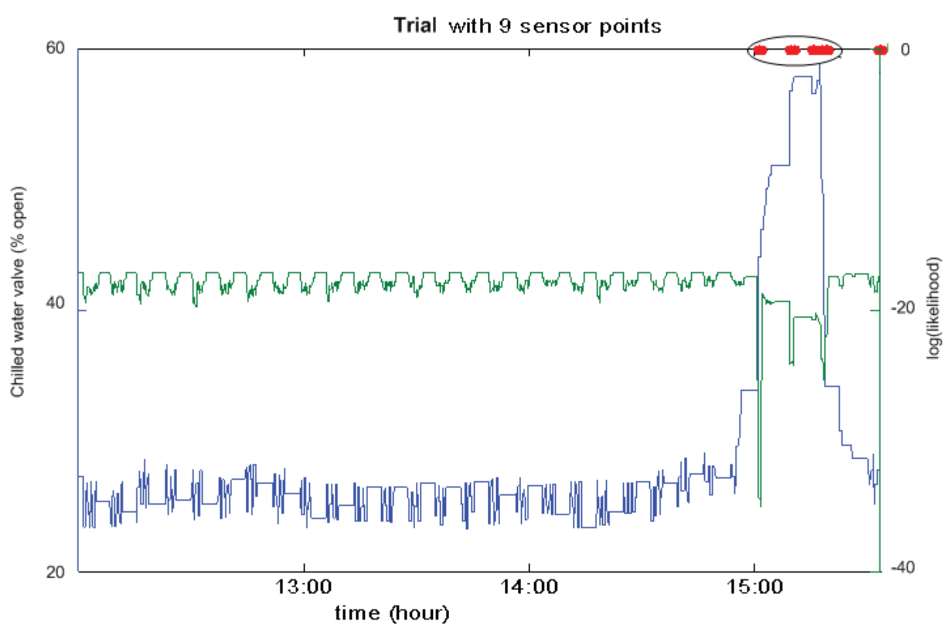

Figure 5: test results using a DBN fault detection process in experiment two

The approach uses the Dynamic Bayesian Networks to model the dynamical sub-systems (agents) and sequence data in HVAC system. These models (agents) are constructed via a learning process from some training data of normal running HVAC systems. The trained models (agents) can then be used for automatic fault detection. Our algorithm can capture the fact that time flows forward by using directed graphical models (agents), it is adaptive to environment changes and is reliable in HVAC systems. The approach has been tested based on real data from commercial HVAC systems. We can successfully detect a number of typical faults. The experimental results are all very positive.

Using the adaptive learning approach requiring only incumbent sensors' data, this technique offers several advantages over rule-based techniques. In particular, the amount of customization, configuration and expert knowledge required to implement such techniques on different HVAC systems is greatly reduced. The results obtained by using the methods prove that this algorithm has the property of robustness and adaptability. Current intention can achieve better fault detection results when the building situation and/or environment properties changes, while traditional approaches are normally not adaptable to these changes automatically.

6. REFERENCES

[1] Myung, I.J. 2003. Tutorial on maximum likelihood estimation. Journal of Mathematical Psychology 47, 90-100.

[2] D. Claridge, M. Liu, W.D. Turner, Whole Building Diagnostics, in: Proceedings of the Diagnostics for Commercial Buildings: Research to Practice Workshop, held June 16th & 17th; 1999, Pacific Energy Center, San Francisco, Lawrence Berkeley National Laboratory.

[3] M.A. Piette, S. Kinney, P. Haves, Analysis of an information monitoring and diagnostic system to improve building operations. Energy and Build¬ings 33 (8) (2001) 783-791.

[4] T1AX, Energy Consumption Characteristics of Commercial Building HVAC Systems, vol. Ill: Energy Savings Potential, Final Report to U.S. Department of Energy, Office of Building Technologies, 2002.

[5] D. Westphalen, K.W. Roth, System and component diagnostics, ASHRAE Journal 45 (4) (2003) 58-59.

[6] Kaler, G.M. Jr., Expert system predicts service, Heating, Piping & Air Conditioning, v 60, n 11, Nov, 1988, p 99-101.

[7] Schein, Jeffrey; Bushby, Steven T.; Castro, Natascha S.; House, John M., A rule-based fault detection method for air handling units, Energy and Buildings, v 38, n 12, December, 2006, p 1485-1492.

[8] Song, Young-hak; Akashi, Yasunori; Yee, Jurng-Jae, A development of easy-to-use tool for fault detection and diagnosis in building air-conditioning systems. Energy and Buildings, v 40, n 2, 2008, p 71-82.

[9] Wagner, J.; Shoureshi, R., Failure detection diagnostics for thermofluid systems. Journal of Dynamic Systems, Measurement and Control, Transactions of the ASME, v 114, n 4, Dec, 1992, p 699-706.

[10] Haves, Philip; Salsbury, Timothy I.; Wright, Jonathan A., Condition monitoring in HVAC subsystems using first principles models. ASHRAE Transactions, v 102, n 1, 1996, p 519-527.

[11] Bendapudi, Satyam; Braun, James E.; Groll, Eckhard A., Dynamic model of a centrifugal chiller system - Model development, numerical study, and validation. ASHRAE Transactions, v 111 PART 1, 2005, p 132-148.

[12] Peitsman, Henk C.; Bakker, Vincent E., Application of black-box models to HVAC systems for fault detection. ASHRAE Transactions, v 102, n 1, 1996, p 628-640.

[13] Najafi, Massieh; Auslander, David; Bartlett, Peter; Haves, Philip, Application of machine learning in fault diagnostics of mechanical systems. Proceedings of the World Congress on Engineering and Computer Science Oct 2008, San Francisco, USA.

[14] Najafi, Massieh; Auslander, David; Bartlett, Peter; Haves, Philip; Sohn, Michael, Modeling and Measurement Constraints in Fault Diagnostics for HVAC Systems. Lawrence Berkeley National Laboratory: Lawrence Berkeley National Laboratory. LBNL Paper LBNL-3903E. 2010.

[15] Rossi, T.M., and J.E. Braun. A statistical, rule-based fault detection and diagnostic method for vapor compression air conditioners. International Journal of Heating, Ventilating, Air Conditioning and Refrigerating Research 3(1), p 19-37.

[16] Yoshida, Harunori; Yuzawa, Hideki; Iwami, Tatsuhiro; Suzuki, Masami, Typical faults of air-conditioning systems and fault detection by ARX model and extended Kalman filter. ASHRAE Transactions, v 102, n 1, 1996, p 557-564.

[17] Peitsman, Henk C.; Soethout, Luc L., ARX models and real-time model-based diagnosis. ASHRAE Transactions, v 103, n 1, 1997, p 657-671.

[18] Yoshida, Harunori; Kumar, Sanjay, ARX and AFMM model-based on-line real-time data base diagnosis of sudden fault in AHU of VAV system. Energy Conversion and Management, v 40, n 11, Jul, 1999, p 1191-1206.

[19] Lee, W. Y., J.M. House, C.Park, and G.E.Kelly. Fault diagnosis of an air-handling unit using artificial neural networks. ASHRAE Transactions 102 (1), p 540-549.

[20] Wang, Shengwei; Jiang, Zhiming, Valve fault detection and diagnosis based on CMAC neural networks. Energy and Buildings, v 36, n 6, June, 2004, p 599-610.

[21] Najafi, Massieh, Auslander, David, Bartlett, Peter, Haves, Philip. Application of machine learning in fault diagnostics of mechanical systems. Proceedings of the World Congress on Engineering and Computer Science Oct 2008, San Francisco, USA.

Approaching Decentralized Demand Side Management via Self-Organizing Agents

Christian Hinrichs Ute Vogel Michael Sonnenschein
Dept. of Environmental Computer Science, C.v.O. University of Oldenburg, Germany
{hinrichs,vogel,sonnenschein}@informatik.uni-oldenburg.de

ABSTRACT

This paper provides an introduction to self-organizing mechanisms in the domain of energy systems. We motivate the use of such mechanisms by outlining the drawbacks of present control systems and providing some advantageous properties of self-organizing systems. We address the problem of a supply and demand matching of a large number of distributed actors and give a mathematical formalisation of this optimization problem. Based on that we propose the algorithm DSS which uses the stigmergy paradigm to form a distributed search heuristic. Afterwards the DSS-T is introduced, a refinement of the DSS algorithm which incorporates mechanisms from the tabu search. To examine the performance of the algorithms, we implemented a simulation framework which lead us to the problem of local mutual exclusion with fairness. Therefore we describe this problem and its origin and then propose the ARP, a protocol which guarantees local mutual exclusion, fairness and starvation freedom.

Finally, we give some first insights into the performance of the DSS and DSS-T algorithms by examining a single simulation scenario for reference.

Categories and Subject Descriptors

I.2 [**Artificial Intelligence**]: Distributed Artificial Intelligence

General Terms

Algorithms

Keywords

self-organization, emergence, distributed heuristics, demand side management, supply and demand matching

1. INTRODUCTION

The term *Demand-Side Management* (DSM) [4] appeared in the last third of the past century. First research focused on manually controlling large electrical loads (like industrial plants) in order to compensate stress in the power grid. On the one hand, these methods incorporated direct control signals which allowed immediate control over the loads. On the other hand, indirect signals (i.e. dynamic pricing) have been studied. These form incentives for appliance owners to regulate their energy use. This was followed by methods to sense and react to stress situations in the grid automatically. For an overview of these strategies the reader may refer to [5, 2, 8]. However, these approaches were centrally organized and not suitable for a large amount of individually controllable units.

Therefore, strategies of decentralized control have been investigated. Intelligent appliances which autonomously monitor properties of the power grid (i.e. frequency or voltage) and independently react to stress situations have been proposed and tested [19, 6, 17]. In these approaches however, the lack of coordination lead to oscillation effects due to synchronization of appliance reactions. So there is a need for communication to allow for intelligent coordination. Some work has been done to address the oscillation problem [9]. Though this – to a certain extent – can be done in a centralized fashion, a more scalable and elegant way would be given by a completely decentralized control and coordination scheme. This requires the appliances being able to autonomously act in a system-stabilizing way.

But how can we design such a distributed artificial intelligence? We investigated this question from a nature-driven perspective. In nature there are many systems that fulfill a specific task without central coordination (e.g. the foraging of ants). Those systems reveal a number of advantageous properties like adaptivity, self-healing, scalability, redundancy and more [16]. If we now examine the above question from a more abstract point of view, we come to the following conclusion: In a decentralized DSM-enabled system, a large number of agents perform autonomous actions and form a collective behaviour to fulfill a global objective. Such a system is called *self-organizing*.

In this paper, we present our ongoing work on self-organization and emergence in energy systems. We believe that such nature-inspired algorithms and models can effectively be applied to the task of stabilizing the electricity grid, i.e. by providing balancing power to compensate fluctuating energy resources. The paper is organized as follows: Section 2 describes the concept of self-organizing systems in detail. In section 3 we present related attempts (own and other) in this area. Section 4 then formalizes the addressed optimization problem of a supply and demand matching and proposes DSS and DSS-T, first approaches to self-organizing algorithms for this problem which form distributed search heuristics based on the stigmergy paradigm. Afterwards, in section 5 the simulation framework and the activity restriction protocol ARP are introduced before finally some tentative interpretations of the performance of the proposed algorithms are given.

2. SELF-ORGANIZATION AND EMERGENCE

There is no common definition of self-organization. In biology, a self-organizing process produces a global pattern purely on a basis of local actions, without knowledge of the global pattern or system [1]. In the field of multiagent systems there is a distinction between self-organization and *emergence*. In [16] the former is defined as a mechanism or process which allows a system to (re)organize during runtime without explicit external instruction. The latter is described as a phenomenon that becomes visible at the global (macro) level of the system through interactions of elements at the local (micro) level. This phenomenon can be a structure, behaviour or functionality of the system. Note that emergence is not an essential property of self-organizing systems. However, engineering such systems often aims at the emergence of specific properties, leaving self-organization being a tool to make emergence possible. In the following, we will use the term self-organization synonymously for self-organizing systems that exhibit emergent properties.

Derived from the definitions above, self-organizing systems consist of at least two distinct activity levels. There is the macro level, which is an observation of the system from an external point of view. The individual system elements may be visible, but their internal processes are unknown. Observing a school of fish, for example, one can perceive the overall structure of the system, as well as its behaviour. That is, it forms a swarm which consists of several fishes. The swarm moves around, with no identifiable central leader to coordinate the movement. It is able to evade obstacles and enemies and may even split up into multiple parts without losing its abilities. These observable characteristics are the emergent properties of the system. To understand the origin of those, we must take a view into the system and inspect the micro level. That consists of fishes with individual properties (i.e. position, direction, velocity) and behaviours. Regarding the fishes as agents, the behaviours are expressed as rules which modify the agent's properties. The rules are *local*: they consider only the properties of the agent itself, and those of nearby agents. That is important, because it implies that no agent is able to perceive the whole system (and therefore: the macro level). The term "nearby", however, has to be defined specifically for each system. The behaviour of swarming animals like fishes, birds and gregarious land animals has been studied in [15]. The author has identified three principal rules that, when simulated in an artificial system, lead to the formation of swarms:

1. Collision Avoidance: avoid collisions with nearby objects.
2. Velocity Matching: attempt to match velocity and movement direction with nearby individuals.
3. Swarm Centering: attempt to stay close to nearby individuals.

Applying these simple behavioral rules (with descending priority) to a number of simulated fishes in parallel, they will

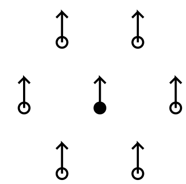

Figure 1: Equilibrium condition of the centered individual in a school of fish. Arrows represent direction/velocity vectors.

form a swarm with the above characteristics. Note that none of these rules takes the whole system into account. But since the action of any indivdual *triggers* reactions of other individuals, a global behaviour emerges. So there is a mechanism of self-organization: every single individual acts on behalf of its rules, aiming for a local equilibrium (the fulfillment of a local goal, here rather: of all its rules). In a school of fish, an individual reaches its equilibrium condition if all its neighbours are equally far away positioned and have all the same movement direction and velocity as itself (see figure 1). In the next timestep however, this condition will be broken, because due to the third rule the outer fishes will try to move towards the swarm center, disturbing the equilibrium of the centered fish.

The example "school of fish" shows the following: There are individuals at the micro level of the system, acting upon simple local rules and each striving after an equilibrium (fulfillment of the local goal). If all individuals were in equilibrium condition, the system would have organized itself. But as soon as an external disturbance occurs, the individuals would try to fulfill their local goals again, effectively reorganizing the system. Since this reorganization occurs without external guidance, it is called self-organization. The emergent properties of the system are either produced by the reorganization process itself (as in the school of fish), or by reaching the equilibrium state. An example for the latter would be the termite nest-building. Here, the emergent property is the optimal protection of the termite queen. This is reached by building the nest equally around the spatial position of the queen. If the queen moves its permanent residence, the worker termites start to rebuild the nest so that the queen will afterwards again be positioned centrally in the nest. Upon reaching this goal, the system has reorganized and will again optimally protect the queen. This example shows an additional aspect: An internal *central* entity in a self-organizing system is allowed, as long as this entity only triggers the self-organization process but does not control the individual actions.

Self-organizing systems can be applied to a number of tasks, especially optimization problems. For a classification of self-organization mechanisms and possible applications, the interested reader may refer to [7, 12].

3. RELATED WORK

In the domain of energy systems, the concept of *virtual market places* has intensively been discussed in recent years. In those systems, electrical loads (consumers) as well as energy resources (producers) are actors at a virtual market place. There they offer or demand blocks of electrical power. The market clearing is reached by adjusting prices and therefore matching the offered and demanded amounts. Because the actors react differently and autonomously to price adjustments, the system is partly self-organized. The emergent property is a *supply and demand matching* of electrical power. Two of such systems are described in [10] and [20]. Because of their statical hierarchical design however, market place systems are also centrally organized.

Another partly self-organized approach has been proposed in [11]. There, communication is carried out indirectly via a global black board. A central broker agent publishes the global optimization goal, while each other participating agent (producers and consumers) publishes its predicted operational mode (amount of electrical power demanded/supplied). These agents will then iteratively modify and publish their operational modes according to the global goal, published information from other agents and local degrees of freedom, until the system converges or is stopped by a special condition. The modification of local operational modes is carried out via a genetic algorithm.

In [14], a system based on *overlay trees* has been proposed. In such a system, the agents organize themselves in a virtual tree structure and perform an iterative bottom-up supply and demand matching. For that task, each agent sends its degrees of freedom (in terms of possible operational modes) to its parent node in the tree. The parent node then selects the locally optimal operational modes and sends the selection back to the children, who have to adhere to the selected modes. These actions are performed at each level of the tree, until the root node has been reached. Coupled with a dynamically built tree structure, this system could be able to emerge an optimal global solution.

In [18] the concept of *holonic virtual power plants* has been studied. The idea is somehow related to the virtual tree structures. A holon can be composed of a single or more agents and even other holons. To the outside, the holon acts like a single agent. Therefore, the terms "agent" and "holon" can be used synonymously here. Additionally, a holon can be part of several holons at the same time. The structure is hierarchical, but overlapping and highly dynamic. In conjunction with distributed scheduling mechanisms, the author proposed a holonic system which is able to fulfill a global optimization goal (i.e. building a virtual power plant) in a self-organizing way.

4. APPROACHING SELF-ORGANIZATION

Most of the presented work in the previous section incorporated some kind of centrality or static organization. In this paper we propose a completely decentralized, highly dynamic and self-organized system of autonomous agents who work cooperatively on local goals in order to emerge a solution to a global optimization goal. For that objective, we presume that the monetary outcome of the system will be distributed equally among the agents. Hence, we are modelling agents which have no interest in competing with each other. Instead, they try to maximize their profit by cooperation. In the following section, we construct a formal model of the considered problem.

4.1 Formalisation

4.1.1 The optimization problem

Let $g : \mathcal{T} \to \mathbb{R}$ be a timeseries of amounts of electrical power. Time is discretized into time slots $t \in T$. The set of agents can be described by $\mathcal{A} = \{a_i \mid 0 \leq i < n\}$ with n being the number of agents. Each agent represents an actor which has the ability to adapt its mode of operation for each time slot. These operational modes therefore describe an either positive (producer) or negative (consumer) amount of electrical power. Additionally, they are rated by the agent with respect to how beneficial it is to choose the specific mode. For example, a CHP might rate its possible operational modes with respect to their efficiency. According to these ratings, the agents cannot be purely altruistic, but rather benevolent while being constrained by their device specific properties. So the set of rated operational modes for an agent a_i for a time slot t is defined as

$$\mathcal{P}_i(t) = \{(p, r)_k \mid 0 \leq k < m_i\} \tag{1}$$

where $p = p_{i,k}(t)$ is an operational mode, $r = r(p_{i,k}(t))$ is the rating for this mode and m_i is the number of possible modes for this agent in the given time slot.

The global goal of the system now is to select an operational mode for each agent for each time step so that a given load profile g is approximated as closely as possible. For the time being, we will neglect any constraints between operational modes.

A vector $s(t)$ of operational modes (one for each agent) with associated ratings for a time slot t is called a solution vector and is defined as

$$s(t) = (y_0, \ldots, y_n)\ , \quad y_i \in \mathcal{P}_i(t) \tag{2}$$

The matrix $s(\mathcal{T})$ then aggregates $s(t)$ for all $t \in \mathcal{T}$ and is called a *solution.* Now assume a rating function f which rates a solution vector $s(t)$ of selected operational modes for a specific time slot, and a second rating function F that aggregates the ratings given by f over time, effectively rating the solution $s(\mathcal{T})$. An example for f could be

$$f\left(g(t), s(t)\right) = \overline{r}_{s(t)} \cdot \left| \left(\sum_{p \in s(t)} p \right) - g(t) \right| \tag{3}$$

where

$$\overline{r}_{s(t)} = \frac{\sum_{r \in s(t)} r}{n} \tag{4}$$

is the mean of all ratings in $s(t)$. So f computes the absolute value of the difference between the superimposed operational modes and the load profile at the given time slot, and multiplies that with the overall rating. Then an example for F could be

$$F(g, s) = \sum_{t \in \mathcal{T}} f\left(g(t), s(t)\right) \tag{5}$$

so that F provides the sum of the ratings of f over time.

To find an optimal set of operational modes for a given load profile g, a solution s_{opt} has to be identified that minimizes F:

$$s_{opt} = \operatorname*{argmin}_{s} \left(F\left(g, s\right)\right). \tag{6}$$

The solution space to this problem is given as

$$\mathcal{S} = \bigtimes_{t \in \mathcal{T}} \bigtimes_{i \in \mathcal{A}} \mathcal{P}_i(t) \tag{7}$$

yielding

$$|\mathcal{S}| = m^{n \cdot |\mathcal{T}|} \tag{8}$$

if we assume $m = m_i \ \forall i \in \mathcal{A}$.

```
repeat:
   if local environment changes:
      for each considered time slot:
         adapt own selection
      if selection has changed:
         publish new selection
```

Figure 2: DSS – A distributed search algorithm based on the stigmergy paradigm.

4.1.2 Communication

In a self-organizing system, agents act upon their local perception. In the present case, we model this perception via communication between agents. Hence, agents can be *connected* to other agents, which means they know each other and can communicate. Agents without a direct connection cannot communicate with each other. So the communication subsystem can be expressed as an undirected graph

$$G = (\mathcal{A}, \mathcal{E}) \tag{9}$$

where the set of vertices is the set of agents $\mathcal{A}$ and the edges $\mathcal{E}$ are communication links. Now, as suggested in section 2, we may define the term "nearby" as a kind of *locality*: Each agent a_i has a defined limited neighbourhood $\mathcal{N}_i$ of other agents with whom it can communicate. So $\mathcal{N}_i$ is defined by

$$\forall a_i \in \mathcal{A}\, \forall a_j \in \mathcal{N}_i : (a_i, a_j) \in \mathcal{E},\ \mathcal{N}_i \subseteq \mathcal{A},\ a_i \neq a_j. \tag{10}$$

Additionally

$$(a_i, a_j) \in \mathcal{E} \Rightarrow (a_j \in \mathcal{N}_i \wedge a_i \in \mathcal{N}_j) \tag{11}$$

guarantees that communication links are bidirectional.
Considering locality, we may introduce the *local solution vector* $s_i(t)$ of an agent a_i which reflects the local view of the agent at the system for a specific time slot t. It is defined as

$$s_i(t) = (y_{i_0}, \ldots, y_{i_l})\ , \quad y_{i_j} \in \mathcal{P}_j(t)\ , a_j \in \mathcal{N}_i \cup \{a_i\} \tag{12}$$

and is the subset of $s(t)$ which is currently visible to a_i with respect to its neighbourhood (including its own selection). Accordingly, $s_i(\mathcal{T})$ is the *local solution* visible to a_i.

4.2 Self-organizating algorithms

4.2.1 A first approach

As a first attempt to a completely distributed and self-organizing solution for the given problem, we have designed an algorithm which incorporates the *stigmergy* paradigm. In this paradigm, coordination is carried out via monitoring and modifying the environment rather than direct agreements between agents [12]. For the given case, we let each agent follow the *Distributed Stigmergy Search* (DSS) algorithm shown in figure 2. The local environment of an agent a_i here corresponds to the local solution without its own selection: $s_i(\mathcal{T}) \setminus \mathcal{P}_i(\mathcal{T})$. So if any of the neighbours changes any of its selected operational modes, the agent a_i will adapt its own selection to the new circumstances. To accomplish that, a_i selects those own operational modes (one for each time slot) which yield the best rating for the current local solution. This corresponds to a local search which yields

$$s_{opt_i} = \underset{s_i}{\operatorname{argmin}}\left(F\left(g, s_i\right)\right)$$

```
repeat:
   if local environment changes:
      for each considered time slot:
         sort op-modes by estimated improvement
         find first element s in the sorted list
                     which is not in the tabu list
         select s and add to tabu list
         if size of tabu list > l:
            remove oldest element from tabu list
      if selection has changed:
         publish new selection
```

Figure 3: DSS-T – Including tabu search in DSS.

(cmp. eq. 6), where only the components of s_i which belong to a_i are modified. If the new selection differs from the former, the agent then publishes that new selection, thus triggering activity in its neighbours.

In this mechanism, agents are in equilibrium condition if they are not able to improve the local solution by selecting other operational modes. If the whole system is in equilibrium condition, each agent has reached a pareto optimum in its local search space. Unfortunately this doesn't mean that the global solution $s(\mathcal{T})$ is a pareto optimal solution, too. For example consider a system $\mathcal{A}$ consisting of three agents connected in a row so that $a_0 \in \mathcal{N}_1$, $a_1 \in \mathcal{N}_2$ but $a_2 \notin \mathcal{N}_0$. Let $\mathcal{A}$ be in equilibrium condition, so no agent may choose an operational mode which improves its local solution. However, as a_0 and a_2 are not connected, there might be operational modes which would in fact improve the global solution of the system while impairing the local solution. Hence, the attraction of these local pareto optima of each agent rapidly leads to global equilibium condition and therefore a termination of the algorithm in a suboptimal state.

4.2.2 Including tabu search

Therefore, our work aims at providing a distributed heuristic that is able to generate "good" solutions in a short amount of time, providing all advantages of self-organizing systems but without guaranteeing a specific solution quality. Such a mechanism should in the long term converge to an optimal solution s_{opt}, while allowing temporary deteriorations like the famous *tabu search* algorithm in order to explore the whole search space. We incorporated the ideas of tabu search into the algorithm DSS-T by providing each agent with a local tabu list with a maximum length l_i (see figure 3). In this algorithm, for each time slot t an agent first calculates a list of its possible operational modes, sorted by their estimated improvement. It then selects the first element of this list which is not contained in the tabu list. Thus, the best operational mode which is not currently marked "tabu" is selected. Upon selection, the operational mode is added to the tabu list. As the size of this list is bounded by l_i and in each iteration it is truncated to this size, the re-selection of the currently added operational mode is effectively excluded from the next l_i iterations. After performing this procedure for all considered time slots, the new selection is published just like in the DSS algorithm, triggering activity in the neighbourhood of the agent.

The tabu list mechanism enables the agent to select operational modes which impair its local solution but might possibly lead to a global optimum. In contrast to the DSS algorithm, the DSS-T will not rapidly terminate. As every agent definitely will make a new selection when triggered by the environment, the system will never reach a global equilibrium condition (unless we set $l_i \geq m_i$, in that case every operational mode would be chosen exactly once before the algorithm terminates). This effect introduces the problem of *termination*: How can agents decide *not* to make a new selection but instead stick to the current one? And what kind of mechanism can guarantee that all agents make this choice when the system currently exhibits a near-optimal global solution? These questions cannot be answered yet, but will be addressed in our future work.

5. EVALUATION

5.1 Simulation framework

To evaluate the proposed (and prospective) self-organized algorithms for solving the given optimization problem, a simulation framework has been implemented. The framework is written in the Python programming language and is able to simulate an arbitrary number of autonomous software agents as defined in section 4.1. Each of these *workers* is identified by an URI of the form (*host:port*) and has the ability to send and receive messages to and from other agents. The messages are sent as network packages, so that the whole simulation can be carried out either locally or distributed. As stated in section 4.1.2, local perception of the workers is carried out via communication with their neighbours. We make the following assumptions about communication:

Assumption 1. *Identity*
All nodes have unique IDs.

Assumption 2. *No lost messages*
All sent messages will eventually be received.

Assumption 3. *Message travelling time*
The message travelling time is constant so all messages are received in the order they have been sent.[1]

5.2 Activity restriction protocol

In the so far outlined paradigm however, conflict situations could arise if two neighbouring workers sense each other and act upon this perception *simultaneously*, effectively hindering each other's goal achievement. For example, let a system $\mathcal{A}$ consist of two agents a_0, a_1 with $\mathcal{T} = \{t_0\}$ and $g(t_0) = 0$. Each agent has the same two equal rated possible modes of operation so that

$$\mathcal{P}_0(t_0) = \{(-1.0, 1.0)_0, (1.0, 1.0)_1\},$$
$$\mathcal{P}_1(t_0) = \{(-1.0, 1.0)_0, (1.0, 1.0)_1\}.$$

Now let an initial solution be

$$s_{init} = ((-1.0, 1.0), (-1.0, 1.0))$$

so that a_0 as well as a_1 have their first operational mode selected. The rating for this solution is $F(g, s_{init}) = 2$. Now both agents get active simultaneously. According to DSS or DSS-T, they will now both perform a local search and select an operational mode which improves the current solution. In their local view, the second operational mode $(1.0, 1.0)_1$ would reduce the rating to 0. Therefore, both agents choose their second mode so that afterwards

$$s_{post} = ((1.0, 1.0), (1.0, 1.0))$$

which still yields $F(g, s_{post}) = 2$. This procedure is then reversed in the next iteration and the system will begin to oscillate: $s_{init}, s_{post}, s_{init}, s_{post}, \ldots$ while the rating constantly remains at $F(g, s) = 2$.

Therefore an activity restriction protocol has been designed which prevents such conflicts. Given a set of interconnected agents expressed as a graph G as defined in section 4.1.2, the activity restriction protocol prevents the simultaneous action execution of neighbouring agents. A related problem is known as *mutual exclusion in networks*. However, in the given case we aim at mutual exclusion only between neighbouring agents. So any number of nodes in the graph is allowed to be active at the same time, as long as they are not directly connected. This is similar to the *generalized dining philosophers problem*. In the following solution we included *fairness* so as to when a set of neighbouring nodes repeatedly and simultaneously request activity, it is granted rotatory to the involved nodes so that all nodes will equally often become active in the long term. To accomplish that, we have assigned each node a_i a state variable $st_i \in \{I, R, A\}$ where $I :=$ *inactive*, $R :=$ *requesting* and $A :=$ *active*, and an activity value $act_i \in \mathbb{N}$. All nodes start with $st_i = R$ and $act_i = 0$. A node a_i may at any time change its state from I to R, signaling an activity request. Upon changing the state from R to A however, the activity value of the node is incremented by 1. Additionally, this transition is bounded by the condition

$$\forall a_j \in \mathcal{N}_i : f_{mutex}(i, j) > 0 \tag{13}$$

where

$$f_{mutex}(i, j) = \begin{cases} 1 & \text{if } st_j = I, \\ 2 & \text{if } st_j = R \wedge f_{act}(i, j) > 0, \\ 0 & \text{else,} \end{cases} \tag{14}$$

and

$$f_{act}(i, j) = \begin{cases} j - i & \text{if } act_i = act_j, \\ act_j - act_i & \text{else.} \end{cases} \tag{15}$$

This condition ensures that no neighbouring node is already active or is requesting with a lower activity value. In case of equal activity values, the node id is crucial. So the transition R to A can be expressed as

$$R \xrightarrow[act_i = act_i + 1]{\forall a_j \in \mathcal{N}_i :\, f_{mutex}(i,j) > 0} A \tag{16}$$

where the condition is placed above the transition arrow, and the required action which has to be taken below.

A node may stay arbitrary but finitely long in the active state so that we assume the following:

Assumption 4. *Finiteness*
A node in the active state will eventually become inactive.

The possible state transitions are shown altogether in figure 4.

[1] This is a quite restrictive assumption which will very likely become a topic in our future work.

$$\mathrm{I} \longrightarrow \mathrm{R} \xrightarrow[act_i = act_i + 1]{\forall a_j \in \mathcal{N}_i : f_{mutex}(i,j) > 0} \mathrm{A}$$

Figure 4: State transition diagram with $\frac{condition}{action}$ annotation for the activity restriction protocol.

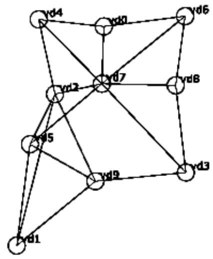

Figure 5: Connection graph of the examined scenario.

The correctness of the activity restriction protocol is shown in appendix A.

5.3 Simulation results

As this is a report on ongoing work, we will only examine a simplistic scenario here. We consider just one time slot

$$\mathcal{T} = \{t_0\}$$

where the global optimization goal is zero:

$$g(t_0) = 0.$$

There are ten agents ($n = 10$) with each four random modes of operation ($m = 4$) in the range [-1000,1000] so that

$$p_{i,k} = random(-1000, 1000),\ 0 \leq i < n,\ 0 \leq k < m.$$

Furthermore, all mode ratings are equal:

$$\forall p_{i,k} :\ r(p_{i,k}) = 1.$$

The connection graph is randomized and is shown in figure 5. And finally, as we use the DSS-T algorithm, the tabu list length is

$$\forall a_i :\ l_i = \frac{m_i}{2} = 2.$$

According to eq. 8 there are 1.048.576 possible solutions in this scenario. We calculated these with a brute force algorithm for reference.

Figure 6 shows the progress of the simulation of the scenario. The y-axis denotes the rating value while the x-axis displays the iterations of the simulation. An iteration is defined as a selection of an operational mode of any single agent. So after each of the ten agents has chosen its initial operational mode, the simulation would have proceeded at least ten iterations. Note that because the agents start without a selected operational mode, each agent must make its initial choice before the solution vector is called valid (see eq. 2). This point in time t_{init}, where all agents have chosen their initial solution, is marked by the vertical bar in the diagram. In the present case $t_{init} = 14$ because some agents became active more than once before all agents did their initial choice. The straight horizontal lines indicate ratings for the best (lower) and worst (upper) reference solutions. The optimal solution in this case has a rating of zero. In between, the progress of the algorithm is visible.

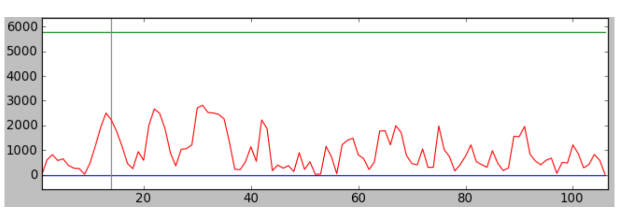

Figure 6: Exemplary simulation run.

It is noticeable that the system oscillates between excellent and average solutions, which is due to the tabu list. The algorithm seems to be able to generate near-optimal solutions in a very short amount of time in comparison to the size of the search space. However, this is based on a very few experiments done so far and therefore quite speculative. We are presently conducting a more formal analysis of a large number of simulation runs with varying random seeds and also different scenarios.

6. CONCLUSIONS

In this paper we have motivated the use of self-organizing mechanisms in the domain of energy systems. We provided an introduction to some relevant topics in this area: emergence, disturbance, equilibrium condition, cooperation, locality. The addressed optimization problem has been formalized in mathematical terms. We proposed the distributed search algorithms DSS and DSS-T which are based on the stigmergy mechanism in combination with a local search. A protocol for local mutual exclusion with fairness in a network of distributed processes has been proposed in order to avoid oscillations of suboptimal solutions. We proved the correctness of this protocol. Finally we gave some tentative interpretations of the performance of the algorithms.

In our future work we will address the formal analysis of the proposed algorithms. We will examine the reduction of problem difficulty which is given by distributing the calculations. Furthermore, we will compare the proposed approach to the well known `Adopt` algorithm [13] from the field of *Distributed Constraint Optimization* as well as to the `Max-Sum` algorithm which has been proposed in [3]. Besides that, we will develop other self-organizing algorithms for the presented problem.

7. ACKNOWLEDGMENTS

We would like to thank the Universitätsgesellschaft Oldenburg e.V. (UGO) for supporting this work. Furthermore, we thank the reviewers of this paper for their very constructive contributions.

8. REFERENCES

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

[2] J. Eto. The past, present, and future of U.S. utility demand-side management programs. Technical Report

LBNL-39931, Lawrence Berkeley National Laboratory, 1996.

[3] A. Farinelli, A. Rogers, A. Petcu, and N. R. Jennings. Decentralizsed coordination of low-power embedded devices using the max-sum algorithm. In Padgham, Parkes, Müller, and Parsons, editors, *Proc. of 7th In. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, pages 639–646, 2008.

[4] C. Gellings and J. Chamberlin. *Demand-Side Management: Concepts And Methods*. The Fairmont Press, Inc., 1988.

[5] C. Goldman and M. Kito. Review of demand-side bidding programs: Impacts, costs, and cost-effectiveness. Technical Report LBL-35021, Lawrence Berkeley National Laboratory, 1994.

[6] D. J. Hammerstrom, J. Brous, D. P. Chassin, G. R. Horst, R. Kajfasz, P. Michie, T. V. Oliver, T. A. Carlon, C. Eustis, O. M. Jarvegren, W. Marek, R. L. Munson, and R. G. Pratt. Pacific Northwest GridWise(tm) Testbed Demonstration Projects: Part II. Grid Friendly(tm) Appliance Project. Technical Report PNNL-17079, Pacific Northwest National Laboratory, Richland, WA, 2007.

[7] L. Hassas, G. Serugendo, A. Karageorgos, and C. Castelfranchi. On self-organising mechanisms from social, business and economic domains. *Informatica*, 30:63–71, 2006.

[8] G. Heffner. Configuring load as a resource for competitive electricity markets - review of demand response programs in the U.S. and around the world. Technical Report LBNL-51496, Lawrence Berkeley National Laboratory, 2002.

[9] C. Hinrichs, U. Vogel, and M. Sonnenschein. Modelling and evaluation of desynchronization strategies for controllable cooling devices. In I. Troch and F. Breitenecker, editors, *Proc. Mathmod 2009 - 6th Vienna International Conference on Mathematical Modelling*, number 35 in Argesim Report, Vienna, Austria, 2009.

[10] K. Kok, M. Scheepers, and R. Kamphuis. *Intelligent Infrastructures*, chapter Intelligence in electricity networks for embedding renewables and distributed generation, pages 179–209. Intelligent Systems, Control and Automation: Science and Engineering Series. Springer, 2009.

[11] J. Li, G. James, and G. Poulton. Set-points based optimal multi-agent coordination for controlling distributed energy loads. In *Proceedings of the Third IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO'09)*, San Francisco, California, USA, 2009.

[12] J.-P. Mano, C. Bourjot, L. Gabriel, and P. Glize. Bio-inspired mechanisms for artificial self-organised systems. *Informatica*, 30:55–62, 2006.

[13] P. Modi, W. Shen, M. Tambe, and M. Yokoo. Adopt: Asynchronous distributed constraint optimization with quality guarantees. *Artificial Intelligence Journal (AIJ)*, 161:148–180, 2005.

[14] E. Pournaras, M. Warnier, and F. Brazier. Local agent-based self-stabilisation in global resource utilisation. *International Journal of Autonomic Computing*, 1(4):350–373, 2010.

[15] C. W. Reynolds. Flocks, herds and schools: A distributed behavioral model. *SIGGRAPH Comput. Graph.*, 21:25–34, August 1987.

[16] G. Serugendo, M.-P. Gleizes, and A. Karageorgos. Self-organisation in multi-agent systems. *The Knowledge Engineering Review*, 20(2):65–189, 2005.

[17] J. Short, D. Infeld, and L. Freris. Stabilization of grid frequency through dynamic demand control. *IEEE Transactions on Power Systems*, 22(3):1284–1293, 2007.

[18] M. Tröschel. *Aktive Einsatzplanung in holonischen Virtuellen Kraftwerken*. PhD thesis, Universität Oldenburg, 2010.

[19] D. Trudnowski, M. Donnelly, and E. Lightner. Power-system frequency and stability control using decentralized intelligent loads. In *Proceedings of the 2005/2006 IEEE PES T&D Conference and Exposition*, Dallas, TX, 2006.

[20] H. Wedde, S. Lehnhoff, C. Rehtanz, and O. Krause. Bottom-up self-organization of unpredictable demand and supply under decentralized power management. In *Proceedings of the 2nd IEEE International Conference on Self-Adaptation and Self-Organization (SASO'08)*, Venice, Italy, 2008. IEEE Press.

APPENDIX

A. CORRECTNESS OF ARP

A.1 Local mutual exclusion

THEOREM 1. *The action restriction protocol guarantees that no two neighbouring nodes are in the active state at the same time, so that*

$$\forall (a_i, a_j) \in \mathcal{E} : (st_i = A \Rightarrow st_j \neq A) \wedge (st_j = A \Rightarrow st_i \neq A).$$

PROOF. Assume two connected nodes are simultaneously in the active state, so that

$$\exists (a_i, a_j) \in \mathcal{E} : st_i = st_j = A.$$

As all nodes initially started in the inactive state, the nodes a_i, a_j had each to perform the transitions $I \rightarrow R \rightarrow A$ at least once, both conforming to the condition of $R \rightarrow A$. Then there are two possible cases:

1. The transitions $R \rightarrow A$ were performed sequentially, so that after the first transition $st_i = A \neq st_j$. Then because of $f_{mutex}(j, i) = 0$ the transition $R \rightarrow A$ of a_j cannot be performed, contradicting our assumption.

2. The transitions $R \rightarrow A$ of both nodes were performed simultaneously. This implies $st_i = st_j = R$ at the time of the transition. Then according to eq. 14:

$$f_{act}(i, j) > 0 \ \wedge \ f_{act}(j, i) > 0$$

and because of eq. 15:

$$\begin{array}{ll} & (act_i = act_j \quad \Rightarrow \quad i > j \ \wedge \ j > i) \\ \vee & \\ & (act_i \neq act_j \quad \Rightarrow \quad act_i > act_j \ \wedge \ act_j > act_i) \end{array}$$

which contradicts our assumption.

Therefore $\nexists (a_i, a_j) \in \mathcal{E} : st_i = st_j = A$. □

A.2 Progress

THEOREM 2. *The action restriction protocol is deadlock free.*

PROOF. Assume a deadlock where an arbitrary number of nodes (at least one) are in the requesting state but none of them is able to perform the $R \rightarrow A$ transition, so that

$$\exists \mathcal{R} \subseteq \mathcal{A}, |\mathcal{R}| > 0 : \quad \forall a_i \in \mathcal{R} : \quad st_i = R \ \wedge \\ \forall a_j \in \mathcal{N}_i : f_{mutex}(i,j) \leq 0.$$

Let a_i be a node which has no neighbour with a smaller act_j so that

$$\exists a_i \in \mathcal{R} : \forall a_j \in \mathcal{N}_i : act_i \leq act_j.$$

There are two cases to be considered:

1. $\exists a_j \in \mathcal{N}_i : act_i = act_j$, then

$$f_{mutex}(i,j) > 0 \quad \text{iff} \quad i < j$$

   and

$$f_{mutex}(j,i) > 0 \quad \text{iff} \quad j < i.$$

2. $\forall a_j \in \mathcal{N}_i : act_i < act_j$, then

$$f_{mutex}(i,j) > 0.$$

Both cases lead to

$$\exists a_k, a_l \in \{a_i, a_j\} : f_{mutex}(k,l) > 0 \ , \ k \neq l$$

which contradicts the assumption. □

A.3 Starvation freedom

THEOREM 3. *Any node in the requesting state will eventually become active.*

PROOF. Let a_i be a node in the requesting state which cannot become active so that

$$\exists a_j \in \mathcal{N}_i : act_j \leq act_j.$$

Considering the progress property (theorem 2) and the fact that every node must increase its activity value upon entering the active state, all neighbouring nodes will eventually get a higher activity value than a_i, enabling a_i to become active. □

Finding Optimal Dispatch Schedule of Power Generation Using Locational Marginal Pricing in a Dynamic Market

Von-Wun Soo, Yi-Je Tsai
National Tsing-Hua University, Taiwan
{soo@cs, g9762645@oz}.nthu.edu.tw

ABSTRACT

During these years, the regulatory framework for the wholesale sector of the electricity industry has been replaced by market competition in many countries. There are many ongoing debates over market design issues such as how to avoid market oligopoly problem, how to properly implement a retail electricity market and how to increase profits of market players. This thesis focuses on the development of a model and the simulation of electricity market and studies the impacts of power losses on electricity market equilibrium states in order to find an optimal dispatch schedule of distributed power supplies to power demand in the power delivery networks. To simulate the interaction among strategic behaviors of market players, we generalize a multi-agent system with the supply function equilibrium (SFE) model and take power losses into consideration. We conduct the experiments on IEEE 30-bus systems to illustrate the performance of the proposed method and present some numerical results that are effective in showing various economic impacts of power losses including system benefits.

Categories and Subject Descriptors
J.7 [Industrial control]: Electricity market, Multi-agent system, Optimal power generation dispatch schedule

General Terms
Algorithm, Performance, Design, Economics, Experimentation

Keywords

Dynamic electricity market, Optimal power generation dispatch, Supply function equilibrium, Locational marginal price

1. INTRODUCTION

In recent years, competitive retail market models in power markets have been introduced to replace traditional wholesale ones in order to increase social welfare and improve market efficiency. By repeated practical experiments, many valuable experiences have been gained [1-6]. However, a perfectly competitive market does not exist in practice. Many issues such as market design flaws, market power and inherent engineering features of power systems can contribute to the market inefficiency. Electricity market has some characteristics that impact the market design. These characteristics include inelastic demand, limited transmission capacity and supply-demand balance requirement. And these characteristics make the electricity market more similar to an oligopoly one where some suppliers can control the market clearing price (MCP) via their strategic bidding behaviors because purchasers never reduce their demand. Therefore, how to improve the bidding strategies of suppliers to maximize system profit of has been a widely discussed problem.

Numerous methods such as genetic algorithm [7], Markov decision process [8], Lagrangian relaxation [9] are brought up to solve the strategic bidding issues. And then, game theoretical approaches [10-13] are also addressed and used to characterize the strategic behavior of market player. Game theoretical approaches include Bertrand, Cournot and supply function equilibrium (SFE) models. Among these methods, only SFE model reflect the bidding rules in electricity market where SFE requires individual supplier to submit its supply function to the market operator. Therefore, SFE is the most widely used to investigate the strategic bidding problem recently [14-19].

The concept of SFE was originally developed by Klemperer and Meyer in [14] and then applied by Green and Newbery as a model for players' behavior in a competitive market in England and Wales [1]. Furthermore, Green introduced a linear version that greatly simplifies the SFE mathematics and applied it to reduce the generation cost of the electricity industry in England and Wales [4]. These literatures show that SFE model can predict prices that match the empirical data reasonably well. And rather than simple price-quantity pairs in Cournot and Bertrand models, the decision variables in SFE model are the parameters of supply function of market players. Furthermore, SFE model can also deal with inelastic demand cases. The above merits prompted the applications and studies of SFE model. In [16], Rudkevich discussed some basic issues about learning through the SFE model. Niu et al. [17] modeled the electric firms' bidding behaviors with SFE model, and studied the effects of forward contracts on ERCOT market. Then Anderson and Xu discussed symmetric duopoly SFE in the case with contracts and price caps in [19].

Although the SFE model works well in implementing electricity market, the result it derived may not be the optimal solution because the underlying physical network also has impacts on efficiency of supplied power. Power losses of transmission network may make energy delivered from the least expensive supplier to a different location become more costly. Thus, it is important to consider network transmission issues such as transmission capacity limit and consume of power while delivering on physical network in our electricity market design. Locational marginal price (LMP) is defined as the incremental cost of supplying per unit power energy to a specific location. It considers not only the cost of producing energy, but also the loss effect of power transmission. Federal Energy Regulatory Commission (FERC) first proposed a notice for a LMP -based market pricing approach as a mechanism to build efficient energy market [20]. Then LMP has become a part of standard market design issue and discussed in many literatures [21-25]. LMP modeling affects the power dispatching and market clearing price. We show that bidding strategies based on the LMP model is more cost-effective than bidding with only generation cost for sellers in this thesis.

However, market clearing price is not only affected by the individual strategic bidding behavior of a market player, but also by the interactions among market players in the repeated auctions. Therefore, it is necessary to consider each player as an individual agent which has its own decision making policy and make them communicate with each other. Von der Fehr and Harbord [26] described a multi-unit auction model for the electricity market of England and Wales. Chen-Ching Liu [6] also developed an agent-based market model that conducting learning behavior during bidding period to improve the bidding strategies of market players.

The remainder of this thesis is composed as follows. Section 2 introduces the SFE model and LMP computation method. In this section, we present a SFE model that incorporates supplier agents and the market operator in our inelastic demand system. The theory and application of LMP are also included in this section. Section 3 provides the algorithm of our system. A multi-agent method of dynamic electricity market structure is presented in this section. Section 4 is the numerical results and the analysis of these results. The multi-agent algorithm described in section 3 is implemented by JADE [27], a JAVA-based agent development framework. We test the system with different cases and illustrate the results in section 4. Section 5 offers conclusions of our work.

2. MARKET ARCHITECTURE

2.1 SFE Electricity Market Modeling

Electricity market is composed of suppliers, load serving entities (LSEs) and market operator. Suppliers and LSEs submit their offer or demand bids to the market operator. Then market operator determines the market clearing price and the dispatch schedule. Suppliers decide their offer bids based on their own marginal cost function. In this thesis, we do not consider the demand bids of LSEs because the load is almost inelastic in the market.

2.1.1 Supplier Agent

A Supplier denotes the agent that generates and sells power to the electricity market. Although supplier agents may have multiple generation plants on the same bus, it can be assumed without losing generality that each supplier agent has only one generation plant for simplicity. Generally, most of electricity market researchers consider the generation cost function of suppliers as a quadratic function. Assume S denotes the set of supplier agents in electricity market, P_{gi} is the provided generation output of supplier agent i. Then the generation cost function of supplier agent i is defined as Equation (1)

$$C_i(P_{gi}) = a_i * (P_{gi})^2 + b_i * P_{gi} \qquad (1)$$

Where $C_i(P_{gi})$ is generation cost of supplier agent i at power generation P_{gi}. Coefficients a_i and b_i are known as strategic variables that supplier agents can adjust the values in order to increase their profits.

The marginal generation cost function $MC_i(P_{gi})$ of supplier agent i is obtained by taking the partial derivative of $C_i(P_{gi})$ with respect to P_{gi}:

$$MC_i(G_i) = \frac{\partial Ci(P_{gi})}{\partial Gi} = 2*a_i*G_i + b_i* P_{gi} \qquad (2)$$

Each supplier agent calculates its optimal amount of supply offer using the marginal generation cost function and then submits to the market.

2.1.2 Market Operator

Market operator implements the market clearing based on the offer bids of suppliers and demand bids from LSEs. Because the demand side bidding is not considered, the intention of market operator is to minimize the cost of supplying the total required demand while satisfying the system constraints such as transmission flow constraint and generators' power generation capacity. In SFE approach, the market operator uses the following model to complete market clearing:

$$C_g = \min \sum_{i=1}^{N} Ci\ (P_{gi}) \qquad (3)$$

s.t.

$$P_k = P_{gk} - P_{dk},\ k=1 \ldots N_B \qquad (4)$$

$$F_{mini} \le F_i \le F_{maxi},\ i=1 \ldots N_L \qquad (5)$$

$$P_{minj} \le P_j \le P_{maxj},\ j=1 \ldots N_G \qquad (6)$$

where

Cg: Generation cost of system

$C_i(P_{gi})$: Generation cost of generator i

N_B: Number of buses in the system

N_L: Number of transmission lines in the system

N_G: Number of generators in the system

P_k: Net power injection at bus k

P_{gk}: Power offer at bus k

P_{dk}: Power demand at bus k

F_i: Power flow on transmission line i

F_{mini}: Lower bound of power flow on transmission line i

F_{maxi}: Upper bound of Power flow on transmission line i

P_j: Power offer of generator j

P_{minj}: Lower power offer bound of generator j

P_{maxj}: Upper power offer bound of generator j

Equation (4) is subject to the real power balance constraints at each bus. The power offer has to be the sum of power demand and net power injection. Equation (5) is transmission flow limit for each line. And Equation (6) is the generator limit of suppliers.

In a competitive market, prices of electricity should not be set by a single regulator but by the auction behaviors of market players. An auction mechanism is a market institution that contains rules for allocating resource and setting price based on the bids of market players. Market operator does not control the offer and price of suppliers. It just adjusts the price that reflects the supply offer and demand bid curves to determine the market clearing price and corresponding power dispatch schedule. The objective, which is represented by Equation (3), of minimizing the cost of supplying is actually obtained by auction of market players rather than market operator.

2.2 Locational Marginal Price

Although the SFE model introduced in Section 2.1 successfully minimizes the generation cost to increase the global welfare, it does not consider physical transmission network. If all generators are located in one bus or are geographically close, power losses is physically reasonable to neglect in calculating the optimal dispatch. However, if generators are spread out on the transmission network, the power losses must be considered, and this will modify the optimal generation dispatch assignments.

To take a simple example, suppose that all the generators on the transmission network are identical. Then, we expect that drawing most heavily from the generators closest to the demands will be cheaper than drawing identically from all generators considering the power losses.

2.2.1 Power Losses Calculation

The standard approach for calculating power losses can be expressed by the law of overall energy balance for a power system:

$$P_{loss} = \sum_{i=1}^{n} Pi = \sum_{i=1}^{m} Pgi - \sum_{i=1}^{n} Pdi \qquad (6)$$

Where P_{loss} is the total power losses and the P_i is the bus powers. Since the demand is fixed, it can be seen from Equation (6) that P_{loss} depends only on the P_{gi}. However, we can consider only m-1 of the P_{gi} are independent variables and the remaining one P_{gj} is a dependent variable found by solving the power flow equation, where bus j is referred to as a slack or swing bus and is usually defined as bus 1 in the transmission network. Thus, for a given system, the functional dependence of P_{loss} can be written as:

$$P_{loss} = P_{loss}(P_G) \triangleq P_{loss}(P_{g2}, P_{g3}, \ldots, P_{gm}) \qquad (7)$$

It should be noted that Equation (7) is not an explicit formula but depends on the solution of implicit equations, the power flow equations.

2.2.2 Penalty Loss Factor

We extend the optimization problem described by Equation (3) and state a new version that involving the power losses as follow:

$$C_g = \min \sum_{i=1}^{N} Ci\,(P_{gi})$$

s.t.

$$\sum_{i=1}^{m} Pgi - P_{loss}(P_{g2}, P_{g3}, \ldots, P_{gm}) - P_D = 0 \qquad (8)$$

$$F_{mini} \leq F_i \leq F_{maxi},\ i=1\ldots N_L$$

$$P_{minj} \leq P_j \leq P_{maxj},\ j=1\ldots N_G$$

Equation (8) is the power balance constraint involving power losses. In order to obtain a formal solution, we use the Lagrange multipliers method to define the augmented cost function:

$$C_g = \sum_{i=1}^{N} Ci\,(Pgi) - \gamma\left(\sum_{i=1}^{m} Pgi - P_{loss}(P_{g2}, P_{g3}, \ldots, P_{gm}) - P_D\right) \qquad (9)$$

Next find a stationary point of C_g:

$$\frac{dCg}{d\gamma} = \sum_{i=1}^{m} Pgi - P_{loss} - P_D = 0 \qquad (10)$$

$$\frac{dCg}{dPg1} = \frac{dC1(Pg1)}{dPg1} - \gamma = 0 \qquad (11)$$

$$\frac{dCg}{dPgi} = \frac{dCi(Pgi)}{dPgi} - \gamma * \left(1 - \frac{\partial Ploss}{\partial Pgi}\right) = 0 \qquad i=2, \ldots, m \qquad (12)$$

The Equation (12) can also express as:

$$\frac{1}{1-\frac{\partial Ploss}{\partial Pgi}} * \frac{dCi(Pgi)}{dPgi} = \gamma \qquad i=2, \ldots\ldots, m \qquad (13)$$

Next, we define a variable for the calculation of power losses on each generator, the penalty loss factor L_i for the generator i:

$$L_1 = 1 \qquad (14)$$

$$L_i = \frac{1}{1-\frac{\partial P_{loss}}{\partial Pgi}} \qquad i=2, \ldots\ldots, m \qquad (15)$$

By Equation (14) and (15), the necessary conditions for the optimization problem which given in Equation (11) and (12) may be expressed by:

$$L_1\frac{dC_1(P_{g1})}{dP_{g1}} = L_2\frac{dC_2(P_{g2})}{dP_{g2}} = \ldots\ldots = L_m\frac{dC_m(P_{gm})}{dP_{gm}} = \gamma \qquad (16)$$

Where $\frac{dC_m(P_{gm})}{dP_{gm}}$ is the marginal generation cost of generator m. Therefore, the optimal dispatch with consideration of system losses is obtained by operating all generators, not at their generation limits, to satisfy $L_i * MC_i(P_{gi}) = \gamma$

3. SYSTEM DESIGN

3.1 Market Structure

Most of literatures consider the electricity market structure as a day-ahead market. The day-ahead market works as follow. During the morning of day D, LSEs submit their demand bids and price of the day D+1, suppliers submit their supply offers as well. These bids are processed using a matching algorithm: the suppliers' bids are ordered by increasing prices; the demand bids are ordered by decreasing prices. This matching algorithm results in aggregated supply curves and demand curves. The intersection of both curves is the market clearing price.

However, the day-ahead market structure is based on a static analysis of market equilibrium. It pays only a little attention to market evolution in the situation of repeated bidding behaviors and the interactions between market players. And problems may occur to the day-ahead market when the transmission network or generators breakdown. Therefore, we implement a different market structure that can dynamically adjust the equilibrium point when generation conditions change.

In this dynamic market structure, suppliers do not submit their supply and price bids in advance and the market operator does not determine market clearing price by using the matching algorithm. Instead, market operator adjusts the market price based on the difference between offers and demands. If offers are more than demands, market operator decreases the market price to give the suppliers less incentive to offer the electricity; On the other hand, if demands are more than offers, market operator increases the market price to give the suppliers more incentive to offer. And suppliers submit their offer bids to the market operator based on their marginal cost function and the market price at that time after every adjustment of market price. This repeated bidding and adjusting behaviors will eventually force the market to become a balanced equilibrium. The equilibrium price is then the market clearing price. Since under this dynamic market structure, suppliers conceal its own supply function. What market operator knows is only the offer bids of suppliers. Therefore, we assume that every supplier is honest. They always offer as much generation as its supply function decides.

3.2 Dynamic LMP Bidding Algorithm

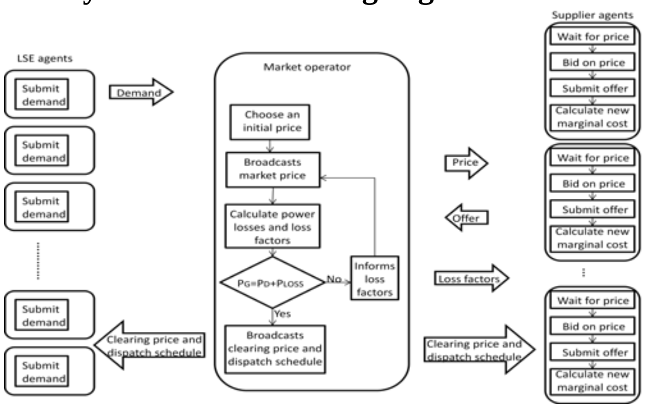

Figure 1. Flowchart of dynamic bidding algorithm

In this section, we present our algorithm for implementing a dynamic LMP bidding algorithm that considers the power losses in a multi-agent system. In our system, there are three types of agent involved in the implementation: Supplier, LSE and Market operator. The following sequence is the algorithm of our system as shown in Fig 1:

1. Initially, LSEs set up their demand bid data. The suppliers determine their marginal generation cost function $MC_i(P_{gi})$. And the market operator creates the framework of the electricity transmission network.

2. LSEs submit their demand bids to market operator, and suppliers set their penalty loss factor L_i as 1.

3. The market operator chooses an initial price p to be the market price for the 1^{st} bidding stage and broadcasts p to suppliers, then starts the market balancing procedure.

4. Suppliers set their marginal cost as the function of penalty loss factor and marginal generation cost $F_i(P_{gi}) = L_i{*}MC_i(P_{gi})$ and bid on market price p, then submit their offer to the market operator.

5. Market operator computes the power losses (P_{loss}) and the penalty loss factor L_i of every supplier. Then the market operator determines the difference between power offer (P_G) and power requirement (P_D+P_{loss}). If $P_G = P_D + P_{loss}$, then the market is balance and go to step 6; If not, go to step 7.

6. The market operator broadcasts the market $_{clearing}$ price and the dispatch schedule to market participants.

7. The market operator adjusts the market price p based on the difference between total power offer and total power requirement. If power offer is larger than power requirement, decrease the market price; otherwise, increase the market price. Then market operator informs every supplier their penalty loss factor and the market price p. Go back to step 4.

Since in our model, the demand is inelastic. LSE agents do not involve in the repeated bidding stages. Therefore, LSE agents only submit their demand bids to the market operator in the beginning of the auction. During the repeated bidding stages, suppliers bid on current market price by their marginal cost function $F_i(P_{gi})$. In Generally speaking, the more electricity suppliers offer, the more profits suppliers get as long as $F_i(P_{gi})$ is smaller than market price. Therefore, suppliers will offer as much electricity P_{gi} as $F_i(P_{gi})$ that equals to the market price in every bidding stage to maximum their profit. This bidding strategy leads to the minimum cost of supplying.

In the first bidding stage, every supplier sets their loss factor as 1 since the power losses has not yet available and every supplier has not generated power losses. Setting the penalty loss factor as 1 make suppliers' marginal cost function $F_i(P_{gi})$ equals to $MC_i(P_{gi})$, suppliers beginning the first bidding stage without the consideration of power losses. In order to compute the loss factor, the knowledge of the global framework such as transmission lines conductance and resistance, generation of each suppliers, and demands on each bus must be established. Since the market operator has all the information we need, suppliers' loss factor are computed by market operator and then broadcasts to each supplier.

The market price adjustment is also executed by the market operator based on the difference between offer and requirement and the price adjustor factor α. Price adjustor factor is a variable that controls the rate of price variation and also the equilibrium rate of bidding curve. The market price of bidding stage n is computed as follow:

$$p_{n+1}=p_n-\alpha{*}(\text{power offer} - \text{power requirement})$$

The value of α has to be selected carefully since it may lead to a divergent bidding if α is too large or slow convergence if it is too small. The impacts of α value will be discussed in Section 4.3.2.

4. SIMULATION RESULTS

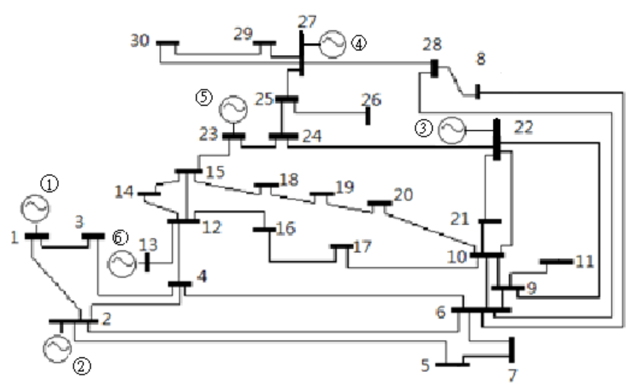

Figure 2. IEEE 30-bus system

We implemented our algorithm in JADE, an agent development framework in JAVA to simulate the multi- agent environment in the electricity market. And we use MATPOWER [30] to simulate the underlay transmission network. MATPOWER is a package of MATLAB for solving power flow problems. To combined MATPOWER into our JAVA based system, we use a toolbox of MATLAB which is called java-builder to make MATPOWER into a JAVA package. We then use this JAVA package in our system. In order to illustrate results of our algorithm, the IEEE six-generator 30-bus system [30] shown in Fig 2 is used for testing with CPU Intel 1.73GHz and 2GB main memory running under the windows operating system. The number in circle indicates the index of a power generator while the number without circle indicates the index of a bus. We will analysis the results between our system and the purely SFE model bidding strategy in section 4.1. And then in section 4.2, we proposed an experiment on observing the effects of our system due to changes of the offer or demand conditions during the bidding stages. In section 4.3, we will discuss how the initial price and price adjustor factor affect the rate of market equilibrium.

4.1 Tests in Randomly Generated cases

We examine our system in three cases which demand bids on each bus are randomly generated and compare the results with purely SFE model. In these cases, the coefficients of marginal generation cost functions and generation limits of generators are not identical, as listed in Table 1. The marginal generation cost of generator i is calculated as $MC_i(P_{gi})=a_i * P_{gi}+b_i$.

Table 1. Marginal generation costs and limits of suppliers

Supplier	ai	bi	Generation limit
1	0.0222222	20	150
2	0.01	30	80
3	0.25	10	50
4	0.0775795	20	55
5	0.01	30	30
6	0.0322581	30	40

In these cases, the initial price is set to 20, and the price adjustor factor α is set to 0.005. The initial price and price adjustor factor do not affect the market equilibrium point but only on equilibrium rate, we will show that in section 4.3.

Experiment results are summarizes in Table 2, 3, and 4. The total demands are 189.2(MW), 242.5(MW), and 312.5(MW) for cases 1, 2 and 3 respectively. In all three cases, the demands are almost offered by supplier 1 since it has the lowest marginal generation cost and also the highest generation upper bound.

Table 2. Computation result of case 1

	Purely SFE	LMP bidding algorithm	Variation ratio(%)
Pg1	132.1	112.66	-14.7
Pg2	0	0	0
Pg3	31.74	37.04	+16.7
Pg4	37.83	50.71	+34
Pg5	0	0	0
Pg6	0	0	0
PG	201.67	200.41	-0.6
Losses	12.47	11.21	-10.1
Cost	4323.8381	4116.2654	-4.8

In case 1, demand bids on bus 7 and 8 are more than other buses. And as the variation ratio of each supplier shows, we have less willing to derive generation offer from supplier 1 since its geographically farness from bus 7 and 8. Therefore, the generation output of supplier 1 is reduced in our system since supplier 1 generates more power losses than supplier 3 and 4. In this experiment, we reduced 10% of power losses and about 4% of generation cost as shown in Table 2.

The same phenomenon also appears in case 2 because demand bids are almost focused on bus 25, 29 and 30 which are far away from supplier 1, 2, and 3. Therefore, the variation ratio of supplier 2 and 3, especially suppler 2, are largely negative. The generation output of these two suppliers in our system is a lot less than in purely SFE model. Table 3 is the result of case 2.

Table 3. Computation result of case 2

	Purely SFE	LMP bidding algorithm	Variation ratio(%)
Pg1	150	132.53	-11.6
Pg2	12.21	0	-100
Pg3	40.48	32.43	-19.8
Pg4	55	49.27	-10.4
Pg5	12.21	22.44	+83.8
Pg6	3.78	20.06	+430.6
PG	273.69	256.73	-6.2
Losses	31.19	14.23	-54.4
Cost	6384.7168	6094.8845	-4.5

As shown in Table 3, the generation output of supplier 2 is all reduced in LMP bidding algorithm. Although suppler 1 also generates more power losses than other suppliers, due to its lowest generation cost, supplier 1 still offers the most generation in LMP bidding algorithm and makes the variation unobvious. And in this case, supplier 5 and 6 are nearest to the bus 25, 29 and 30. In addition, these two suppliers have a large generation cost and thus provide less generation in purely SFE model. These reasons make the variation ratio of supplier 5 and 6 a large positive number. In particular, supplier 6 offers about 4 times of generation output of purely SFE model. In this case, LMP bidding algorithm reduced 4.5% of the generation cost and half of the power losses.

Table 6. Computation result of case 3

	Purely SFE	LMP bidding algorithm	Variation ratio(%)
Pg1	150	150	0
Pg2	43.28	39.72	-8.2
Pg3	41.73	43.09	+3.2
Pg4	55	55	0
Pg5	30	30	0
Pg6	13.41	10.87	-18.9
PG	333.42	329.49	-1.2
Losses	20.92	16.99	-18.8
Cost	8317.5556	8201.7683	-1.3

The demands in case 3 are a lot more than those in case 1 and case 2, and the demands are evenly distributed on every bus. Thus, the variation ratio of each supplier except supplier 6 is relatively small as shown in Table 6. It means that the distribution of power offers makes no much different between two models. Therefore, the variation ratio of generation cost is only 1.2%, but LMP bidding algorithm still reduced 18% of power losses in this case.

As shown in Table 2, 3 and 4, our system is more efficient in both generation output and generation cost than purely SFE model in all three cases. In case 1 and 2, the demand bids of system are unbalanced. Therefore, purely SFE model generates more excess losses in these two cases and highlights the efficiency of LMP bidding algorithm.

4.2 Reactions for the changes in conditions

The electricity market is composed by suppliers, LSEs, a market operator, and the transmission network. Any change in these components will cause variation in the whole system and has to be fixed immediately to remain stable. In this section, we examine the reactions of our system when any generation condition changes in terms of the market bidding curve. Market bidding curve is a record of the market price at each bidding stage.

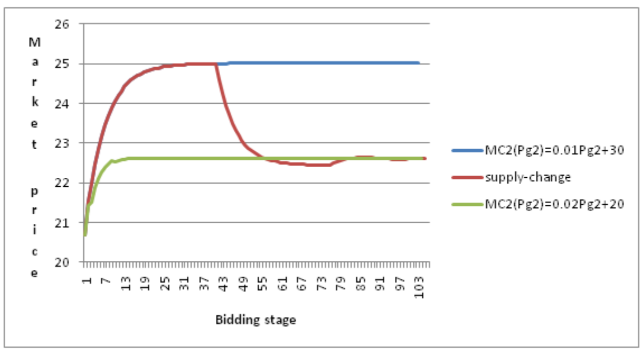

Figure 3. Market bidding curves of supply-change test

First, we assume at 40th bidding stage of case 1, supplier 2 suddenly decreases its marginal generation cost function from $MC_2(P_{g2})=0.01P_{g2}+30$ to $0.02P_{g2}+20$. And the resulting market bidding curve is shown in Fig 3. Before the 40th bidding stage, the market bidding curve of the supply-change experiment is identical to the market bidding curve of $MC_2(P_{g2})=0.01P_{g2}+30$. After that, the market bidding curve of the supply-change experiment begins to decline and finally reaches an equilibrium at the same price as the bidding curve of $MC_2(P_{g2})=0.02P_{g2}+20$.

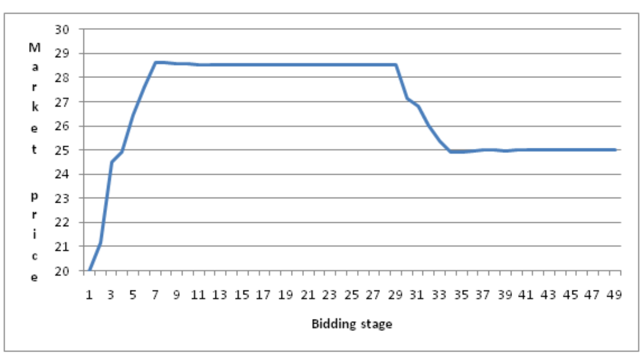

Figure 4. Market bidding curves of demand-change test

Next, we examine the reaction of our system with respect to the change in demand distribution. Although the demand is inelastic in electricity market in general, the demand bid of each LSE might change every 5-minute, half-hour or an hour depending on the real market nature. In the Day-ahead market structure, LSEs should submit their demand bid for each hour in a day to the market operator.

In our dynamic system, LSEs can change their demand bid at any time and market operator starts to adjust the market in order to remain balanced when LSEs change their demand bids. Fig 4 indicates the market bidding curve of the demand-change experiment. In this demand-change experiment, the market reaches an equilibrium point at 14th stage. After a period of time, LSEs changes their demand bids at 30th stage and the market operator begins to balance the market immediately. At 43rd stage, the market reaches a new equilibrium point. The resulting market price is equal to the clearing price in case 1 where the demand bids of LSEs at 30th stage is exactly the same as in case 1.

4.3 Influence of market operator variables

In our algorithm, the generation offers and demands are controlled by suppliers and LSEs. What market operator does is balancing the market according to the offer and demand bids. Therefore, the market operator cannot affect the equilibrium point of the market. Even though, the market operator has the ability to control the process of balancing the market to the equilibrium point by two controllable variables: initial price and price adjustor factor. In the remaining of this section, we discuss the effects of different initial prices and price adjustor factors.

4.3.1 Initial price

Initial price is the market price that market operator chooses as the beginning price of the market balancing procedure. The market operator then adjusts the market price at next bidding stage based on the market price of current bidding stage until it reaches the equilibrium point.

To examine the effect of the initial price, we use case 1 as example. We proposed four experiments in case 1 and set the initial price of these four experiments at different constant values: 0, 10, 20, and 30 respectively. The numbers of bidding stages in the four experiments are listed in Table 7.

Table 7. Results of static initial price tests

Initial price	0	10	20	30
Bidding stages	80	70	59	61

As shown in Table 7, the experiments which initial prices is 0 and 10 take more time to reach the equilibrium point than those 20 and 30. And since the clearing price is $25, we can draw a conclusion that it takes less time to clearing the market if the initial price is closer to clearing price.

The conclusion then leads to a new question: How does the market operator choose an initial price that is close to the clearing price? Apparently, using a fixed initial price in every instance is not a good answer. It is better if we can adjust the initial price to fit in different cases. And we notice that the clearing price at equilibrium point is affected by the demand and offer bids. Since the demand is inelastic, it is simpler to adjust the initial price depends on the demand bids.

To dynamically choose an appropriate initial price, we conduct an adaptive scheme to market operator. The market operator memorizes the clearing price and the demand bids of different cases, and uses the clearing price of the closest case as initial price every time the balancing procedure starts. But it is inefficient to record the whole demand bids condition. Therefore, the market operator only memorizes the clearing price and the total demand as the training data. We examine the effect of the adaptive scheme in case 1 comparison with that at the fixed initial price 20.

Table 8 shows the resulted bidding stages to reaches the equilibrium point for the adaptive scheme in comparison with that of fixed initial price scheme. As shown in Table 8, the equilibrium rate of adaptive initial price scheme is faster than fixed initial price scheme. Since the demand distribution can be different even if total demand is the same it will still result in different clearing prices, we keep the most recent clearing price as the initial price. This total demand adaptive scheme is still more efficient than the fixed initial price scheme.

Table 8. Results of adaptive initial price tests

Initial price	static	Adaptive
Bidding stages	59	53

4.3.2 Price adjustor factor

The price adjustor factor α is a variable that controls the rate of price variation as we mentioned in section 3.2. It is another factor that affects the market price of next bidding stage besides the current market price.

Table 9. Results of static price adjustor tests

α	0.05	0.01	0.005	0.001
Bidding stages	∞	35	58	147

Again, to examine the influence of the price adjustor factor, we propose case 3 as the test system. We investigate the bidding process under four different static price adjustor factor conditions: 0.05, 0.01, 0.005 and 0.001. Table 9 is the results of each condition.

Table 9 shows that the market reaches equilibrium point sooner if α is larger, but it may result in infinite bidding circumstance if α is too large. Since it is inefficient when α is too small and may be endless when α is too large, it is important to find an appropriate α for every instance. Although in the above case, the market bidding curve which α=0.01 seems to be efficient, it may also result in a divergent bidding curve in other cases. Thus, we may improve the efficiency of the system if we can adjust the value of α according to the bidding circumstance.

By observing the previous market bidding curves, we notice that in a convergent market bidding curve, the market price of next bidding stage is usually closer to clearing price than previous bidding stage, and the amplitude of the curve will become smaller as the balancing procedure progress. Therefore, it may be better to decrease α gradually until reaching the equilibrium point.

We examine the α-decreasing mechanism in two different methods: the periodically α-decreasing method and the dynamically α-decreasing method. In periodically α-decreasing method, the value of α is decreased each fixed number of stages. For example, we halfway down the value of α each 5 stage and each 10 stage for testing.

As we mentioned before, the amplitude will become smaller as the balancing procedure progress in a convergent curve. Thus, if the amplitude of the market bidding curve does not become smaller, we have to deduce the rate of price variation. The price variation rate is affected by the price adjustor factor α and the difference between power offer and power requirement. Since the market operator can only control the value of α, the value of α is decreased if the power offer-requirement difference does not become smaller in the dynamically α-decreasing method.

The results of two periodically α-decreasing testing and dynamically α-decreasing testing are listed in Table 10 with the result of static experiment as comparison.

Table 10. Results of α-decreasing tests

α	Bidding stages
Static 0.01	35
Periodically 5	138
Periodically 10	25
Dynamically	23

Table 8 shows that if we decrease α every 5-stage, the equilibrium rate is slowest since α decrease to 0.001 at the early stage. In the other hand, if we decrease α every 10-stage, the resulting equilibrium rate is faster than fix α in 0.01. And the equilibrium rate of adjusting α depending on the difference of power offer and power requirement is the fastest.

5. DISCUSSION AND CONCLUSION

Above simulation is running on a personal computer with CPU Intel 1.73GHz and 2GB main memory under the windows operating system. Since even under identical system complexity, the number of bidding stages is still different if demand bids are different, we measure the system efficiency by the process time per bidding stage instead of total process time. Averagely, it takes 2.76 (Sec.) to per bidding stage. In this process period, bidding procedure takes 0.12 (Sec.), power losses and loss factor calculation takes 0.1(Sec.), remaining 2.54(Sec.) are all wasting on communication between JAVA and MATLAB platforms. Therefore, if we can find another way to efficiently combine these two platforms together, the process efficiency will be greatly increased.

In this thesis, we have presented a multi-agent system to simulate the bidding behaviors and the interactions of participants in a retailed electricity market. The locational marinal pricing model we used to determine the optimal dispatch schedule of the generation output of suppliers in the market makes the system more cost-effective and power-effective. The numerical results support the analytical conclusion very well. Moreover, since we adapt the dynamic market structure, our system can handle changing conditions of the market in real time.

We have studied the influence of initial price and price adjustor factor α in our system bidding process. We proposed several experiments to show how to choose and adjust the value of initial price and price adjustor factor. The rate to reaches the equilibrium point will be fastest if we using the adaptive initial price and the dynamical adjust of a α-decreasing mechanism.

To make our system more powerful and efficient, the negotiation mechanism would be an interesting issue in future work. If we allow suppliers and LSEs negotiate directly with each other about the power distribution, it may speed up the equilibrium rate of the market but it needs to also enhance the complicated bidding strategies of market players.

REFERENCE

[1] Green R. and Newbery D. M. (1992): 'Competition in the British Electricity Spot Market', Journal of Political Economy, 100, 929–953.

[2] Green R. (1999) 'The Electricity Contract Market in England and Wales', J. Ind. Econ., vol. 47, no. 1, pp. 107–123.

[3] Borenstein S. and Bushnell J. (1999): 'An Empirical Analysis of the Potential for Market Power in California's Electricity Industry', The Journal of Industrial Economics 473, 285–323.

[4] Green R. (1996): 'Increasing Competition in the British Electricity Spot Market', J. Ind. Econom., vol. 44, no. 2, pp. 201–216.

[5] Bushnell J.B. and Wolack F.A. (2000): 'Regulation and the Leverage of Local Market Power in the California Electricity Market', Competition Policy Center. Paper 00-013.

[6] Liu C.C., Bose A. and Cardell J. (2008): 'Agent Modeling for Integrated Power System, Power and Fuel Market Simulation', PSERC Project M-13 Status Report.

[7] Richter C.W. Jr., Sheblé G.B. and Ashlock D. (1999): 'Comprehensive Bidding Strategies with Genetic Programming/Finite State Automata', IEEE Trans. Power Syst., 14, pp. 1207–1212.

[8] Song H., Liu C.C., Lawarrée J. and Dahlgren R.W. (2000): 'Optimal Electricity Supply Bidding by Markov Decision Process', IEEE Trans. Power Syst., 15, pp. 618‑624.

[9] Zhang D., Wang Y., and Luh P.B.(2000): 'Optimization Based Bidding Strategies in the Deregulated Market', IEEE Trans. Power Syst., 15, pp. 981‑986.

[10] Ferrero R.W., Rivera J.F., and Shahidehpour S.M. (1998): 'Application of Games with Incomplete Information for Pricing Electricity in Deregulated Power Pools', IEEE Trans. Power Syst., 13, pp. 184‑189.

[11] Park J.B., Kim B. H., Jung M.H. and Park J.K. (2001): 'A Continuous Strategy Game for Power Transactions Analysis in Competitive Electricity Markets', IEEE Trans. Power Syst., 16, pp. 847‑855.

[12] Song H., Liu C.C., and Lawarrée J. (2002): 'Nash Equilibrium Bidding Strategies in a Bilateral Electricity Market', IEEE Trans. Power Syst., 17, pp. 73‑79.

[13] Bai X., Shahidehpour S. M., Ramesh V. C. and YU E. (1997): 'Transmission Analysis by Nash Game Method', IEEE Trans. Power Syst., 12, 1046–1052.

[14] Klemperer P.D. and Meyer M.A. (1989): 'Supply Function Equilibria in Oligopoly under Uncertainty', Econometrica, vol. 57, pp. 1243–1277, 1989.

[15] Baldick R, Grant R., and Kahn E.(2000): 'Linear Supply Function Equilibrium: Generalizations, Application, and Limitation', POWER Working paper.

[16] Rudkevich A. (2003): 'Supply Function Equilibrium: Theory and Applications', Proceedings of the 36th Hawaii International Conference on Systems Science.

[17] Niu H., Baldick R. and Zhu G.(1859-1867): 'Supply Function Equilibrium Bidding Strategies with Fixed Forward Contracts', IEEE Trans. on Power Systems, vol. 20, no. 4, pp.

[18] Baldick R., Grant R., and Kahn E.(2004): 'Theory and Application of Linear Supply Function Equilibrium in Electricity Markets', Journal of Regulatory Economics, 25(2):143-167.

[19] Anderson E.J. and Xu H.(2005): 'Supply Function Equilibrium in Electricity Spot Markets with Contracts and Price Caps', Journal of Optimization Theory and Applications, 124, (2), 257-283.

[20] FERC(2003), Notice of White Paper, U.S. Federal Energy Regulatory Commission.

[21] Li H., Sun J., and Tesfatsion L.(2009): 'Separation and Volatility of Locational Marginal Prices in Restructured Wholesale Power Markets', ISU Economics Working Paper #09009.

[22] Litvinov E., Zheng T., Rosenwald G., and Shamsollahi P.(2004): 'Marginal Loss Modeling in LMP Calculation', IEEE Trans. on Power Systems, vol. 19, no. 2, pp. 880-888.

[23] Fu Y. and Li Z.(2006): 'Different Models and Properties on LMP Calculations', Power Engineering Society General Meeting, 2006. IEEE.

[24] Cardell J. B.(2007): 'Marginal Loss Pricing for Hours with Transmission Congestion', IEEE Transactions on Power Systems, Vol. 22, No. 4

[25] Conejo A.J., Castillo E., Minguez R. and Milano F.(2005): 'Locational Marginal Price Sensitivities', IEEE Transactions on Power Systems, vol. 20, no. 4, pp. 2026–2033.

[26] von der Fehr N. H. M. and Harbord D.(1993): 'Spot Market Competition in the UK Electricity Industry.', The Economic Journal, 103, 531-546.

[27] Java Agent Development Framework (JADE). Available at: http://jade.cselt.it/.

[28] Bellifemine F., Caire G., and Greenwood D.(2007): 'Developing Multi-agent Systems with JADE', John Wiley & Sons.

[29] Yao J., Adler I., and Oren S.(2008): 'Modeling and Computing Two-settlement Oligopolistic Equilibrium in a Congested Electricity Network', Operations Research, vol. 56, pp. 34-47, 2008.

[30] Zimmerman R.D., Murillo-Sánchez C.E., and Gan D. 'MATPOWER:A MATLAB Power System Simulation Package', CornellUniversity. Ithaca, NY. [Online]. Available: http://www.pserc.cornell.edu/matpower.

Towards Optimal Planning for Distributed Coordination Under Uncertainty in Energy Domains

Jun-young Kwak, Pradeep Varakantham*, Milind Tambe, Laura Klein, Farrokh Jazizadeh, Geoffrey Kavulya, Burcin B. Gerber, David J. Gerber

University of Southern California, Los Angeles, CA, 90089

*Singapore Management University, Singapore, 178902

{junyounk,tambe,lauraakl,jazizade,kavulya,becerik,dgerber}@usc.edu,

*pradeepv@smu.edu.sg

ABSTRACT

Recent years have seen a rise of interest in the deployment of multiagent systems in energy domains that inherently have uncertain and dynamic environments with limited resources. In such domains, the key challenge is to minimize the energy consumption while satisfying the comfort level of occupants in the buildings under uncertainty (regarding agent negotiation actions). As human agents begin to interact with complex building systems as a collaborative team, it becomes crucial that the resulting multiagent teams reason about coordination under such uncertainty to optimize multiple metrics, which have not been systematically considered in previous literature. This paper presents a novel multiagent system based on distributed coordination reasoning under uncertainty for sustainability called SAVES. There are three key ideas in SAVES: (i) it explicitly considers uncertainty while reasoning about coordination in a distributed manner relying on MDPs; (ii) human behaviors and their occupancy preferences are incorporated into planning and modeled as part of the system; and (iii) the influence of various control strategies for multiagent teams is evaluated on an existing university building as the practical research testbed with actual energy consumption data. We empirically show the preliminary results that our intelligent control strategies substantially reduce the overall energy consumption in the actual simulation testbed compared to the existing control means while achieving comparable average satisfaction level of occupants.

Categories and Subject Descriptors

I.2.11 [**ARTIFICIAL INTELLIGENCE**]: Distributed Artificial Intelligence

General Terms

Algorithms, Human Factors

Keywords

Sustainability, Multi-Objective Optimization, Energy, Satisfaction, Multiagent Systems

1. INTRODUCTION

Over the decades, energy issues have been getting more important. In the U.S., about 40% of energy consumption is from buildings (shown in Figure 1), of which 25% is associated with heating and cooling [21] at an annual cost of $40 billion [21]. Furthermore, on an annual basis, buildings in the United States consume 73% of its electricity. Recent developments in multiagent systems are opening up the possibility of deploying multiagent teams to achieve complex goals in such energy domains that inherently have uncertain and dynamic environments with limited resources.

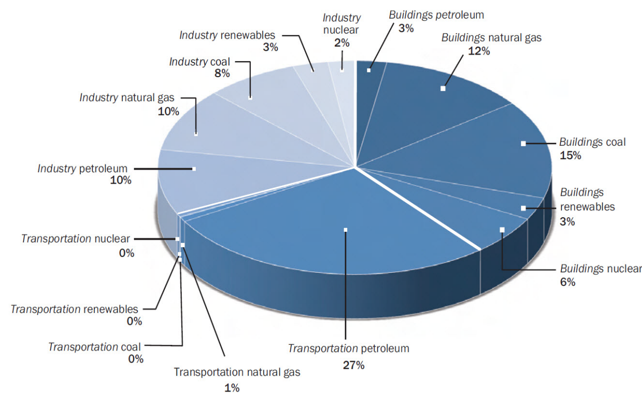

Figure 1: Distribution of US energy use in 2006, grouped by end-use sector (transportation, buildings, industry). Annual consumption for 2007 was 101.6 quads (10^{15} BTU)

This paper focuses on a novel planning method for distributed coordination under uncertainty (regarding agent negotiation actions) to optimize multiple competitive objectives: i) amount of energy used in the buildings; ii) occupant's comfort level; and iii) practical usage considerations. There have been some trials to balance energy consumption and enhancement of building services and comfort levels [15, 19, 23, 25] and to monitor and collect energy consumption data [15, 16] in energy domains. Other works have explicitly focused on design optimization and use of multiagent systems [13, 17, 20] in different domains. In addition, some multiagent systems [5, 6, 8, 9, 10, 11, 22] and the underlying theory for their decision supports [12] have been employed to model home automation systems. Unfortunately, past work in the energy domain has three key weaknesses. First, they do not consider uncertainty while reasoning about coordination and mostly rely on deterministic plans. Second, they limitedly incorporate intelligence of occupancy or occupancy preferences into the system and thus occupants are not explicitly modeled as agents in the system. Third, their works are mostly evaluated in their own simulation environments, which are not constructed on the actual energy data and occupants' responses in the buildings. Thus, their assumptions may not be realized in real-world problems.

This paper presents a novel multiagent system based on distributed coordination reasoning under uncertainty for sustainability called SAVES (Sustainable multi-Agent systems for optimiz-

Figure 2: **Testbed - Educational Building at USC**

ing Variable objectives including Energy and Satisfaction). SAVES provides three key contributions to overcome limitations in past work. First, we explicitly consider uncertainty while reasoning about coordination in a distributed manner. In particular, we rely on MDPs (Markov Decision Problems) to model agent interactions, specifically focusing on rescheduling meetings, which will be extended to decentralized MDPs. Second, human behaviors and their occupancy preferences are incorporated into planning and modeled as part of the system. As a result, SAVES is capable of generating an optimal plan not only for building usage but also for occupants. Third, the influence of various control strategies for multiagent teams is evaluated on an existing university building as the practical research testbed with actual energy consumption data in the simulation. Since the simulation environment is based on actual data, this result can be easily deployed into the real-world. Preliminary results show that our intelligent control strategies substantially reduce the overall energy consumption in the actual simulation testbed compared to the existing control means while achieving comparable average satisfaction level of occupants.

2. MOTIVATING DOMAINS

This work is motivated by energy domains where multiagent coordination can be the key issue. To pin down the domain problem, we consider an actual educational building (shown in Figure 2) as a representative test case to measure and collect the energy consumption and responses of occupants because it is a multi-functional building of sufficient size and activity for research.[1] Furthermore, the building is representative in that it has been designed with a building management system, and it provides a good environment to test various control strategies to mitigate energy consumption. The research can be easily generalized to other building types, where we can observe many different types of energy-use awareness based on the behavioral patterns of occupants in the buildings.

Our research testbed is focused on testing different operation optimization strategies based on the scope of occupant behaviors and schedules. The simulation component will include the building, its human occupants, and its facility management. It will then interact with the occupants and management via proxy agents [24] to advise them on how to reduce energy use while measuring occupant comfort level. More specifically, human occupants are divided into two main categories — permanent and temporary building occupants. Permanent building occupants include office resident such as faculty, staff, and researchers and laboratory residents like researchers in web labs, structural labs, etc. Temporary building occupants have scheduled occupants who include students or faculties attending classes or meetings and unscheduled occupants who are students or faculty using common lounge or dining space. In this domain, proposed human energy behaviors are entering/leaving a room, turning on/off light sources, turning on/off computers and other electronics, adjusting thermostat (heating or cooling), adjusting window shading, opening an operable door or window, adjusting personal clothing, and adjusting activity level. Building components and equipments that are another type of agent in the buildings include HVAC systems, which are composed of air handler units, VAV boxes, temperature sensors, and thermostats, lighting systems, office electronic devices such as computer and AV equipments, and laboratory equipments. Measurement of energy consumption for each equipment action may be estimated from design specifications. In our work, we choose and implement a subset of agents and their energy-related behaviors listed above.

[1]The size and other parameters of the building are given in the evaluation section.

3. RELATED WORK

With rising energy costs, the need to design and integrate scalable energy consumption reduction strategies in buildings calls for novel approaches. There are numerous challenges associated with energy resources such as supply and depletion of energy resources and heavy environmental impacts [19] (ozone layer depletion, global warming, climate change, etc.). The rise in energy consumption in buildings can be attributed to several factors such as enhancement of building services and comfort levels [15, 19, 23, 25], through heating, cooling and lighting needs and increased time spent indoors [19].

To model and optimize buildings' energy consumption, building owners and facility managers are demanding robust, intelligent and adaptable performance monitoring techniques. These techniques are important in energy consumption data collection [15, 16] and ambient environmental conditions control [15]. Existing heating, cooling, ventilation, and lighting systems generally operate with no intelligence of occupancy or occupancy preferences and therefore are unable to optimize operations. Even more, no feedback is available to occupants about how their actions and schedules impact building energy consumption. To realize both tangible benefits such as energy and operation savings, value property, reduction in occupant complaints as well as the intangible benefits such as occupant comfort and satisfaction, designers must develop energy adaptive capabilities within the building environmental control systems.

Abras *et al.* [5], Conte *et al.* [9] and Roy *et al.* [22] have employed multiagent systems to model home automation systems (or smart homes) and simulating control algorithms to evaluate performance. While there is relevance in terms of the problem domain and employing multiagent systems, our representation and approaches are different in having to account for human preferences and decisions directly.

Research by Fahrioglu *et al.* [14], Mohsenian-Rad *et al.* [18] and Caron *et al.* [7] provide incentive compatible mechanisms for distribution of energy among interested parties. This thread of research is complementary, especially in designing incentives for humans to reveal their true energy preferences. However, these approaches assume a centralized controller with whom all the members interact, which is not present in our domain. Instead, there are peer-to-peer negotiations between humans regarding their energy consumption and comfort level.

Figure 3: Overall System Design

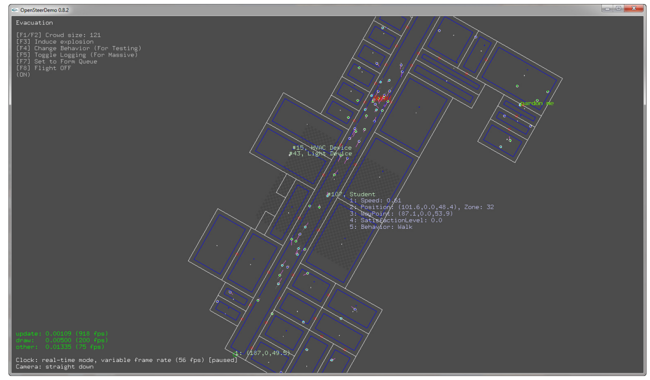

(a)

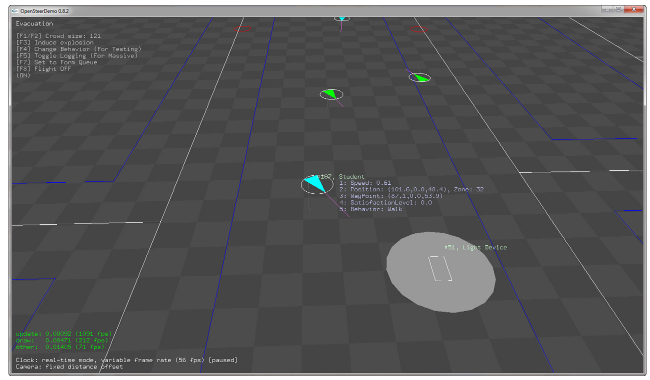

(b)

Figure 4: Screen capture

4. DESIGN DECISIONS

The SAVES system consists of a simulation module, an input/output module to communicate with agents, and an underlying reasoning and planning module. Figure 3 shows a generic loop of the system. In particular, the input/output module first collects data and constructs the world model. Given the world model, the reasoning and planning module generates policies to achieve the given objectives in the context of coordination. With these world model and generated policies, the simulation module models agents' physical and behavioral interactions in the system and realize the coordination in the actual world via the input/output module. We now describe the modules as well as the particular instantiations of these modules in the energy domain.

The simulation module provides a 2D, OpenGL environment based on the open-source project OpenSteer [2] as shown in Figure 4(a) & 4(b). The simulation module consists of two different types of agents as described below, modeling their physical and behavioral interactions. It can be used for efficient statistical analysis of different control strategies in buildings before deploying the system to an actual physical world.

The input/output module makes a connection among different modules in the system by collecting actual data in the domain, transferring data to the reasoning module, sending output results to either the simulation or deployed module in the world to represent outputs, and providing means to communicate with agents via proxy and handheld devices.

The coordination and planning module generates optimal policies to achieve the given team missions considering multiagent interactions in the context of coordination in the mutiagent setup.

Here we describe the design issues regarding agents, first introducing building component agents and human agents, then detailing the method to calculate the properties of agents, and finally discussing different control strategies considering agent interactions.

4.1 Building Component Agents

We consider three building component agent categories: a HVAC (Heating, Ventilating, and Air Conditioning) agent, a lighting agent, and an appliance agent. The HVAC agent (Figure 5(a)) is modeled based on the principles of thermodynamics, fluid mechanics, and heat transfer. We assume that this agent mainly controls the temperature of the assigned zone. The lighting agent (Figure 5(b)) controls the lighting level of the room. For the appliance agent, we only include the computer device including the desktop and laptop computers in this work. These agents have two possible actions: "on" [2] and "standby". When the lighting or appliance agents are "on", they consume some fixed amount of energy. We measure the average amount of energy used by these agents, which will be detailed in Section 6.

Since the energy consumption of HVAC agents relies on a set of parameters including the temperature change in the space, air flow, and number of people, and etc., the average value cannot be simply measured. Instead, we describe how to compute the energy use by HVAC agents below.

[2]Note that the "on" action can be divided into several actions with different output levels, e.g., "on" with 30%, 60%, and 100% power.

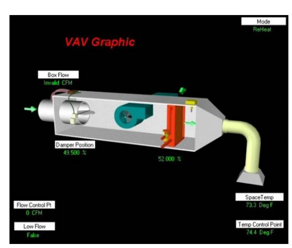

(a) HVAC Agent

(b) Lighting Agent

Figure 5: Agents

Calculating Total Energy Consumption: Since the building is composed of a large number of HVACs and they are the main consumers of the energy, it is important to choose the right set of parameters and reasonable values for them. In particular, the energy consumption of HVAC agents is calculated as following [1, 4] mainly based on changes in air temperature and airflow speeds:

$$Q = \frac{1.1 \times CFM \times \Delta T}{3412.3},$$

$$\Delta T = log(\frac{CFM}{C}),$$

where Q is the amount of energy used (kWh), CFM is an air volume flow rate (ft^3/min), which is typically ranged between 500–1500 (ft^3/min), ΔT is the temperature change in a zone (°F), and C is a scale factor.

4.2 Human Agents

There are four different types of human agents such as a faculty, staff, graduate student, and undergraduate student. Each agent has access to a subset of the six available behaviors according to their types — wander, attend the class, go to the meeting, teach, study, and perform research, any one of which may be active at a given time, where the behavior is selected via the given class and meeting schedules.

During execution of these behaviors, individual travelers may move at integer speeds from 0 to 3. Each agent also has specific levels of emotions and information about the environment. Specifically, every agent has a property about the satisfaction level based on the current environmental condition and knows his or her current location without any noise. A more extended discussion of the satisfaction property will take place below.

Calculating Satisfaction Level: The satisfaction level (SL) of an individual human agent is modeled as a percentage value between 0 and 100 (0 is fully unsatisfied, 100 is fully satisfied). SL of the individual occupant is calculated as following:

$$SL = 100.0 - PPD,$$

$$PPD = 100.0 - 95.0 \cdot \exp^{-(0.03353 \cdot PMV^4 + 0.2179 \cdot PMV^2)},$$

where SL is the satisfaction level (%), PPD is the Predicted Percent Dissatisfied (%), PMV is the Predicted Mean Vote. The PMV index is calculated from an equation of thermal balance for the human body in ASHRAE Standard [3], involving the parameter values shown in Table 1.

The PMV model uses heat balance principles to relate the seven key factors for thermal comfort listed in Table 1 to the average response of people on the above scale. The PPD index is calculated using the PMV as defined in [3]. It is based on the assumption that people respond about their comfort level with a number between -3 and +3 (-3 is cold, +3 is hot and 0 is neutral) on the thermal sensation scale and that the simplification that PPD is symmetric around a neutral PMV.

Table 1: Parameters for the Satisfaction Level

Parameter	Value Range
Clothing	0.5 – 1.0 (light to heavy clothing)
Metabolic Rate	1.0 – 2.0 (low to high activity)
External Work	0
Air Temperature	20 – 28 (°C)
Radiant Temperature	20 – 28 (°C)
Air Velocity	0 – 0.2 m/s
Relative Humidity	30 – 60 %

5. CONTROL STRATEGIES

In a given scenario, all agents within the simulation will use the same strategy. Possible strategies include: i) manual control strategy, ii) reactive control strategy, iii) proactive control strategy, and iv) proactive control strategy based on multiagent coordination.

5.1 Manual Control

The manual control strategy simulates the current building control strategy maintained by USC facility managers. Specifically, we assume that HVAC agents are not controlled by human agents and that appropriate temperature points are centrally set/given by facility managers. For HVAC agents, the CFM values are fixed throughout the simulation. In this control setting, HVAC agents always try to reach the pre-set temperature using the fixed CFM value regardless of the presence of human agents in the specific space and their preferences in terms of temperature. Lighting agents are controlled by only human agents. Control actions (i.e., turning on/off the light) of human agents are either deterministic or stochastic according to the type of action. In particular, when human agents enter the space, they always turn on the light. When they leave the space, they stochastically turn off the light. For appliance agents, we simply assume that they are always on.

5.2 Reactive Control

Since the manual control strategy simply follows the pre-defined policy provided by the facility managers, it is fairly easy to come up with action plans of building component agents. However, it does not adapt the given policies based on actual schedules or preferences of occupants in the building, and thus the building component agents are limited to adapt their control policies appropriately according to the dynamic changes. Particularly, HVAC agents keep operating to reach the desired point, even though the space is empty, which ends-up wasting energy. At the same time, since they do not consider occupants' preferences in the space and instead prioritize the pre-determined points, the average satisfaction level of occupants can decrease.

Here we discuss about another control strategy that building component agents reactively respond to the behaviors of human agents. In this setting, we assume that HVAC agents are not controlled by human agents and that appropriate temperature points are measured based on the average preference of human agents in the specific space. HVAC agents automatically turn on and off according to the presence of people and temperature set points, and the CFM values are adjusted based on the desired temperature point. In the reactive control strategy, the lighting and appliance agents are now automatically controlled. In particular, they are turned on and off according to the presence of people. For instance, when people enter the specific room, the lighting and appliance agents are automatically turned on, and when people leave the room, they are turned off.

While human agents follow their given schedules, with the reactive setting, the building component agents can act more intelligently than the manual policy as they operate based on human agents' actual needs. As a result, we can reduce the cases where the energy is wasted for unnecessary spaces, which will contribute to the reduction of the overall energy consumption.

5.3 Proactive Control

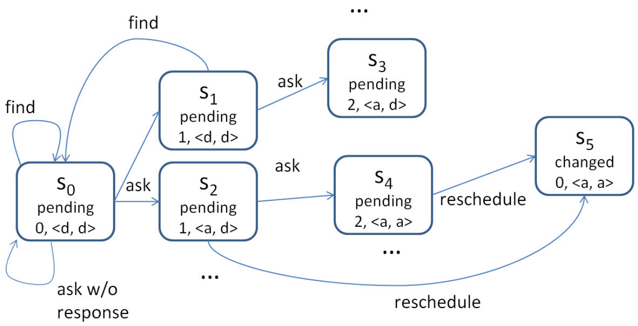

Figure 6: Simplified MDP model — d: disagree, a: agree

Although the reactive control strategy can adapt their policies based on actual needs of occupants in the building, this approach is still limited in a sense of optimality. In practice, there is a delayed effect in changing temperature. In other words, HVAC agents can only change a certain amount of degree in temperature per hour. This property exposes the weakness of the reactive control strategy. Although HVAC agents know the desired temperature of human agents at a specific time point, it takes a certain amount of time to reach the desired temperature point from the current air temperature, and the satisfaction level of occupants in the space will decrease during that time.

To overcome limitations of the reactive setting, we suggest a third control strategy operated in a proactive manner. Given the meeting and class schedules of human agents, the building component agents can predict: i) what their preferences are in terms of temperature, ii) how long it will take to reach the preferred temperature point from the current air temperature, iii) what CFM value is required, etc. In this setting, the building component agents can access the meeting/class schedules of human agents. Based on that prior knowledge, they now generate more optimal policies to reduce the overall energy consumption while maximizing the average satisfaction level of occupants in the building. For instance, a HVAC agent knows that the current air temperature is 55°F and the preferred temperature of the group of human agents who will use the space in 2.5 hours is 70°F. If the maximum possible temperature change by HVAC is 10°F/hr with the maximum CFM value, the HVAC agent predicts that it needs to change the temperate by 6°F per hour with a smaller CFM value which will use less energy. With this proactive plan, when human agents get to the space, the air temperature is already 70°F, which is the desired temperature point of people, and thus their satisfaction level increases.

5.4 Modeling Multiagent Coordination: MDP representation

With the existence of human agents, agent interactions are a fundamental aspect of our energy simulation. In SAVES , all agents share a common architecture based on MDP (Markov Decision Problem) frameworks, possessing varying degrees of knowledge about the world and other building agents (i.e., local knowledge).

MDPs have been used to tackle such real-world multiagent planning and coordination problems under transition uncertainty, which are described by a tuple $\langle S, A, T, R\rangle$, where

- $S = \{s_1, ..., s_k\}$ is a finite set of world states.
- A is the finite set of actions of an agent.
- $T : S \times A \times S \mapsto \mathbb{R}$ is the transition probability function, where $T(s'|s, a)$ denotes the transition probability from s to s' if an action a is executed.
- $R : S \times A \times S \mapsto \mathbb{R}$ is the reward function, where $R(s, a, s')$ denotes the reward that an agent gets by taking a from s and reaching s'.

We denote a policy computed by MDP $\pi : S \mapsto A$ is a mapping from world state to action. Our goal is effective multiagent team coordinations to minimize the total energy consumption while maximizing occupant's comfort level.

This section describes our MDP representation in the energy domain for illustration. The MDP model represents a class of MDPs covering all types of meetings for which the agent may take rescheduling actions. In our work, we construct a MDP for each meeting as shown in Figure 6. Alternatively, we can model all meetings in the building as a single MDP. However, if we consider a gigantic MDP model for rescheduling all meetings together, the number of states and actions exponentially explodes as the number of agents increases. In addition, the complexity to handle all possible coordinations among agents significantly increases, which is burdensome to handle within a reasonable amount of time.

As preliminary work, we construct a simplified MDP model for rescheduling meetings. For each meeting, a meeting agent can perform any of three actions — reschedule, find another slot [3], and ask. For the "ask" action [4], an agent can autonomously reduce its own autonomy and ask a human agent whether he or she agrees with rescheduling the existing meeting. The human agent can respond to the meeting agent with "agree" or "disagree".

The agent may choose to perform any of these actions in various states of the world. State is composed of three features: $\langle f_1, f_2, f_3\rangle$, where f_1 is the status whether meeting location and time is changed (i.e., pending or changed), f_2 is the number of "ask" actions invoked so far, and f_3 is a set of responses from all meeting attendees: $\langle rp_{i,1}, rp_{i,2}, ..., rp_{i,n}\rangle$, where n is the number of attendees of meeting i and $rp_{i,k}$ is a response of agent k to rescheduling meeting i (i.e., agree or disagree).

The MDP's reward function has its maximum value when the meeting agent invokes the "reschedule" action in the state where all meeting attendees agreed to reschedule. We denote the component of the reward function that focuses on the expected energy gain by rescheduling the meeting as r_{energy}. However, there is clearly a high team cost incurred by forcing all of the attendees to rearrange their schedules. This team cost is incorporated into the MDP's reward function by adding a negative reward, $r_{rearrange}$. The magnitude is also an increasing function in the number of attendees (e.g., rescheduling a meeting of a large group is more costly than rescheduling a one-on-one meeting). The overall reward function for taking the "reschedule" action, $a_{reschedule}$, in a state s is a weighted sum of these components:

$$R(a_{reschedule}, s) = \alpha \cdot r_{energy} + (1 - \alpha) \cdot r_{rearrange}$$

, where $0 \leq \alpha \leq 1$. In addition, a small amount of cost is incurred to invoke actions of "ask" and "find another slot".

The MDP's transition probabilities represent the likelihood over possible action outcomes. Specifically, the transition function is defined considering four factors: i) meeting constraints of attendees; ii) level of energy consciousness, which determines how much they care about energy; iii) degree of intimacy among occupants; and iv)

[3] This action can be modeled differently, e.g., delay 1 hr, delay 2 hrs, ..., delay n hrs, delay 1 day, ..., delay n days, cancel the meeting, change location (but same time), etc.

[4] The ask action can be divided into several actions with different amount of incentives, e.g., ask with 5% incentive, ask with 10% incentive, etc.

Table 2: **Parameter Values for the Experiments**

	Temperature	CFM	Likelihood
Manual	65–70°F	1500.0	50%
Reactive	Preference	500.0–1500.0	Automatic
Proactive	Preference	500.0–1500.0	Automatic
Proactive w/ MDP	Preference	500.0–1500.0	Automatic

Table 3: Parameter Values for the SL Calculation

Parameter	Value
Clothing	1.0
Metabolic Rate	1.2
External Work	0
Air Temperature	Zone temperature
Radiant Temperature	65°F
Air Velocity	0.1 m/s
Relative Humidity	40 %

the current status of responses, which can be related to emotional contagion within the group. Since we store the current set of responses from individual agents and number of "ask" actions called so far, the repeated "ask" action may result in different transitions. In particular, the "ask" action, by which the agent queries the human agent, has $2^{n_i} + 1$ possible outcomes, where n is the number of attendees of the meeting i. First, the human agent may not respond at all, in which case, the agent is performing the equivalent of a "wait" action for a given timeout. Other set of possible outcomes are decided depending on responses of meeting attendees as illustrated in Figure 6. We assume that the "find" action reset values of features in the state.

One possible policy, generated for a subclass of possible meetings, specifies "ask" and then "wait" in state S_0 of Figure 6, which prompts the agent to give up its autonomy. If the agent then reaches state S_1, the policy specifies "find", so the meeting agent figures out another available location or time for rescheduling. However, if the agent then reaches state S_2, the policy again chooses "ask", which asks the human agents once more to collect their responses. Similarly, if the agent reaches S_4, the "reschedule" action is chosen according to the policy.

Based on this MDP model, the agent reasons about different tradeoffs in team costs. This reasoning follows a fundamental tenet of teamwork in our system, that the individual team members act responsibly towards the team.

6. EMPIRICAL VALIDATION

We evaluate the performance of SAVES in our energy domain

Table 4: Average Energy Consumption of Lighting & Appliance Agents

Agent Type	Category	On	Off/Standby
Lighting	Office	0.128 kW/hr	0 kW/hr
	Conference room	0.192 kW/hr	0 kW/hr
	Classroom	0.768 kW/hr	0 kW/hr
Appliance	Desktop	0.150 kW/hr	0.010 kW/hr
	Laptop	0.050 kW/hr	0.005 kW/hr

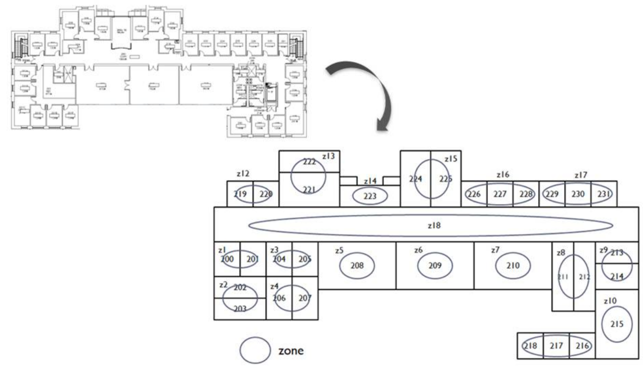

Figure 7: **Floor Plan - Educational Building at USC**

and compare four different control techniques: 1) manual control, 2) reactive control, 3) proactive control, and 4) proactive control with MDP. We focus on measuring two different criteria — total energy consumption (kWh) and average satisfaction level of occupants (%). The parameter values used in the experiments are shown in Table 2. In Table 2, column 2 shows the desired temperature for HVAC agents (Note: Preference in rows 3–5 means that the desired temperature is decided based on the average preference values of building occupants) and column 4 displays the likelihood value for the "turn off" action for the lighting agent. To calculate the energy consumption by the HVAC agent, we set the scale factor to 100.0. For the satisfaction calculation, we used the same parameter values in Table 3 while performing the experiments across four different control strategies. The experiments were run on Intel Core2 Quad-core 2.4GHz CPU with 3GB main memory. All techniques were evaluated for 100 independent trials throughout this section. We report the average values.

6.1 Experimental Domain Description

We have identified an educational building in conjunction with USC Facilities Management Services, as our practical testbed. This campus building is composed of classrooms, offices for faculty and staff, and conference rooms for meetings. Specifically, we use one floor of the actual university building in the experiments, which has 18 zones and 33 rooms as illustrated in Figure 7. There is one HVAC agent for each zone, and one lighting agent for each room. We also assume that each person in the office has either one desktop or laptop computer, and conference room and class room has two computers, respectively. There are four human agent categories: faculty, staff, graduate student and undergraduate student. Throughout the entire simulation, we consider a typical winter season in southern California (i.e., starting indoor temperature is 55°F in the simulation). During the simulation, indoor temperature goes down by -1°F per timestep, where each time step is 30 minutes. Possible temperature range in the building is between 50 and 90°F. Students follow 2010 Fall class schedule, and we generated the arbitrary meeting schedules for faculty, staff, and student agents. The measurement is performed during a working hour (i.e., 8:00am – 7:00pm), and the preference value of each occupant in temperature is randomly drawn from the uniform distribution between 60 – 70°F. To calculate the energy consumption of the lighting and appliance agents, we collect actual energy consumption data in the testbed building and used the average values shown in Table 4.

6.2 Comparison: Total Energy Consumption

We compared the cumulative total energy consumptions mea-

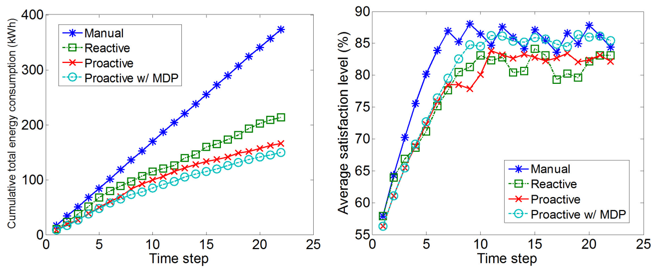

(a) Total Energy Consumption (b) Average Satisfaction Level

Figure 8: Comparison

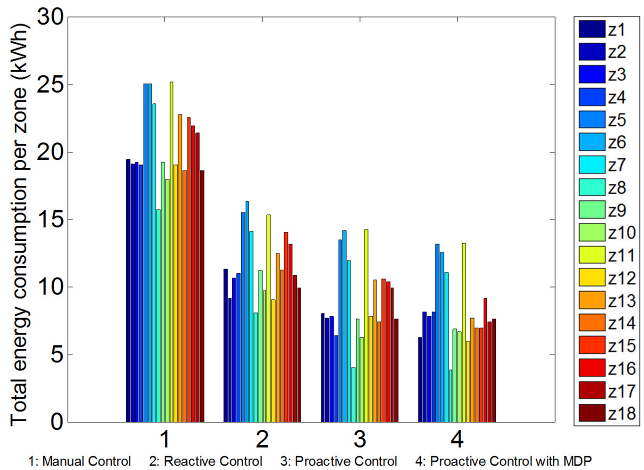

Figure 9: Energy Consumption Distribution

sured during work hours for all control strategies in the energy domain. Figure 8(a) shows the cumulative total energy consumption on the y-axis in *kWh* and the time step on the x-axis. Time step 1 indicates 8:00am and each time step increases by 30 minutes. As shown in the figure, the manual control strategy showed the worst result since it does not take into account behaviors or schedules of human agents and building component agents simply follow the predefined policies. The reactive and proactive control strategies showed lower energy consumptions than the manual setting by 43.0% and 55.6%, respectively. The proactive control strategy with the MDP model showed the best results among all different control strategies and statistically significant improvements (via t-tests) in terms of energy used in the testbed building, relative to other control strategies. Specifically, the proactive control with MDP reduced the energy consumption by 59.9% than the manual control strategy.

Although we did not tune the parameter values and only applied the simplified MDP model, with considering multiagent coordination in SAVES, we could achieve significant improvements. These outcomes are still preliminary results and yet only tested in the simulation environment, all experimental results were measured based on the actual data and testbed. Later, we will be able to show even more improvement with the optimally tuned parameters and extend our work to deploy it into the actual building with proxy agents. Furthermore, as we revise the equations shown in Section 4.1, we will be able to get more exact results for analysis.

Now, we analyze how various control strategies can cause different results. Figure 9 shows the energy consumption distribution over zones for all control strategies. In the figures, the x-axis shows the group number of data obtained by each control strategy and the y-axis displays the total energy consumption for each zone in *kWh*. The floor plan we used in the simulation has four different types of zones, which decides the total energy consumptions. Specifically, zones 1–4 (blue), 9 (green), and 12 (yellow) have two offices per zone, zones 5–7 (light blue or cyan) are class rooms, zones 13–15 (orange or red) are conference rooms, and zones 11 (yellow), 16 (light red), and 17 (red) have three offices per zone. As shown in the first group of Figure 9, the manual control strategy results in the similar level of energy consumptions according to the different types of zones. This result clearly indicates that the manual setting is only impacted by the physical constraints of the building space itself, which never considers the interactions among agents. The normalized standard deviation was 0.134. In the reactive (the second group in Figure 9) and proactive setting (the third group in Figure 9), it now started showing the difference in terms of the amount of energy used even within the same type of zones since those methods consider the actual behavioral patterns and schedules of human agents, and building component agents respond and adapt their policies based on them. Their normalized standard deviations are 0.205 and 0.312, respectively, which are higher than the value of the manual setting. Lastly, the proactive control strategy with the MDP model considers rescheduling of meetings. The target meetings to reschedule are ones with less than 4 people in the conference rooms in zones 13–15. We only considered the location reallocation and did not assume the meeting time can be also changed. New candidate locations are small faculty offices in zones 1–4. As shown in the fourth group of Figure 9, it showed increased energy consumptions in zones 1–4 due to the reallocated meetings, but simultaneously showed much more reduction in zones 13–15, as a result the overall energy consumption decreased. The normalized standard deviation was 0.313, which was the highest among different control strategies. These results give us a lesson that multiagent coordination/negotiation can benefit our model in SAVES, and by considering higher degree of coordination among agents, we will be able to achieve the significant energy reduction in this domain.

6.3 Comparison: Average Satisfaction Level

Here, we compare the average satisfaction level of human agents under different control strategies in the simulation. We used the equations discussed in Section 4.2.

Figure 8(b) shows the average satisfaction level in percentage on the y-axis and time step on the x-axis, which are the same as mentioned in the previous section.[5] As shown in the results, all methods were able to achieve at least 80% or higher results on average, and the manual and proactive with MDP settings showed the best results among them. Note that the equations to calculate the individual satisfaction level are based on the average model about the responses according to different environmental conditions, which is mostly related to air temperature, and they do not consider individual preferences. Thus, although the reactive and proactive control strategies act more intelligently by additionally considering the preferences of human occupants, we could not obtain explicit benefits to improve the satisfaction level and even in some cases, the solution quality may be harmed. On the other hand, the manual setting just make HVAC agents attempt to reach the desired temperature set point over time. Once HVAC agents get to the desired point, they are turned off, which will decrease the satisfaction level. If the tem-

[5] Note that the starting indoor temperature of the building is 55°F in the simulation, which causes the low average satisfaction level for a while.

perature is again away from the scope of desired temperature point, HVAC agents are turned on and the satisfaction level increases. As a result, the manual setting shows a race condition in the graph, which means it eventually cannot go over a certain point in terms of the satisfaction level. With revised equations considering more factors from the coordination perspective such as preferences, energy awareness, emotional contagion effect, etc., we expect more significant improvements in terms of the satisfaction level.

In our work, we still only separately consider two different optimization criteria — the energy consumption and the satisfaction level since this is still preliminary work. However, as we will eventually optimize multiple objectives in SAVES, we will be able to achieve effective multiagent team coordinations to minimize the total energy consumption while maximizing occupant's comfort level.

7. CONCLUSION

This paper aims to open a new area of research for multiagent systems: in many real-world problems, specifically in energy domains, we see many different levels of agent interactions and coordinations involved, and hence multiagent systems must address such complex situations to achieve the given objectives under uncertainty. In this work, we presented a new framework called SAVES based on distributed coordination reasoning for sustainability. There are three major new ideas in SAVES. SAVES: (i) explicitly considers uncertainty while reasoning about coordination in a distributed manner relying on MDPs; (ii) incorporates human behaviors and their occupancy preferences into planning and models them as part of the system; and (iii) evaluates various control strategies for multiagent teams on an existing university building as the practical research testbed with actual energy consumption data. We justified our design decisions in SAVES through a preliminary empirical evaluation and showed that SAVES can provide solutions to significantly reduce the energy consumption while achieving the comparable satisfaction level of building occupants. For future work, we will consider opportunities for direct occupant participation and incentivization via handheld devices and deploy our system to the real-world.

8. REFERENCES

[1] Method of testing for rating room air conditioner and packaged terminal air conditioner heating capacity. *ANSI/ASHRAE Standard 58*.

[2] Steering behaviors for autonomous characters. Website. `http://opensteer.sourceforge.net/`.

[3] Thermal environmental conditions for human occupancy. *ANSI/ASHRAE Standard 55*.

[4] Ventilation for acceptable indoor air quality. *ANSI/ASHRAE Standard 62*.

[5] S. Abras, S. Ploix, S. Pesty, and M. Jacomino. A multi–agent home automation system for power management. In *ICINCO*, 2006.

[6] V. Callaghan, G. Clarke, A. Pounds-Cornish, and S. Sharples. Buildings as intelligent autonomous systems : A model for integrating personal and building agents. In *International Conference on Intelligent Autonomous Systems (IAS)*, 2000.

[7] S. Caron and G. Kesidis. Incentive-based energy consumption scheduling algorithms for the smart grid. In *SmartGridComm*, 2010.

[8] W.-H. Chen and W.-S. Tseng. A novel multi-agent framework for the design of home automation. In *ITNG*, 2007.

[9] G. Conte and D. Scaradozzi. Viewing home automation systems as multiple agents systems. In *Multi-agent system for industrial and service robotics applications, RoboCUP*, 2003.

[10] D. Cook, M. Youngblood, and S. Das. A multi-agent approach to controlling a smart environment. *Designing Smart Homes*, 4008:165–182, 2006.

[11] D. J. Cook. Multi-agent smart environments. *Journal of Ambient Intelligence and Smart Environments*, 1:51–55, 2009.

[12] M. Dibley, H. Li, J. Miles, and Y. Rezgui. Towards intelligent agent based software for building related decision support. *Advanced Engineering Informatics*, 2010.

[13] J. Dijkstra and H. Timmermans. Towards a multi-agent model for visualizing simulated user behavior to support the assessment of design performance. *Automation in Construction*, 11(2):135 – 145, 2002.

[14] M. Fahrioglu and F. L. Alvardo. Designing incentive compatible contracts for effective demand managements. *IEEE Trans. Power Systems*, 15:1255–1260, 2000.

[15] Y. Gao, E. Tumwesigye, B. Cahill, and K. Menzel. Using data mining in optimization of building energy consumption and thermal comfort management. In *International Conference on Software Engineering and Data Mining (SEDM)*, 2010.

[16] Q. Khalid and R. Langhe. Evaluation and monitoring of energy consumption patterns using statistical modeling and simulation: Case study of commercial and residential buildings. In *International Conference on Emerging Technologies (ICET)*, 2010.

[17] K. Kim and K. J. Kim. Multi-agent-based simulation system for construction operations with congested flows. *Automation in Construction*, 19(7):867 – 874, 2010.

[18] A.-H. Mohsenian-Rad, V. Wong, J. Jatskevich, and R. Schober. Optimal and autonomous incentivebased energy consumption scheduling algorithm for smart grid. In *ISGT*, 2010.

[19] L. Perez-Lombard, J. Ortiz, and C. Pout. A review on buildings energy consumption information. *Energy and Buildings*, 40:394–398, 2008.

[20] Z. Ren, F. Yang, N. Bouchlaghem, and C. Anumba. Multi-disciplinary collaborative building design–a comparative study between multi-agent systems and multi-disciplinary optimisation approaches. *Automation in Construction*, 2011.

[21] A. H. Rosenfeld, J. J. Romm, H. Akbari, and A. C. Lloyd. Painting the town white – and green. *Technol. Rev.*, 100:52–59, 1997.

[22] N. Roy, A. Roy, and S. K. Das. Context-aware resource management in multi-inhabitant smart homes: A nash h-learning based approach. In *PerCom*, 2006.

[23] M. Santamouris, A. Argiriou, E. Dascalaki, C. Balaras, and A. Gaglia. Energy characteristics and savings potential in office buildings. *Solar Energy*, 52:59–66, 1994.

[24] P. Scerri, D. Pynadath, L. Johnson, P. Rosenbloom, M. Si, N. Schurr, and M. Tambe. A prototype infrastructure for distributed robot-agent-person teams. In *AAMAS*, 2003.

[25] H. S. Sun and S. E. Lee. Case study of data centers' energy performance. *Energy and Buildings*, 38:522–533, 2006.

On the integration of power and reserve trade in electricity networks

Nicolas Höning[1], Han Noot[1] and Han La Poutré[1,2]
[1]CWI, Science Park 123, Amsterdam, The Netherlands
[2]Utrecht University, Princetonplein 5, Utrecht, The Netherlands
{nicolas,han.noot,hlp}@cwi.nl

ABSTRACT

Demand and (with the advent of renewable energy) supply in electricity networks are not perfectly predictable, but still need to match exactly. Many modern power markets trade reserve capacity a day ahead, next to a traditional power exchange market, with the goal to lower the price volatility caused by this uncertainty. Several technical solutions which can tackle uncertainty in the electricity markets, like Demand Response or storage systems, are expected to increase the trade volume of reserve power significantly. We argue that this prospect increases the need for a tight integration of power and reserve markets, which neither real-world implementations nor computer science research has provided so far. This integration poses an important design challenge in the field of multi-unit auctions. We propose a novel bidding format to achieve this integration and provide a clearing mechanism, which can ensure the most efficient power generation and transmission. We then put forward how to extend the well-known agent-based electricity simulation framework AMES with our model. Furthermore, it is still unclear how the participants should be engaged in the difficult process of hedging, especially in the light of the ongoing, worldwide liberalisation of power markets. The service to keep reserve power available produces opportunity costs, which should be part of bidding strategies.

Categories and Subject Descriptors

I.6 [**Computing Methodologies**]: Simulation And Modeling—*Applications, Distributed*

General Terms

Management, Design, Measurement, Economics, Reliability, Experimentation

Keywords

Energy and Emissions, Simulation, Electronic Markets, Artificial economies/markets, Public policy and Economics

1. INTRODUCTION

Our power systems are on the verge of a technology change, moving away from fossil fuels. During this transition, the markets are increasingly faced with uncertainty of intermittent supply. At the same time, power markets throughout the world are being deregulated. This can help to put prices on uncertainty, but can also lead to more volatile prices.

On the wholesale level, a power market consists of traders representing supply (generators) and demand (load-serving entities (LSE)). In addition, we need to consider a transmission grid and its operator, who often functions as the market maker. More and more power markets around the world are designed as centralised auctions [13], attempting to cover efficiency aspects better than monopolised or bilateral markets. The currently most popular design is to conduct an ahead-market, followed by a real-time spot market. The ahead-market is often accompanied by a reserve market, in which optional amounts of power are traded (an option being the right, but not the obligation, to exercise a deal [7]). The market held in real-time then functions as a balancing market, in which the difference between the outcome of the day-ahead market and reality is settled by executing parts of the intervals sold in the reserve market.

Until recently, the flexibility of supply as well as demand has exclusively been negotiated in bilateral long-term contracts and was as such already a significant means to price and allocate reserve power. These "ancillary services" incurred almost the same cost impacts as transmission in 1998 in the US (12 billion versus 15 billion per year) [6]. The trade volume of reserve power is expected to grow, as we are now faced with decreasing certainty of supply caused by the advent of intermittent generation, i.e. renewables like solar and wind: The requirements for reserve power are expected to increase proportionally with growing wind power penetration [17]. The trade of reserve power will become one of the major design issues in the centralised power exchanges of the next generation. A well-designed reserve market should not only allocate reserve power, but also minimise the costs which reserve power incurs through generation and transmission. It should enable traders to include the costs of lost opportunity in their bids for reserve power. With regard to the technology mix in future grids, the reserve market should also be tailored to proposed solutions for the uncertainty problem: price-sensitive demand and storage systems can make the market clearing more flexible and decrease price volatility.

The challenge is that neither power nor reserve markets are trivial to design efficiently, let alone their integration. Current approaches by system operators over the world vary in their approaches, especially in the question whether power and reserve markets should be bundled and if so, how to achieve a satisfactory integration. While the prices in most power markets can be minimised using increasing marginal bid functions, the reserve markets are often simplified and accept only constant prices as bids. This will lead to power markets becoming less efficient as the trade volume of reserve power becomes more significant and some traders, such as storage systems, almost exclusively deal with reserve power. Research in the computer science community (that we know of) has as of yet not paid attention to the design and integration of reserve markets. Both research and policy making would benefit much from well-defined and well-studied reserve markets that are able to allocate reserves efficiently among traders, employing realistic settings and scenarios.

This paper explores this new research opportunity. Its main contribution is the proposal of a novel bid format and market clearing mechanism for integrated power- and reserve markets. This enables the use of one unified market, which expresses dynamic marginal prices for both power and reserves. It makes reasoning over expected opportunity costs possible for generators and enables the integration of transmission pricing for executed reserves. Because multi-unit markets, especially in combination with options, are very complex, we use agent-based computational economics (ACE) [16] to model our approach.

We proceed as follows: After Section 2 reviews basic power and reserve market setups and background literature, Section 3 discusses design issues in the integration of traditional power markets and reserve markets. We propose a joint bid format in Section 4. Section 5 then formalises a computational mechanism to clear bids to the day-ahead market, minimising generation costs during both the day-ahead and balancing phase. Within this mechanism, Section 6 discusses a strategy space for the generators to express opportunity costs. In Section 7, we show how this mechanism can be implemented as an extension within the AMES power market simulation framework [14] (thereby incorporating pricing for transmission as well as for reserves) and sketch design variables for experiments. As an outlook, Section 8 sketches how competition can be increased by allowing bids for reserves from both supply and demand, for upward regulation as well as downward regulation. Finally, Section 9 concludes and reviews future work.

2. BACKGROUND

2.1 Power markets

Even in the light of deregulation, power markets are highly regulated: the government is heavily involved in securing steady supply and consumer-friendly price properties. This is why most deregulated markets and the studies that model them work with a centralised model [13]. These models are multi-unit auctions, in which goods (units of power) are exchangeable. The System Operator (SO)[1] is the market maker of a market M populated by two types of players: generators $g \in G$ and load-serving entities (LSEs) $l \in L$. He collects their bids and computes a marginal clearing price p_M.

Power markets are also subject to a given distribution structure: The distribution has to be adhere to the constraints of a hierarchical, sometimes meshed, grid and obey the laws of physics, e.g. Kirchhoff's laws (always matching supply and demand).

To model the costs of generation per amount, both (piecewise) linear or quadratic functions can be used (generation cost functions are non-decreasing, since cheaper generation is selected first). Marginal cost functions are stepwise in the former, and linear in the latter case. Though the issue is not settled, most of the current research in power markets uses quadratic cost functions, as they can be described with few parameters, provide a good trade-off between approximation of real generation costs and good performance in calculus-based clearing algorithms. They also produce a smoother state transition behaviour in the market [2]. We will use quadratic cost modelling, similar to that used in in AMES: $C(P_g) = aP_g + bP_g^2$, where P_g is the amount of power provided by g. Hence, $a + 2bP_g$ is the marginal cost function.

LSEs are commonly equipped with no or very low demand price elasticity. With the advent of Demand Response and storage systems though, this can only be an early modelling assumption for any market model.

The usual optimisation goal is to minimise the generation costs $\sum_g^G P_g p_M$. This is done so that in a market where generators bid close to their true costs (or at least uniformly far from it), lowering total generation costs increases generation efficiency and thus uses less energy input.

2.2 Reserve markets

Reserve markets trade reserve capacity during the day-ahead phase, which is used as a backup for the real-time balancing market. In this work, we look at the case of generators bidding on hourly intervals of reserve capacity, which the SO can "execute" to compensate for unexpected demand in real time[2].

Reserve markets should compensate generators for their (lost) opportunity costs of withholding reserve capacity. Since the computation of opportunity costs by the SO is non-trivial [4], most approaches (see [17]) let generators incorporate them into their bids, either by using one-time *availability costs* (\$/MW), or *activation costs*, only pricing the units of reserve capacity that were actually activated in real time (\$/MWh). Bidding on activation costs requires bidders to estimate the expected execution of reserve capacity.

It is apparent that until now no design has emerged which gives the correct economical incentives. For instance, the Dutch reserve market puts no price on the opportunity costs [1]. Instead, activation prices are used, being averages of bids or determined marginally (where most authors agree on the latter being more economically sound). In addition, the Dutch SO sometimes adds incentives and/or penalties in order to induce a desired bidding behaviour of generators. In the New England (USA) market, both opportunity and availability costs have been used [4], while the recent implementation of the Midwestern market uses only availability costs and declares opportunity costs 'implicit' [11]. It has been shown that under several current pricing regimes, strong incentives may exist to withhold capacity from balancing markets and balance privately owned assets only [17].

2.3 Agent-based market models

There has been recent research interest to model the energy markets by means of multi-agent systems. Most models (e.g. the AMES package [14] and the Trading Agents approach [18]) address two-stage auctions which trade energy

[1]In practice often called Independent System Operator (ISO) or Transmission System Operator (TSO)

[2]Note that in this work, no switching costs between hours are modelled.

for every hour of one day, with a day-ahead market and a real-time balancing stage. Often, the prices include the implications of transmission pricing. These aspects are also core features of the 2003 proposal [3] of the Federal Energy Regulation Commission (FERC) in the US and are based on established research into spot pricing of energy [12].

While the pricing of ancillary services has been under discussion among power system experts (e.g. [9, 19, 17]), conclusive modelling of a market for them has (to the authors best knowledge) not been undertaken. The aforementioned agent-based approaches have not had reserve pricing in the focus of their research. While the AMES package has yet to implement a balancing phase, the Trading Agents approach simply falls back to the unmatched day-ahead bids when balancing. Neither of those approaches uses an explicit approach for pricing reserves.

3. DESIGN ISSUES IN THE INTEGRATION OF POWER AND RESERVE MARKETS

Before we elaborate on our reserve mechanism design and implementation, this section considers the coexistence of traditional power markets with reserve markets, with particular regard to two design issues: Assuming we prefer to hold power and reserve auctions simultaneously (versus sequentially), should we bundle the two markets into one? How can we achieve higher bid expressiveness in order to reach near-equilibrium states quicker?

When a generator g, having an upper capacity limit of power P_g^U, bids both on the power market and the reserve market, he sells a default amount P_g^{def} and an optional interval P_g^{opt}, from which any amount $P_g^{exe} \in [0, P_g^{opt}]$ may be executed by the SO in the balancing phase (of course, $P_g^{max} = P_g^{def} + P_g^{opt} \leq P_g^U$). Thus, the generator offers two goods (power of the amounts P_g^{def} and P_g^{exe}, which are decided upon in day-ahead and balancing phase, respectively) and one service (keeping up to P_g^{opt} in reserve).

3.1 Operation of power and reserve markets

To explain how supply prices in power and reserve markets would develop, we now describe simplified day-ahead markets for power M_p and reserves M_r, each with the same set of players (the generators $g \in G$ and LSEs $l \in L$) and both cleared by the same SO. In this simplified model, the bid of a generator is a pair of amount and unit price: $b_{g,p} = (P_g^{def}, \gamma_{g,p})$ in M_p and $b_{g,r} = (P_g^{opt}, \gamma_{g,r})$ in M_r.

For every hour, there exists demand D_{M_p} for supply bids in M_p and demand $D_{M_r,opt}$ for reserves in M_r. In real time, we must consider a demand $D_{M_r,exe}$ for balancing. The SO now produces a schedule from the generator bids, which means he assigns each generator values for P_g^{def}, P_g^{opt} and (later, during balancing) P_g^{exe} with $P_g^{exe} \in [0, P_g^{opt}]$, conditional on D_{M_p}, $D_{M_r,opt}$ in the day-ahead stage for power and reserves, respectively, and $D_{M_r,exe}$ in the real-time balancing phase. The SO clears M_p in the day-ahead phase at the marginal clearing price γ_p and clears M_r later in the balancing phase at the marginal clearing price γ_r. In this example, we use only activation costs, so the actual payment from M_r comes about only through execution of reserves during the real-time phase - it amounts to $P_g^{exe} * \gamma_r$. This means that γ_r needs to compensate g for any opportunity costs which have been incurred by withholding P_g^{opt} from market M_p. The main optimisation goal of the SO is to minimise generation costs (see Section 2.1), which in our two-market case are $\sum_g^G P_g^{def}\gamma_p + P_g^{exe}\gamma_r$. Of course, the market clearing also needs to satisfy demand with $D_{M_p} \leq \sum_g^G P_g^{def}$, $D_{M_r,exe} \leq D_{M_r,opt} \leq \sum_g^G P_g^{opt}$.

In this model, a discussion of equilibria is defined by cost functions of generators and by probability distributions over D_{M_p}, $D_{M_r,opt}$ and $D_{M_r,exe}$. We can assume generation cost functions to be stable, so the means of these distributions define winning bids and their standard deviations define the stability of equilibria. Furthermore, because $D_{M_r,exe} \leq D_{M_r,opt}$, we can assume that $\gamma_p \leq \gamma_r$, since $\gamma_r - \gamma_p$ needs to compensate generators for opportunity costs.

For now, we confine ourselves here to a one-sided market with static demand and assume no transmission constraints. It is an added level of complexity for the SO to satisfy transmission grid constraints (thermal constraints of branches and enforcing Kirchhoff's laws), considering all amounts of power, i.e. $\sum_g^G P_g^{def}$ and $\sum_g^G P_g^{exe}$.

3.2 The case for simultaneous markets and bundled bidding

How should power and reserve markets operate next to one another? Many implementations (e.g. in the Netherlands [15]) operate sequentially, trading reserve capacity after the day-ahead auction for power has finished. However, most scientific literature favours simultaneous bidding over sequential bidding, as the re-commitments in later auctions can lead to inefficient allocations [6]. In a perfect competition, sequential bidding should lead to socially efficient procurement, but in imperfect competitions and in particular in the special case of power and reserve market stages, generators can (and have been observed to) show strategic behaviour, such as holding back bids in earlier stages [9]. It also makes sense to hold the two auctions for $\sum_g^G P_g^{def}$ and $\sum_g^G P_g^{opt}$ simultaneously if we consider that they share the same constraint for each generator: $P_g^{def} + P_g^{opt} \leq P_g^U$. Each unit of power can only be offered in one market or the other. If g decides to bid less $(P_g^{def} - x)$ in M_p, he can immediately bid more $(P_g^{opt} + x)$ in M_r, and vice versa. Furthermore, each generator needs to recover its costs of operation, determined by the cost function $C_g(P)$. Imagine that the bid $b_{g,p} = (P_g^{def}, \gamma_{g,p})$ is accepted first. Then g faces the risk that $C_g(P_g^{def} + P_g^{exe}) - P_g^{def} * \gamma_p > P_g^{exe} * \gamma_r$, i.e. even if g knew P_g^{exe} during the day-ahead phase, he might have to accept a bid $b_{g,r}$ that fails to recover the remainder of his overall costs $C_g(P_g^{def} + P_g^{exe})$.

If we merge $b_{g,p}$ and $b_{g,r}$ into a unified bid, which is accepted or rejected in its entirety, and thus turning two market clearing processes into one, the only contingencies generators need to consider during the cost allocation problem now exist *within* each bid. The cost allocation problem is still dependent on the actual demand in the balancing stage ($D_{M_r,exe}$) but not on contingencies between bids.

Such bundling has until recently seldom been proposed (examples of exceptions from the Electrical Engineering community are [19] and [8]). Interestingly, the first market (to our best knowledge) in which generators bid for supply of power and reserves in a bundled fashion is in operation since 2009 in the US Midwest region [11]. In addition to technical specifications of power plants, a bid of a generator in this market consists of a piecewise linear cost function for the power market and a constant unit price for all availability costs of reserve energy, the amount of which is scheduled by the SO. The scheduling of both power and reserves is co-optimised (but the effectiveness of this mechanism has not been analysed or modelled in any computer science approach).

3.3 The case for high expressiveness of bids

We continue to use our example market formalisation of Section 3.1 and assume bundled bidding as proposed in Section 3.2. Generators are now using quadruples as bids (two amounts and two unit prices) and are faced with market dynamics which are dependent on three probability distributions of demand (see Section 3.1). If each generator submits only one bundled bid at a time, the time until the market reaches a near-equilibrium point can be expected to be quite long and the market thus to be less efficient than it could be.

To design an efficient market, the bidding process should include bids which express more alternatives per generator. We can achieve this in two ways: over time (via repeated bidding) or via the number of submitted bids per generator. While repeated bidding is a much-researched topic (e.g. Continuous Double Auctions), we prefer a one-shot auction, where all information is available to the mechanism at one moment. We mentioned in Section 3 that the integrated day-ahead market trades one good and one service and that these two commodities share constraints within each generator and on the network. A one-shot auction seems to offer more control over these complexities.

For submitting several bids per generator, there is a proven approach called Supply Function Equilibria [10], in which firms bid supply functions. A supply function as bid connects any amount to a price for the market to choose from. The form of bids pertains to the discussion of cost functions in Section 2.1. A supply function combines Cournot and Bertrand modelling approaches, and in addition allows an individual characteristic of costs or utility to be expressed in the curve shape, which is a useful property in power markets, where generation and consumption of power technologies increasingly display different cost and utility profiles. It is also especially useful in the light of uncertainty [10] and has since recently been considered often for traditional power market design (e.g. by the Federal Energy Regulation Commission in the US [3], which led to implementations in the US Midwest and New England markets). Certain properties like monotonicity and convexity are commonly assumed, in order to aid the clearing process.

Even in the advanced market design of the US Midwest region, supply bids for P_g^{opt} to the power market are simply constant unit prices, independent from both the amounts P_g^{opt} and P_g^{opt}, at which the generator is already scheduled to operate. Our goal is to express as much pricing information for P_g^{opt} in the reserve trade as for P_g^{def} in the power trade.

4. BUNDLED BID FUNCTIONS

In this section, we propose a novel bid format which enables the expressiveness of the Supply Function approach for bids on reserve power, while bundling power and reserve markets as proposed in Section 3.2. A bid consists of the function parameters - for instance, when using quadratic functions of the format $b(P) = aP + bP^2$, the bid consists of the parameters a and b. For each amount of power P, a bid can describe both the total price by $b(P)$ and the marginal price by $b'(P)$.

In traditional power auctions with supply functions, the SO (as the auctioneer) allocates P_g^{def} for every generator g by announcing a marginal clearing price. In a bundled power- and reserve market, we need to allocate both P_g^{def} and P_g^{opt} (we denote $P_g^{max} = P_g^{def} + P_g^{opt}$). To also express bidding for reserve capacity P_g^{opt} within these supply functions, we will now propose that in each bid, the ratio $r = P_g^{opt}/P_g^{max}$ is fixed, such that knowing P_g^{def} determines P_g^{opt}, since $P_g^{opt} = P_g^{def}\frac{r}{1-r}$. Thus, the reserve interval $[0, P_g^{opt}]$ is determined by the market clearing, which allows the generators to include opportunity costs in the function. For example, with $r = \frac{1}{3}$ we denote that P_g^{def} will certainly be sold and $[0, P_g^{opt}] = [0, \frac{1}{3}P_g^{def}/\frac{2}{3}]$ is the optional interval.

There is one interesting observation to make: We concluded in Section 3.1 that $\gamma_p \leq \gamma_r$. Then, if $P_g^{opt} > 0$, it is clear that, given an non-decreasing marginal cost function, a generator g should consider the interval $[0, P_g^{def}]$ for γ_p and the interval $[P_g^{def}, P_g^{max}]$ for γ_r, since the first interval will be cheaper to generate than the second.

The ratio r denotes flexibility, in the use case of this work the flexibility of supply to increase production. At $r = 0$, no flexibility is offered and the generator has full certainty how much he sells ($P_g^{def} = P_g^{max}, P_g^{opt} = 0$). At $r = 1$, everything is flexible and the SO will assume full flexibility over P_g^{exe} in the balancing phase ($P_g^{def} = 0, P_g^{opt} = P_g^{max}$). Which values for r are preferred by a generator g is influenced by his own true costs and his residual demand[3] from D_r and $D_{r,exe}$. He needs to learn the prices in both the power and reserve market and how he should bid for them.

Each generator g can submit bid functions for several possible option-ratios r. Each bid connects all possible amounts for $P_g^{def} + P_g^{exe} \in [0, P_g^U]$ to a total bid price for some r, which includes opportunity costs if $r > 0$. Thus, in our model, traditional bid functions denote $r = 0$, and in addition generators can include an offer for reserve capacity by using any $r \in [0, 1]$. From these functions, the SO can allocate both P_g^{def} and $P_g^{exe} \in [0, P_g^{opt}]$. The next section will propose a mechanism to do that.

5. THE MARKET MECHANISM

In this section we develop a mechanism to clear an integrated power and reserve market using bids as explained in section 4, both in the day-ahead and in the real-time balancing phase.

5.1 General outline

Each generator $g \in G$ reports his lower and upper generation limits P_g^L and P_g^U. In addition, he can submit several bid functions $b_{g,r}$ to the SO, each using a distinct r. Each LSE $l \in L$ submits only the requested fixed amount P_l^{def} and reserve amount P_l^{opt}.

During the day-ahead auction, our mechanism operates as a one-shot auction, for which the SO collects bids from generators and LSEs for every hour of the next day. From every generator g, the SO will choose one bid function b_{g,r_g}. He will announce a marginal market clearing price γ_{def}, which defines how much each unit in $\sum_g^G P_g^{def}$ will be paid for. Via

[3]A generators residual demand describes the portion of market demand that is not supplied by other generators.

γ_{def}, each generator can look up on b_{g,r_g} how much power P_g^{def} he is scheduled to supply and this also tells him how much reserve capacity P_g^{opt} he needs to keep available (see Section 4).

During the real-time phase for each hour of the following day, LSEs announce their balancing requirements $P_l^{exe} \in [0, P_l^{opt}]$, and the SO computes an optimal allocation to settle the balancing demand $\sum_l^L P_l^{exe}$. In order to allocate each g a value for $P_g^{exe} \in [0, P_g^{opt}]$ and to find a marginal market clearing price γ_{exe} for the balancing phase, the SO translates the winning bids from the day-ahead auction into another cost minimisation problem. Our mechanism does currently not aim at solving balancing demand for $P_l^{exe} > P_l^{opt}$.

Thus, each LSE will have to pay $P_l^{def} * \gamma_{def} + P_l^{exe} * \gamma_{exe}$ and each generator will earn $P_g^{def} * \gamma_{def} + P_g^{exe} * \gamma_{exe}$.

5.2 The price range for balancing power

One of the advantages of a reserve market is that it can reduce the volatility of the price of balancing power (in our case, this is γ_{exe}). Making use of the bundled bids for both markets as described in Section 4, the clearing mechanism is able to make certain guarantees about the range of possible values for γ_{exe}. Sections 5.3 and 5.4 will explain how this is achieved.

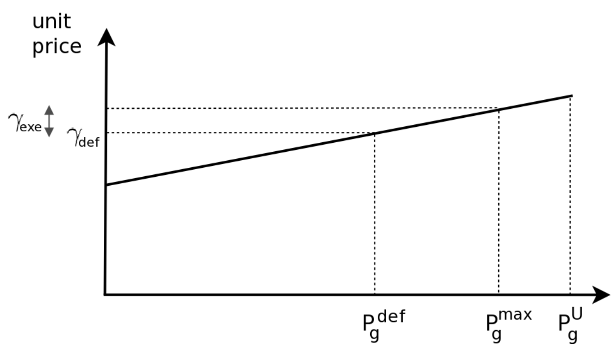

Figure 1: A marginal bid function with $r_g = 1/3$, denoting γ_{def} and γ_{exe}

Meanwhile, Figure 1 shows the simple case of one generator g. After the day-ahead scheduling, γ_{def} denotes the lower bound for γ_{exe}. Without P_g^{exe} being known yet, our mechanism ensures that $\gamma_{exe} \in [\gamma_{def}, b'_{g,r}(P_g^{max})]$, as $P_g^{exe} \leq P_g^{max} - P_g^{def}$. With more than one generator, the SO can announce the upper bound for γ_{exe} to be the highest possible price over all the successful bids.

5.3 Optimal dispatch in the day-ahead trade

We now formulate a Constraint Satisfaction Problem to clear the bundled day-ahead market. The optimisation goal of the SO is to minimise generation costs:

$$\min \sum_g^G P_g^{def} \gamma_{def} \tag{1}$$

Note that the generation costs for P_g^{def} can implicitly include opportunity costs for reserve power, as far as values of $r_g > 0$ have been among the winning bids (see Section 4). Furthermore, we of course require that demand is satisfied:

$$\sum_g^G P_g^{def} = \sum_l^L P_l^{def} \tag{2}$$

Secondly, the SO needs to make sure that each generator will stay within his generation limits:

$$P_g^L \leq P_g^{def} < P_g^U(1 - r_g) \tag{3}$$

Finally, as a result of the scheduling, the reserve capacity P_g^{opt} which each generator will hold back needs to match the reserve capacity that was requested by the LSEs. Hence, we add the constraint:

$$\sum_{g \in G} P_g^{opt} \geq \sum_{l \in L} P_l^{opt} \tag{4}$$

We could have used $=$ instead of $\geq$, but this is not necessary, as the cost optimisation is ensured by minimising generator costs. We add some flexibility to the process this way, e.g. bid curves with $r_g > 0$ can also take part in the competition for demand with $P_l^{opt} = 0$.

In addition, as each generator g is allowed to bid x functions $b_{g,r} = a_{g,r}P + b_{g,r}P^2$ (each with a different r), the SO needs to make sure to use only one per generator (the winning bid b_{g,r_g}). The mechanism iterates over the N^x possible sets $\{b_{g_1,r_{g_1}}, b_{g_2,r_{g_2}}, ..., b_{g_N,r_{g_N}}\}$ for all generators in $\{g_1, g_2, ..., g_N\}$ and their respective bids, finding the best configuration among them for scheduling. Thus, x is a parameter for the SO to control. By restricting generators in the number of bids they can submit, the SO defines the trade-off between the time complexity of finding a solution and the freedom of the generators to submit as many bids for different r as they want. Note, however, that the day-ahead scheduling is done offline and not time-critical. Nonetheless, to efficiently allow for high values of x will be a consideration of future work.

5.4 Optimal dispatch during balancing

To minimise generation costs during the balancing phase, the SO runs another, simpler optimisation to find the most efficient allocation that can satisfy $\sum_g^G P_g^{exe} = \sum_l^L P_l^{exe}$ and make the guarantees we discussed in Section 5.2. For this additional optimisation, the SO submits one bid function b_g^{bal} on behalf of each generator g. To get b_g^{bal}, the SO translates the interval $[P_g^{def}, P_g^{max}]$ of each successful bid b_{g,r_g} from the day-ahead phase (compare Figure 1) into a new bid function in the interval $[0, P_g^{opt}]$. For example, we use the bid format $b_{g,r_g}(P) = aP + bP^2$ - the SO sets the unit price a in each b_g^{bal} to the partial derivative of b_{g,r_g} at the point P_g^{def}: $a = \frac{\partial b_{g,r_g}}{\partial P}(P_g^{def})$. Furthermore, the SO sets b in b_g^{bal} to the value of b in b_{g,r_g}. This new optimisation problem is quite a bit simpler than the day-ahead optimisation: Constraint 3 has been taken care of in the day-ahead market, so for this calculation r can be assumed to be 0. Constraint 4 can also be disregarded. The result is the clearing price γ_{exe}, which determines both how much the LSEs have to pay for each unit of balancing power and the allocated output of balancing power P_g^{exe} for each generator $g \in G$.

6. OPPORTUNITY COST ASSESSMENT STRATEGIES

When increasing r, the bid curves need to include more costs of (probably) lost opportunity. The starting point of any generator strategy is to calculate the expected opportunity costs via some function $\phi_g(r)$, the internals of which we disregard for now. Opportunity costs could be lost profits or simply the lost contributions of the unsold units to fixed costs.

We will begin by introducing two strategies for $r > 0$, Availability and Activation, corresponding to the two ways that have been used in reserve markets to model opportunity costs of generators in their bids (see Section 2.2). Let us assume that the generator uses $aP + bP^2$ as total price bid function for $r = 0$, and for the sake of simplicity we assume b to be very small (which is often the case), so a denotes the unit price. The Availability strategy simply shifts the bid function upwards by $\phi_g(r)$. It thus uses the bid function $b^{Av}_{g,r}(P) = aP + bP^2 + \phi_g(r)$. The Activation strategy instead increases the unit price a by some amount a', so we write $b^{Ac}_{g,r}(P) = (a+a')P + bP^2$. The value of a' is chosen such that over the interval of possible outcomes for $P^{exe}_g \in [0, P^{opt}_g]$, the expected total revenue equals that of the Availability strategy:

$$
\begin{aligned}
& b^{Ac}_{g,r}(P^{def}_g) + \int_{P^{def}_g}^{P^{max}_g} D(P) * b^{Ac}_{g,r}(P)\, dP \\
&= b^{Av}_{g,r}(P^{def}_g) + \int_{P^{def}_g}^{P^{max}_g} D(P) * b^{Av}_{g,r}(P)\, dP
\end{aligned} \tag{5}
$$

With $D(P)$, we denote the probability of $P^{exe}_g = P$, drawn from the probability distribution D over P^{exe}_g which the generator expects ($P^{exe}_g \sim D$). Note that only amounts $\geq P^{def}_g$ are possible outcomes of the balancing phase. Thus, the Activation strategy also automatically includes some availability payment, namely $b^{Ac}_{g,r}(P^{def}_g) - C_g(P^{def}_g)$.

In Figure 2, we exemplify both strategies for bid functions $b^{Av}_{g,r}$ and $b^{Ac}_{g,r}$ with $r = 0.5$. Note that we do not show the part between P^{max}_g and P^U_g that was not selected in the market clearing. In this example, g expects a uniform probability distribution over the demand for his reserve power in the balancing phase: $P^{exe}_g \sim N(0, 1)$.

With the Availability strategy, the generator carries the risk of underestimating P^{exe}_g and the demand side carries the risk of him overestimating it, while for the Activation strategy it is the other way around.

We now consider the full class of functions in between both strategies. A strategy in this mixed class consists of any value a'' and any function $\phi'_g(r)$, with $0 < a'' < a'$ and $0 < \phi'_g(r) < \phi_g(r)$, so we write $b^{Mix}_{g,r}(P) = (a+a'')P + bP^2 + \phi'_g(r)$. The value for a'' and the function $\phi'_g(r)$ are determined by matching the expected total revenue to the Availability strategy (using a probability distribution), similar to Equation 5. Seen graphically (see Figure 2), this class of mixed strategies occupies the space between the pure Availability and Activation bid functions.

7. THE ROAD TO EXPERIMENTS

This sections discusses our approach at implementing the proposed mechanism in AMES and setting up experiments. This is ongoing work.

7.1 Mechanism implementation in AMES

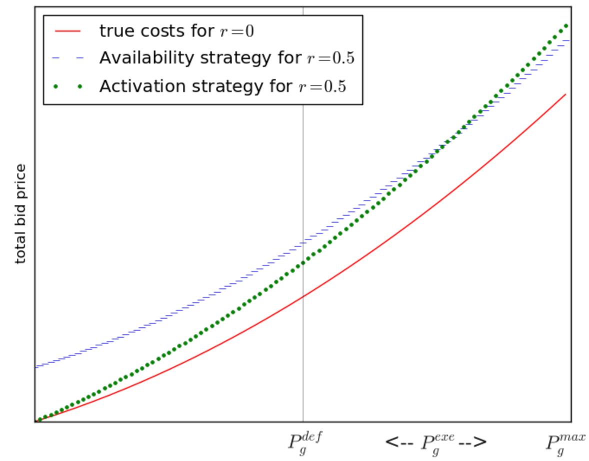

Figure 2: Bid curves for $r = 0.5$ with two pricing strategies for opportunity costs

Our day-ahead problem formalisation in Section 5.3 is a Strictly Convex Quadratic Programming problem [5] (SCQP). It optimises over quadratic functions, subject to constraints formalised as linear equalities and inequalities. We can implement this problem within AMES, because AMES also uses an SCQP. AMES also follows the main goal of minimising quadratic generation functions and of course implements Constraint 2 and (a simpler version of) Constraint 3. AMES translates an optimal power flow calculation using a direct current approximation of the grid (DCOPF) into a SCQP. The equality and inequality constraints within the DCOPF problem in AMES incorporate real physical constraints (e.g. in order to formulate Kirchhoff's laws) of the generators and the network.

The SO in AMES collects supply function bids from generators and (randomised) demand bids P^{def}_l from LSEs. Then, the SO clears the market using his SCQP formulation. The result is an optimal set of power commitments (amounts of P^{def}_g) and Locational Marginal Prices (LMP), which include transmission costs incurred by the schedule. The LMP mechanism (e.g. [12]) does not charge for transmission directly. It affects the selection of generators. Thus, the upper limit for γ_{exe} (see Section 5.2) still holds.

First, we alter the agent models: LSEs now also specify $P^{opt}_l \geq 0$ in the day-ahead phase and correspondingly $P^{exe}_l \in [0, P^{opt}_l]$ in the balancing phase. As for strategies of generators to include opportunity costs in their bids, the SCQP used in AMES assumes that bid functions have no constant term, which is why we are restricted to using the pure Activation strategy 6. The shape of the bid function can be adapted by changing a and b.

In order to incorporate our economical constraints into AMES, we simply add them within the inequality matrix of its SCQP. To incorporate Constraint 3, we simply replace the similar constraint AMES uses (which is $P^L_g \leq P^{def}_g \leq P^U_g$). For the extension with Constraint 4, we only need to extend the inequality matrix by a row.

As AMES has no implementation of a balancing phase, we implement this as a new task of the SO. For the optimisation problem which needs to be solved during the balancing phase, the transmission constraints of the grid are decreased by the overall amounts $\sum_g^G P_g^{def}$, which the cables are already scheduled to carry after the day-ahead scheduling.

7.2 Complexity of evaluation

Power grids and markets are complex systems. A simulation framework should capture the important factors of this complexity to enable realistic simulations. We believe AMES does this quite well. However, evaluating the results is not trivial. We have currently finished the implementation of our extensions to AMES and will soon design simulation experiments with the platform. To this end, it is important to know the influential factors in experiment design.

First, each experiment needs to specify a generator setup: The capacity constraints P^L and P^U as well as the parameters a and b for the true cost function. In addition, it is of course crucial to find reasonable economical strategies which the generators will use to control their bidding behaviours (adapting a and b and choosing which r to bid on). LSEs have fixed demand and (traditionally a lot less) price-sensitive demand, for which they also use a quadratic function ($cP + dP^2$). The difference between fixed demand and available supply is important for the price formation. Most crucial for the reserve aspect, however, is which values of P^{opt} and P^{exe} are requested by the LSEs. Currently, the LSE behaviour is non-strategic and serves as experiment parametrisation. Both P^{opt} and P^{exe} will vary over time, which can be described with statistical properties. The grid setup adds further constraints on the market outcome. Every power line has a maximum capacity and a reactance, in order to parametrise the DCOPF formalisation.

In this abstraction of power systems and markets, we combine Locational Marginal Pricing with an integrated trade of power and reserves. Both are complex problems on their own, and assuming self-interested traders, it is especially hard to minimise the costs of generation, while at the same time preventing attempts at gaming the system. For instance, a study using AMES, which implemented the FERC power market design proposal, successfully minimised generation costs, but found that implicit collusion among generators who adapted their offers by reinforcement learning increased prices substantially [14]. We plan to compare our market mechanism with the current setup of two day-ahead markets (for power and reserves) via simulations.

8. OUTLOOK: EFFICIENT BALANCING MARKETS

In this work, we have only paid attention to the use case of generators offering incremental reserves (i.e. the option to supply more power during balancing). This section sketches market clearing during the balancing phase, assuming our mechanism supported supply functions as well as demand functions for balancing power, incremental as well as decremental. We argue that such a balancing setup offers much room for efficiency gains through increased competition.

We first consult the terminology on financial derivatives, which is of great help in describing not only the possibility to draw more power from reserved capacity, as we have done so far, but also the possibilities to draw less, as well as to consume more or less: Options [7] are derivatives that enable the owner to exercise some right regarding the underlying asset. An option to buy the underlying asset is a "call"; an option to sell is a "put", both can be "long" (the option is lucrative when the value of the underlying asset increases) or "short" (if it decreases).

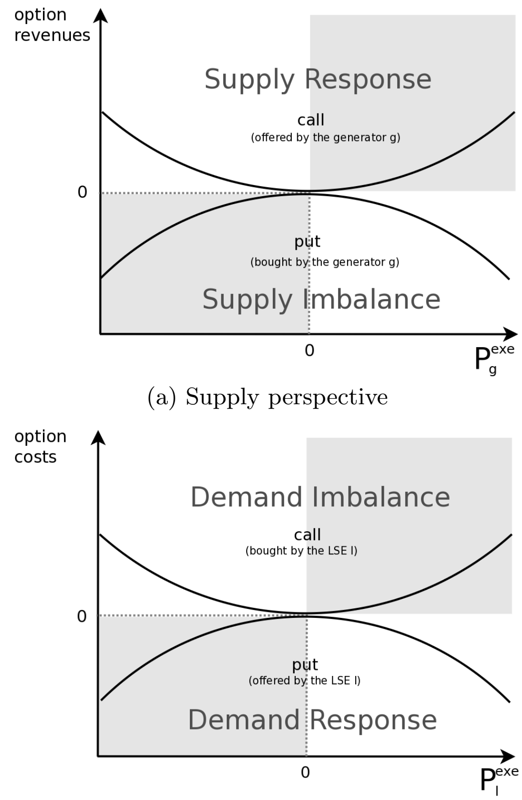

(a) Supply perspective

(b) Demand perspective

Figure 3: All possible call and put options for all amounts of P^{exe} on sale in the balancing market and exemplary revenue/cost valuations for them.

We separately look at the supply and demand perspective, where both can be in a state of imbalance (trader t having to exercise options on P_t^{exe}) or responding to such an imbalance (selling amounts of P_t^{exe}). Figure 3 illustrates a centrally cleared balancing market, in which all traders take part with parts of their successful day-ahead bids who have either bought or sold amounts of P_t^{opt} in the day-ahead market. Options can be incremental (selling/buying more) or decremental (selling/buying less). Options in grey areas denote long positions of the option holder, white ones denote short positions. Both supply and demand can increase their revenue in this market when their calls or puts, respectively, are accepted, and will encounter negative revenues when having to sell puts or calls, respectively.

The curve shapes are only exemplifying - it is possible for the SO to announce a price function for each 'Response' square, which minimises costs for executions of that option. The shape and length of these functions depend on the successful bid functions of the day-ahead market (see Section 5.4). In our proposed mechanism, the marginal clearing price for balancing power is determined by continuous functions and thus very price-sensitive. In addition, the mechanism can make some guarantees about maximal clearing prices for each 'Response' square (see Section 5.2).

As we will exemplify below, each imbalance can be responded to by both demand or supply. Thus, the cost functions in the 'Imbalance' squares depend on the price functions from the two 'Response' squares which can respond to them. This shows that a market design can increase competition by including Demand Response and incorporating Supply Imbalance, both of which will rise significantly with the introduction of renewable power supply.

In this work, we have only considered the case of generators offering long calls and LSEs bidding on short puts (upper right square of both Figure 3a and 3b). By extending on our concept of bundled bids, all the other use cases can be made possible and integrated. Because power on the grid is exchangeable, the market is able to look for the most efficient options to level out power flows. For example, imagine that instead of satisfying short puts of LSEs by long calls at generators, the SO could also execute long puts at other LSEs (Figure 3b, lower left square). In fact, long puts offered by LSEs are often discussed under the name of (decremental) Demand Response: lowering demand rewarded by price reductions. Decremental Demand Response can also sell long puts to intermittent generators who have to buy long put options (Figure 3a, lower left square), because they are uncertain about wind or sun supply. If, in yet another example, a generator produces more than he sold (Figure 3a, lower right square), e.g. when there was unexpectedly much wind for the wind mills, the SO could exercise some short puts at LSEs (Figure 3b, lower right square). The LSEs would now buy more power than planned, but at a discount on the marginal price. In such cases, storage technology will play an important role. Alternatively, the SO could also buy short calls at other generators (Figure 3a, upper left square). These generators would produce less on short notice, but be paid a higher marginal price for their output.

9. CONCLUSIONS AND FUTURE WORK

In power markets, the uncertainty of intermittent generation and the flexibility of storage systems present new market design challenges. In this work, we propose a novel bidding mechanism for the integration of power and reserve markets and discuss how to implement it in the well-known electricity network simulation framework AMES. We discuss how the balancing market can integrate the imbalance and ability to respond from both supply and demand.

In future work, we plan to evaluate the proposed bidding process against an implementation of a contemporary two-market solution. A promising extension to the day-ahead market clearing is to include an estimation over $\sum_l^L P_l^{exe}$ and thus also minimise the expected costs of the balancing market $\sum_g^G P_g^{exe}\gamma_{exe}$. This could lower the impact of inefficiencies, which might become known only during balancing, on generation costs.

10. ACKNOWLEDGMENTS

This work is a part of the IDeaNeD project and sponsored by Agentschap NL, a research funding agency of the dutch ministry of economic affairs, in the IOP-EMVT program.

11. REFERENCES

[1] R. a.C. Van der Veen and L. J. De Vries. The impact of microgeneration upon the Dutch balancing market. *Energy Policy*, 37(7):2788–2797, July 2009.

[2] M. Cain and F. Alvarado. Implications of cost and bid format on electricity market studies: linear versus quadratic costs. In *Large Engineering Systems Conference on Power Engineering*, number July, 2004.

[3] Department Of Energy. Report to Congress: Impacts of the Federal Energy Regulatory Commission's Proposal for Standard Market Design. 2003.

[4] D. Gan and E. Litvinov. Energy and reserve market designs with explicit consideration to lost opportunity costs. *Transactions on power systems*, 18(1):53–59, 2003.

[5] D. Goldfarb and A. Idnani. A numerically stable dual method for solving strictly convex quadratic programs. *Mathematical Programming*, 27(1):1–33, Sept. 1983.

[6] E. Hirst and B. Kirby. Ancillary Services: The Neglected Feature of Bulk-Power Markets. *Electricity Journal*, 11(3):50–57, 1998.

[7] J. C. Hull. *Options, futures, and other derivatives.* Prentice Hall, Upper Saddle River, 3rd edition, 1997.

[8] A. Jokic. *Price-based Optimal Control of Electrical Power Systems.* Phd thesis, TU Eindhoven, 2007.

[9] R. Kamat and S. Oren. Rational buyer meets rational seller: Reserves market equilibria under alternative auction designs. *Journal of Regulatory Economics*, 2002.

[10] P. Klemperer and M. Meyer. Supply function equilibria in oligopoly under uncertainty. *Econometrica*, 57(6):1243–1277, 1989.

[11] Midwest System Operator. Energy and Operating Reserve, 2010.

[12] F. Schweppe. *Spot pricing of electricity.* Kluwer academic publishers, 1988.

[13] S. Stoft. *Power Market Economics - Designing Markets For Electricity.* Wiley-Interscience, 2002.

[14] J. Sun and L. Tesfatsion. Dynamic Testing of Wholesale Power Market Designs: An Open-Source Agent-Based Framework. *Computational Economics*, 30(3):291–327, Aug. 2007.

[15] TenneT. The Imbalance Pricing System, 2005.

[16] L. Tesfatsion. Agent-based computational economics: growing economies from the bottom up. *Artificial life*, 8(1):55–82, Jan. 2002.

[17] L. Vandezande, L. Meeus, R. Belmans, M. Saguan, and J.-M. Glachant. Well-functioning balancing markets: A prerequisite for wind power integration. *Energy Policy*, 38(7):3146–3154, July 2010.

[18] P. Vytelingum, S. Ramchurn, T. Voice, A. Rogers, and N. Jennings. Trading agents for the smart electricity grid. In *Proceedings of the ninth international joint conference on Autonomous agents and multiagent systems*, pages 897–904, 2010.

[19] T. Wu, M. Rothleder, Z. Alaywan, and a.D. Papalexopoulos. Pricing Energy and Ancillary Services in Integrated Market Systems by an Optimal Power Flow. *IEEE Transactions on Power Systems*, 19(1):339–347, Feb. 2004.

Multi-Agent System For Nuclear Condition Monitoring

Christopher J. Wallace
Institute for Energy And Environment
University Of Strathclyde
Glasgow, UK
cwallace@eee.strath.ac.uk

Gordon J. Jahn
Institute for Energy And Environment
University Of Strathclyde
Glasgow, UK
gjahn@eee.strath.ac.uk

Stephen D.J. McArthur
Institute for Energy And Environment
University Of Strathclyde
Glasgow, UK
smcarthur@eee.strath.ac.uk

ABSTRACT

This paper presents experience of Multi-Agent Systems (MAS) within the nuclear generation domain, with a particular emphasis on Condition onitoring (CM). A prototype system, which used a MAS for the management of CM observations and the generation of a user interface, was developed however it was decided that this application lacked the required speed and efficiency. Extending the developed ontology, an alternative MAS system was developed to perform CM analyses and storage and management of data, without the user directly interacting with the MAS.

The paper discusses the development of both systems, relating key features of MAS which meet the particular needs of CM applications, but also highlights remaining issues, such as a lack of an existing ontological framework and the problem of data security, that are particularly prominent in the nuclear domain.

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: Miscellaneous

General Terms

Design, Reliability, Security

Keywords

condition monitoring, nuclear, generation

1. INTRODUCTION

The use of Intelligent Agents for condition monitoring has been shown for a variety of applications in power systems [7][13][15][10][12] due to the flexible and distributed nature of Multi-Agent Systems (MAS). As lifetime extensions are granted to nuclear power stations around the world, there has been a increased interest in the implementation of Condition Monitoring (CM) to satisfy regulatory bodies and supplement knowledge of reactor state between statutory outages and inspections.

This paper describes the development of MAS for addressing nuclear condition monitoring challenges and considers the requirements, benefits and challenges that arise from the use of agent technologies for continuous offline monitoring of the cores of the UK's Advanced Gas-cooled Reactors (AGRs).

1.1 Nuclear Agents

Previous work on intelligent agents in the nuclear industry has largely considered the use of agents for the automated control of future power stations [18][9] and for online monitoring of systems related to the reactor [19]. The application described in this paper however, is designed for engineering support in order to detect anomalous behavior in reactor cores before there are safety concerns or adverse impacts on operations. The use of a MAS to achieve this, it will be argued, is a way for agents to *prove themselves* in safety critical applications, while the MAS platform allows for the addition of further AI analysis techniques encapsulated as agents.

There are, however, key practical challenges associated with the adoption of new technologies such as MAS within the nuclear engineering domain. These are addressed in section 7 before the paper concludes that whilst condition monitoring in the nuclear domain with MAS is possible and beneficial, there remain challenges to be overcome.

2. NUCLEAR CONDITION MONITORING

The continued operation of nuclear power plants beyond their original design lives is dependent on the ability of the operator to demonstrate the structural integrity of critical components. In the UK, the graphite cores of AGRs are the principal life limiting components [17] and are inspected every three years. To supplement these inspections, a condition monitoring strategy has been introduced at some stations, the results of which are considered at quarterly Monitoring Assessment Panel (MAP) meetings, where the observations are graded and recorded. There are a number of analyses associated with key core components, which form the basis of the MAP discussion. These analyses generally involve the extraction of a data set prior to the MAP, analysis of the data by a suitably qualified engineer and presentation of results at the MAP for consideration by a panel of experts. These experts then consider the data presented with a view to identifying temporal and spatial colocation that might indicate a root-cause issue with the reactor core itself.

The initial requirement for reactor core CM within the AGR fleet was a solution to the problem of the management of CM analysis results and observations. The nature of slightly different hardware configurations at different sta-

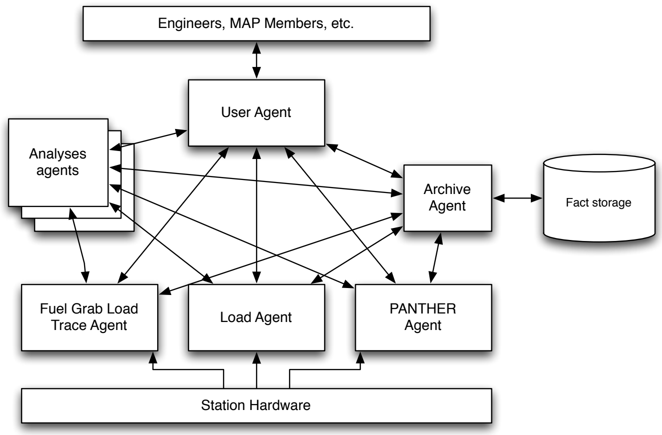

Figure 1: IMAPS High-Level Block Diagram

tions requires that the approach be flexible to allow different analyses to be undertaken.

3. MAPS AND IMAPS

The MAP process was trialed at two power stations from 2005, and is a formal process for collating and analyzing data. As a new process, there were no supporting systems and it was recognized that there was the opportunity to trial MAS within this environment. This project was known as IMAPS, the Intelligent Monitoring Assessment Panel System, which delivered an agent-based platform for the storage and analysis of AGR observations.

The primary means of supporting the MAP process is by allowing observations of different types to be recorded and plotted on a graphical plan view of the reactor, originally using a whiteboard. The simplicity of this approach for allowing the identification of spatially colocated events and the easily understandable relationship between the actual core and data plotted on such a view meant that this basic interface was carried forward into the MAS.

As the system was being developed alongside the formal station-based process, there was also a need to allow existing manual interpretation to be included within the system as soon as possible, allowing data to be recorded and analyzed as soon as possible. The initial prototype of this system involved the creation of an archiving agent along with the user agent. This met the initial brief of providing a working system for storing the data and presenting this back to engineers.

The high-level block diagram of the system envisaged is shown in figure 1.

Due to the specific requirements in this case, the final delivery to end-users was to be provided through a web-browser that posed the challenge of allowing asynchronous user behavior to interact with an agent with the MAS that would be regularly communicating with other agents.

The approach taken to solve this problem was to implement a MAS based on the Foundation for Intelligent Physical Agents (FIPA) standards. This was carried out using the Java Agent DEvelopment framework (JADE)[2] and an agent for each user was executed within a web server; this user agent provided an interface between user inputs provided via a web browser and the agent based system and also with the rendering of information back to the user.

The initial system contained simply the user agents and the archive agent which provided a translation between FIPA Semantic Language (SL) and the Structured Query Language (SQL) used by the database fulfilling the fact storage requirement.

An issue that arose with this technique was of the amount of data requiring to be presented to users. The screenshot in figure 2 shows a typical rendering for a single reactor. As each station has two reactors, the normal case is to present two reactors simultaneously.

Each reactor plan shown has around 400 channels including larger fuel channels and smaller interstitial channels used for control and other purposes. This means that at any given time, the system is rendering around 800 data points. Later iterations of the system also used the border color to represent the presence of inspection information taking each view as above to representing around 1,600 data points.

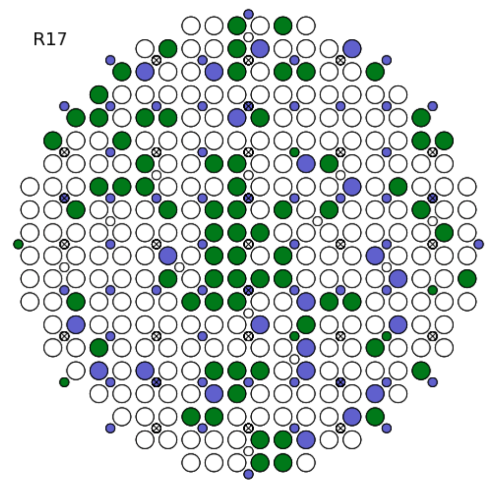

Figure 2: Sample data rendering from IMAPS. Two cores per station are available to the user.

IMAPS [11] has a relatively small domain model to contend with. There is a type hierarchy associated with the various observation types and other concepts must be defined to support the reactors, the station and the different grades that can be assigned to the observations. A base ontology was sought for this application, but no suitable candidate ontology was found and, instead, one was developed specifically for the application. The frames-based ontology shown in figure 3 was used for the initial testing of the system.

Following experimentation with this system, the prototype structure was found to be unable to meet response time expectations. Upon investigation, this was found to be due to the number of data movements having to occur in order that the information be rendered to the user. The following steps were taking place:

- The user agent parses the request from the web browser
- The user agent translates the parsed request into a FIPA SL query

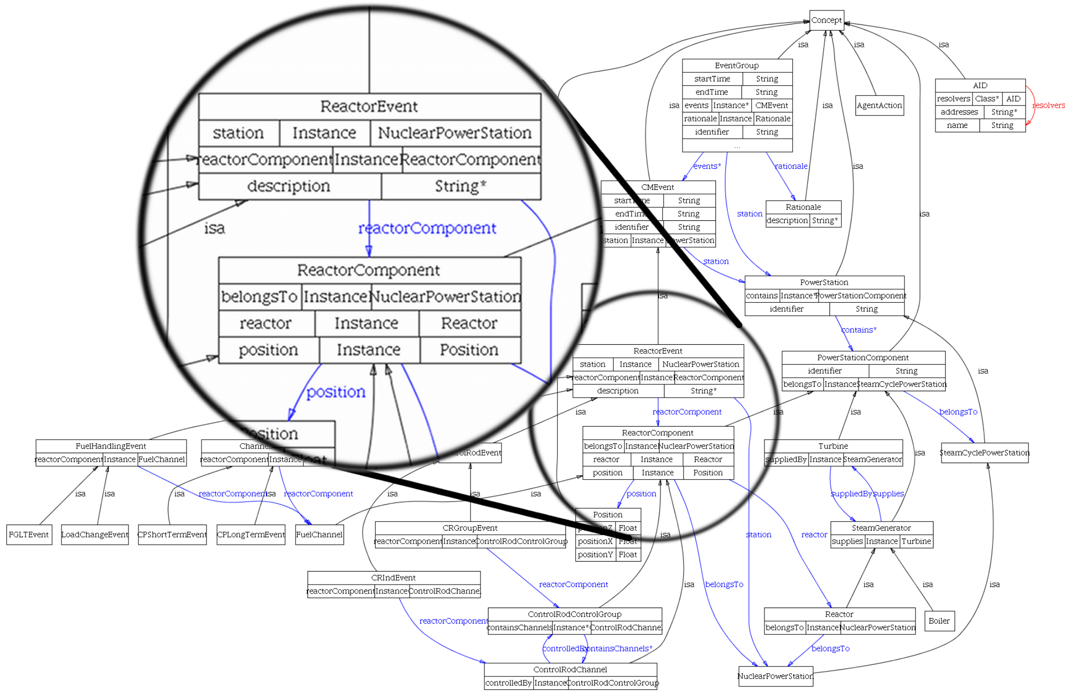

Figure 3: Initial IMAPS ontology, with two zoomed examples of concepts and how they relate to concepts by predicates such as *isa*.

- The user agent envelopes the FIPA SL query and sends it to the archive agent
- The archive agent parses the SL query translating the content and predicates into an SQL query
- The archive agent executes the SQL query and re-creates abstract objects
- The abstract objects are used to render an SL response to the query
- The SL response is enveloped and sent back to the user agent
- The user agent parses the data back into objects for use by the renderer

Due to the structure of the messaging within the agent, this process involves repeated handling of the same data in slightly different forms. This both increases the use of system resources and slows down the data handling. Whilst the advantages of this approach tend to outweigh the disadvantages by making it available in an agent-based society, the fact that this use case is so important to the end-user experience meant that further investigation was warranted.

Analysis revealed that the agent-based approach to retrieving the data ws taking around seven times longer than directly retrieving this data from the database. This problem was ultimately caused by JADE operating on very long strings whilst building the SL response and enveloping the message.

In order to meet the design goal, the agent interaction was removed from the process with the server directly parsing the query, translating to SQL and executing in the database before streaming results – without ever storing them – direct to the client. Render times in this case were brought to around 300msec for the same data on the same hardware platform. Whilst this meant that the prototype was now without the multi-agent design in terms of retrieving data, it was retained for data logging and this hybrid approach was essential to allow the research and development of further analyses.

4. EXTENDING IMAPS INTO ANALYSIS

In order to consider how the prototype system described in the previous section can be extended to provide and support routine analyses, the requirements for the analysis processes are considered along with the changes that are then required in the IMAPS MAS.

4.1 Requirements

Existing CM of the AGR fleet is based on a model of extracting data or information from reactor systems and data stores as required, in order to perform analyses or explain anomalous analysis results according to a monitoring schedule. By delegating tasks and analyses currently performed by engineers to intelligent agents, it is possible to automate the collection, analysis, correlation and verification of data as it becomes available.

This design of CM system allows for faster completion of analyses and requires considerably less engineer time, and allows more time for investigation of non-trivial anomalies, which can then be considered in more detail prior to the MAP meeting.

The current CM regime manually records anomalous observations for particular components in the core, however a more comprehensive monitoring system could monitor the data for all components for further analysis and trending of anomalies not currently detected by existing monitoring analyses, or results dismissed by current analysis limits that may be changed later to meet an enhanced safety case.

Finally, the monitoring of AGR cores is still a relatively young process, compared to the existing inspection protocols, and as such it is essential that a CM system be flexible enough to include additional data sources and further analyses as required. The use of a MAS with a carefully defined ontology allows for this extensibility by encapsulating any relevant data interface or analysis within an agent that utilizes the ontology. This provides a platform within which Artificial Intelligence (AI) techniques can be deployed to perform further analyses with very minimal impact on the existing system. Furthermore, such analyses may extend the current work performed by engineers by not only considering known causal relationships between events, but also statistical relationships using clustering or structural health monitoring.

4.2 Ontology and Development

Based on the success of previous MAS developed for power applications[5], and the original prototype agents and ontology developed for IMAPS along with the maturity of the platform, JADE was selected for further development.

The ontology developed for the initial version of IMAPS provides the basis of a nuclear ontology for this work, however since the initial emphasis of the project was the management of CM observations, expansion of the ontology was required in order to facilitate the management, storage and analysis of reactor data. The ontology was created and developed using the Protégé ontology editor[3] and a plugin which allowed the automatic generation of Java code, an example of which is shown in figure 4. Each concept in the CM system is therefore represented by a Java class and takes advantage of object oriented concepts such as inheritance, resulting in a class structure of the same form as the ontology.

The Protégé bean generator used has also been extended to support the Java Persistence API (JPA), allowing it to add annotations which are included in the generated code and are used in mapping of the ontology to a database for storage and retrieval of data and observations.

4.3 Storage

The archive agent is essentially a small messaging and translation layer added to a standard SQL database system providing storage. For database storage, an object relational mapping library called Hibernate is used, which maps Java classes (generated from Protégé) directly onto an SQL database. Data and observations that are created by the MAS and stored in the database can then be retrieved as Java objects, which allows for simplified querying using the Hibernate Query Language (HQL) which allows object oriented queries against the methods available in the objects stored. For example, the query `select r from Reactors r` returns all instances of the Reactor fact in the database, along with providing access to related objects.

This format is significantly simpler than a typical SQL query which would depend on the table structure of the database and may involve complex operations to retrieve the same data. The use of object relational mapping therefore adds portability to the system and with only slight modification, the database can be changed to any supported SQL database.

Coupled to an SL translator that can understand arbitrary SL queries and map predicates onto SQL queries, this approach supports other agents presenting the archive agent with any valid SL for processing.

```
/* Protege name: ReactorComponent */
@Entity
@DiscriminatorValue("ReactorComponent")
public class ReactorComponent
          extends PowerStationComponent{

   /* Protege name: reactor */
   @ManyToOne(cascade = {CascadeType.ALL})
   @JoinColumn
   private Reactor reactor;
   public void setReactor(Reactor value) {
    this.reactor=value;
   }
   public Reactor getReactor() {
     return this.reactor;
   }
```

Figure 4: An example of the Java Bean code generated using the Protege ontology development application

Data Collection and Distribution

Service: *Collect and Distribute Data*
Inputs: *Data*
Outputs: *Data, Notifications*
Pre-Conditions: *Data exists*
Post-Conditions: *Data sent to database, analysis agent alerted*

Figure 5: An example of an agent role defined using the Gaia Service model

4.4 Design Methodology

The design process of MAS has been discussed in the literature[20], and in particular for power systems[14]. The initial methodology for the MAS design roughly followed the same general path as McArthur et al[14] and several aspects of the Gaia methodology[21], in particular for the development of the Service Model, which defines the roles for which agents are created and the capabilities they can offer the system. The emphasis on defining the services offered by the agents reflects the desire of the system, initially at least, to automate the tasks and procedures already defined and performed by engineers. These roles were defined formally using the structure of the Gaia services model, an example of which is shown in figure 5.

The approach of using this service model to map between tasks and agents within the MAS is now demonstrated through a brief case study showing how a simple analysis has been implemented.

5. CASE STUDY: CHANNEL POWERS

An example of an existing analysis which can be conveniently performed by an agent based system, is the regular comparison of channel power models, which quantifies and attempts to explain differences between thermal and neutronic models of channel power[8]. Excessive deviations between the thermal and neutronic models of channel power could potentially indicate inadequate cooling of fuel, however these can far more commonly be attributed to innocuous local events such as refueling or other normal reactor operations that affect the flux profile across the core.

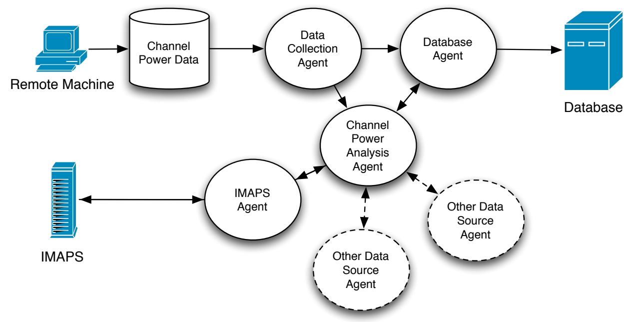

Figure 6: An example of an engineers monitoring procedure mapped onto a set of agents.

Using a set of pre-defined limits, the power discrepancies are graded initially, and then an engineer is tasked with determining whether any reactor operations, such as refueling, adequately explain the discrepancy. Using a looser set of limits, the engineer then re-grades the discrepancies as appropriate using any available information.

5.1 Agent Approach

Mapping the procedure followed by the engineer to a set of algorithms and with the appropriate knowledge of available data sources, a set of agents as shown in figure 6:

- Parses the data (*Data Collection Agent*)
- Stores the data (*Database Agent*) in a database in a form consistent with the ontology.
- A notification is sent to the *Channel Power analysis Agent*
- The *Channel Power analysis Agent* retrieves the data
- The *Channel Power analysis Agent* sends queries to available data source interface agents (the *IMAPS Agent* in this case) for any data or observations within a defined spatial or temporal range of each fuel channel power discrepancy for which it analyzes.
- After correlating all possible events with the data,the gradings are calculated and then stored in the database for verification by the engineer.

Other data sources may be utilized in this particular type of analysis, the availability and nature of which may vary between stations, however provided the analysis agent is aware of relevant data sources and how the data from each should be interpreted, an arbitrary number can be added.

6. FUTURE FUNCTIONALITY

The future development of the system will include the development of agents to perform other MAP analyses as well as an investigation of other design approaches to the AGR CM requirements. The design presented in this paper is very task focused, with the agents designed around the required functionality already provided by the engineer. It can be argued however that this approach is inherently limited by attempting to emulate the engineer, as it places very little emphasis on the autonomy of the agents (outside of the prescriptive CM they perform) and by extension any emergent behavior that could potentially enhance the CM regime.

6.1 Component Monitoring

For this reason, recent work has begun investigating the viability creating agents for major components in the core (fuel channels, control channels, coolant pumps etc) that can be delegated the responsibility of managing their own data and analyses. Though this would involve the creation of around 300-400 agents per reactor core, it has been shown[6] that the JADE platform scales quite well to systems of several hundred agents. It is likely in the interests of performance however that multiple agent containers would be used, for example an analysis container on one machine that performs numerically intensive algorithms and component agents in a separate container.

These agents would be able to request appropriate analyses as data becomes available to them, keep snapshots of the most recent states of the component and track changes of various parameters. They would also have the ability to share this data with other components, and in certain circumstances, such as an anomalous result in an analysis, suggest to other component agents particular analyses that they might wish to perform. This *rolling* model of CM would allow for greater depth of analysis of parameters for which there are often secondary causes and effects, for which the availability of data is not synchronized. The system would of course still be required to perform at least the minimum number of analyses required for the MAP.

Individual component health monitoring will become more important as the core ages and a more detailed view of the state of groups of components or the whole core is required.

It is envisaged that this may present itself as an agent hierarchy with different agent societies responsible depending on the component grouping. This could manifest itself as a layered architecture as shown in figure 7, substituting layered analyses into the original vision outlined in figure 1.

This method of performing analyses at different levels of component grouping should allow for both mirroring the manner in which analyses are carried out by experts – increasing trust and transparency within the system – but should also allow for competing theories regarding localized and global hypotheses for each reactor. These can be presented and assessed by the reactor agents to present a consolidated case, with a risk profile, enabling the engineer to make an informed judgment based on the outputs of the MAS.

7. CHALLENGES

The use of ontologies is intended to allow for uniformity of development and communication, however the question of how to manage revisions of ontologies or the interaction of different ontologies[4] is a problem to which there is not a

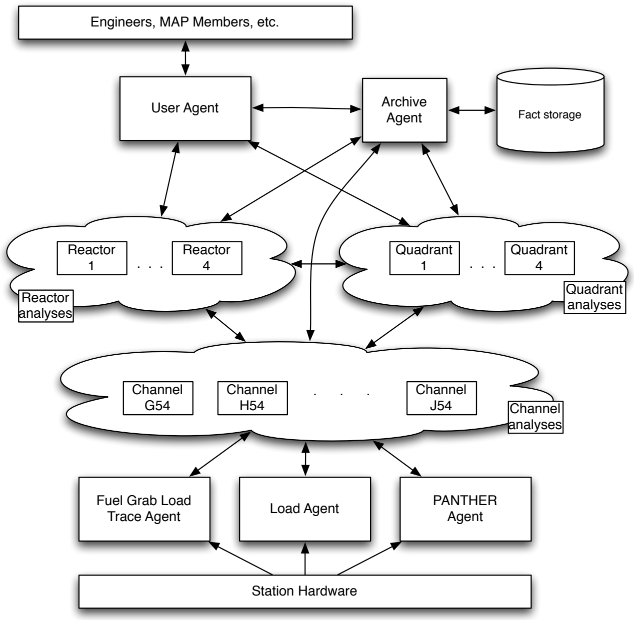

Figure 7: Proposed High-Level Architecture

well defined solution. The IEEE MAS working group[1] has already proposed the adoption of a common upper ontology for MAS in the field of power engineering, and though ontologies are application and domain specific, it remains unclear as to whether a *nuclear ontology* merits development, as no appropriate re-usable ontology was available during the course of this work.

Each reactor data set or CM observation object that the MAS handles, uniquely describes some fact, however there is a significant overhead in checking the data store each time a new fact is recorded to check whether the fact is already stored. Similarly, where datasets overlap in time, the archiving agent is required to determine how much duplicate data exists and how to store the new data set. How to achieve this without adding significant overhead to the processing of data remains an issue.

7.1 Nuclear Domain Challenges

The particular importance of safety in a Nuclear Power Station means that deployed technologies are required to be proven and their impact well understood. This results in a relatively slow adoption of new technologies, so while a MAS may be in principle capable of performing some of the tasks of an engineer, it is likely that any initial deployment of a nuclear CM MAS will be very much in a supporting role and will only be used for verification of the existing engineers analysis. It will also likely be held separately from the main IT infrastructure with several well defined data interfaces for the purposes of verifying data integrity and security.

The need to protect the operational plant and associated sensor and control systems from any interference by unauthorized parties has historically been met by complete segregation of operational and corporate networks. Whilst this ensures malware and control signals cannot pass from, for example, Internet e-mail to the plant, it also reduces interoperability within the plant and results in increased difficulties in accessing online data for analysis.

With the implementation of CM across fleets of nuclear power stations, the question of data security has become increasingly prevalent as operators endeavor to provide access to relevant data while maintaining the integrity of their IT security infrastructure. MAS by their very nature depend on the availability of network access to facilitate communication between agent containers on different remote machines. For the purposes of the work described in this paper, the test system was not deployed and the collection of data was simulated through access to remote systems (such as IMAPS), databases and remote file shares. In an operational environment, security is of major importance and access to data will be much more stringent. One possible solution involves the definition of information security zones with the plant and using "data diodes" that often claim to provide 100% one way data passage. As explained in [16] however, such devices are far from a "silver bullet" due to the inability of the receiver to provide feedback, such as acknowledgements, to the sender and this could result in missing data.

8. CONCLUSIONS

This paper has described the application of MAS to the condition monitoring of a nuclear reactor. Initial work implemented an agent based system for the management of condition monitoring observations, however it was found that using agents to generate a user interface was inefficient and caused excessive delays in updating the display. An alternative use of MAS was developed, by extending the existing ontology, to perform time consuming analyses currently performed by engineers. This system can provide a consistent structure for the addition of many analyses and distributed data sources, which are common features of the monitoring of older nuclear reactors. A MAS was found to be a good choice of platform for this problem, however some issues associated with the application of an agent based system in the nuclear domain remain, particularly the lack of an existing ontological framework and issues associated with data security.

9. ACKNOWLEDGMENTS

This work was funded by EDF Energy, the views presented here do not necessarily represent the views of EDF Energy.

10. REFERENCES

[1] http://ewh.ieee.org/mu/pes-mas/.

[2] http://jade.tilab.com/.

[3] http://protege.stanford.edu/.

[4] V. Catterson, E. Davidson, and S. McArthur. Issues in integrating existing multi-agent systems for power engineering applications. In *Proceedings of the 13th International Conference on Intelligent Systems Application to Power Systems*, 2005.

[5] V. Catterson and S. McArthur. The industrialisation of a multi-agent system for power transformer condition monitoring. In *Applications and Innovations in Intelligent Systems XII*, pages 165–178, 2005.

[6] K. Chmiel, D. Tomiak, M. Gawinecki, P. Karczmarek, M. Szymczak, and M. Paprzycki. Testing the efficiency of the jade platform. In *Third International*

Symposium on/Algorithms, Models and Tools for Parallel Computing on Heterogeneous Networks, 2004, pages 49–56, July 2004.

[7] A. Dimeas and N. Hatziargyriou. A multi-agent system for microgrids. *Methods and Applications of Artificial Intelligence*, pages 447–455, 2004.

[8] S. Haddock and G. Parks. Agr fuel management using panther. In *Proceedings of Fuel management and Handling*, 1995.

[9] J. Heo and K. Lee. A multi-agent system-based intelligent identification system for plant control and fault diagnosis. In *Power Engineering Society General Meeting*, 2006.

[10] J. A. Hossack, J. Menal, S. D. J. McArthur, and J. R. McDonald. A multiagent architecture for protection engineering diagnostic assistance. *IEEE Transactions on Power Systems*, 18(2), May 2003.

[11] G. Jahn, S. McArthur, and D. Towle. Staged implementation of an agent based advanced gas cooled reactor condition monitoring system. In *IEEE Power Engineering Society General Meeting*, 2007.

[12] S. McArthur, C. Booth, J. McDonald, and I. McFadyen. An agent-based anomaly detection architecture for condition monitoring. *IEEE Transactions On Power Systems*, 20(4):1675 – 1682, November 2005.

[13] S. McArthur, E. Davidson, V. Catterson, A. Dimeas, N. Hatziargyriou, F. Ponci, and T. Funabashi. Multi-agent systems for power engineering applications part i: Concepts, approaches, and technical challenges. *IEEE Transactions On Power Systems*, 22(4):1743 – 1752, November 2007.

[14] S. McArthur, E. Davidson, V. Catterson, A. Dimeas, N. Hatziargyriou, F. Ponci, and T. Funabashi. Multi-agent systems for power engineering applications part ii: Technologies, standards, and tools for building multi-agent systems. *IEEE Transactions On Power Systems*, 22(4):1753 – 1759, November 2007.

[15] T. Nagata, H. Watanabe, M. Ohno, and H. Sasaki. A multi-agent approach to power system restoration. In *Power System Technology, 2000. Proceedings. PowerCon 2000. International Conference on*, volume 3, pages 1551 –1556 vol.3, 2000.

[16] L. Piètre-Cambacédès and P. Sitbon. Deconstruction of some industrial control systems cybersecurity myths. In *Proceedings of the 6th International Topical Meeting on Nuclear Plant Instrumentation, Control and Human-Machine Interface Technologies*, 2009.

[17] A. G. Steer. Agr core design, operation and safety functions. In G. Neighbour, editor, *Management of ageing in graphite reactor cores*, pages 11–18. Royal Society Of Chemistry, 2007.

[18] R. E. Uhrig and L. H. Tsoukalas. Multi-agent based anticipatory control for enhancing the safety and performance of generation iv nuclear power plants during long-term semi-autonomous operation. *Progress In Nuclear Energy*, 42(1-4), 2003.

[19] H. Wang and C. Wang. Intelligent agents in the nuclear industry. *IEEE Computer Magazine*, 1997.

[20] M. F. Wood and S. A. DeLoach. An overview of the multiagent systems engineering methodology. *Lecture Notes in Computer Science*, 1957:207–220, 2001.

[21] M. Wooldridge, N. Jennings, and D. Kinny. The gaia methodology for agent-oriented analysis and design. *Autonomous Agents and Multi-Agent Systems*, 3(3), 2000.

Adaptive Home Heating Control Through Gaussian Process Prediction and Mathematical Programming

Alex Rogers, Sasan Maleki, Siddhartha Ghosh, Nicholas R. Jennings
School of Electronics and Computer Science
University of Southampton, SO17 1BJ, UK
{acr,sm9g09,sg2,nrj}@ecs.soton.ac.uk

ABSTRACT

In this paper, we address the challenge of adaptively controlling a home heating system in order to minimise cost and carbon emissions within a smart grid. Our home energy management agent learns the thermal properties of the home, and uses Gaussian processes to predict the environmental parameters over the next 24 hours, allowing it to provide real-time feedback to householders concerning the cost and carbon emissions of their heating preferences. Furthermore, we show how it can then use a mixed-integer quadratic program, or a computationally efficient greedy heuristic, to adapt to real-time cost and carbon intensity signals, adjusting the timing of heater use in order to satisfy preferences for comfort whilst minimising cost and carbon emissions. We evaluate our approach using weather and electricity grid data from January 2010 for the UK, and show our approach can predict the total cost and carbon emissions over a day to within 9%, and show that over the month it reduces cost and carbon emissions by 15%, and 9%, respectively, compared to using a conventional thermostat.

Categories and Subject Descriptors

I.2.11 [**Computing Methodologies**]: Distributed Artificial Intelligence

General Terms

Algorithms, Energy, Control, Learning

Keywords

Smart grid, machine learning, Gaussian process, mixed-integer optimisation, energy feedback, carbon emissions

1. INTRODUCTION

The creation of a smart electricity grid represents one of the greatest engineering challenges of this century, as countries face dwindling non-renewable energy sources and work to minimise the adverse effects of climate change due to carbon emissions [1, 2]. Key components of this vision include ambitious targets for renewable generation, the roll-out of smart meters to domestic consumers in order to facilitate real-time pricing of electricity, and a shift toward the electrification of heating through the use of air and ground source heat pumps.[1]

However, these developments present a potential challenge for householders, since the increased use of intermittent renewable generation means that both the price of electricity, and also the carbon intensity of the electricity (quoted in terms of gCO_2/kWh and signifying the amount of carbon dioxide emitted when one unit of electricity is consumed), will vary in real-time. These real-time signals will likely be passed to consumers, through their smart meters, who will be expected to respond rationally in order to reduce demand for expensive and carbon intensive electricity at peak times, and to make better use of low carbon renewable energy when it is available. However, the links between heater system control settings and energy consumption is already poorly understood by consumers [4], and these changes are likely to compound this issue. Thus, it is essential that future home heating systems are able to provide real-time feedback to householders concerning the implications (in terms of both cost and carbon emissions) of their heating preferences. Going further, these systems should adapt to these real-time signals, adjusting the timing of heater use in order to satisfy the home owners' preferences for comfort while also minimising cost and carbon emissions.

Now, the idea of individual homes responding to real-time signals from an electricity grid is not a new one, and indeed, was first proposed by [10], who discussed the scheduling of loads, and the prediction of both demand in the home, and local weather conditions. However, their work dealt mainly with predicting the overall system behaviour, and used closed form solutions that required approximating and homogenising the behaviours of the individual actors within the systems. As such, it does not actually provide a solution that can be implemented within any individual home. More recently, a number of researchers have revisited these ideas, and proposed solutions that address parts of the challenge described above. For example, [7] use artificial neural networks to predict the heat demand of individual homes using several external factors, but make no attempt to then optimise energy use. Conversely, [11] use mixed integer programming to optimise energy storage within a home which receives real-time price signals, but they consider electrical loads, rather than heating, and hence they do not need to consider the thermal characteristics of the home, nor the external factors that affect it.

Thus, to address these shortcomings, in this paper we develop a home energy management agent, or smart controller,

[1] The UK government has committed to reducing carbon emissions by 80% by 2050, and aims to install smart meters to all 26M UK homes by 2020 [6].

that works on behalf of the householder, and is able to learn the thermal characteristics of the home, and predict the additional factors that affect the cost and carbon emissions of heating use. It is then able to provide householders with the required real-time feedback, and also to autonomously control the heating system on their behalf. In more detail, using internal and external temperature sensors, and by monitoring the activity of the home's heating system, the smart controller finds model parameters that describe the heat output of the heating system and the thermal leakage rate of the home. In addition, using a Gaussian process model that exploits periodicities and correlations between time series, the smart controller predicts the local external temperature over the next 24 hours by combining local measurements from a sensor, with predictions from an online weather forecast. In doing so, it creates a *site-specific* forecast for the next 24 hours. Similarly, using available predictions of future demand, the smart controller predicts the carbon intensity of the electricity supplied to the home. Using these factors it can predict the consequences, in terms of cost and carbon, of any heater control setting and provide this information to the home owner through a graphical user interface. Going further, the smart controller is then able to fully optimise the use of heating (using either a mixed integer solver or a computationally efficient greedy heuristic). In doing so, it seeks to provide the same level of comfort as a standard thermostat operating at the same set-point temperature (evaluated using a comfort model based on the ANSI/ASHRAE Standard 55-2010) whilst also minimising either cost or carbon emissions.

Thus, in more detail, in this paper we make the following contributions to the state of the art:

- We present two novel formalisms that uses Gaussian processes (GP) to predict external temperature and carbon intensity by exploiting (i) periodicities within each single time-series, and (ii) correlations between each time series and another time series (in this case, the temperature at a nearby location and total electricity grid demand) for which a prediction is made available by an outside agency (e.g. a local weather centre and the grid operator). The first extends an existing multi-output GP model that has previously only been used for regression [8], rather than the prediction that we do here. The second uses one time-series to represent the mean value of the second, and applies a single-output GP to model the difference. We empirically evaluate both and indicate their strengths and weaknesses in our domain.

- We present a novel approach to optimising heating whereby the smart controller models the comfort that would result from using a thermostat (set to any particular set-point temperature) to control the heating, and then, optimises heating to ensure the same comfort level, at minimum cost or carbon. We show that this problem has quadratic constraints introduced by the comfort model used, and integer constraints determined by the heater output. Thus, we present both a mixed-integer quadratic program, and a computationally efficient greedy heuristic, to solve it. We show that the greedy heuristic finds solutions that are as competitive as those of a commercial solver (IBM ILOG CPLEX), but in less time.

- We combine these within our smart controller and empirically evaluate it using real UK weather and electricity grid data from January 2010, and show our approach can predict the total cost and carbon emissions over a day to within 9%, and show that over the month it reduces cost and carbon emissions by 15%, and 9%, respectively, compared to a conventional thermostat.

The remainder of this paper is organised as follows. In Section 2 we present our model of the home and heating system. In Section 3 we describe how the smart controller provides real-time feedback, and in Section 4 we describe how it optimises heating use. We conclude in Section 5.

2. HOME THERMAL PROPERTIES

We consider a standard thermal model in which heat leaks from the home (by thermal conduction and ventilation losses) at a rate that is proportional to the temperature difference between the air inside and outside of the home [3]. In more detail, let $\phi \in \mathbb{R}^+$ be the thermal leakage rate of the house measured in W/K. We divide the day into a set of discrete time slots, $t \in T$, and denote the internal temperature of the home at time t as $T^t_{in} \in \mathbb{R}^+$ and the external temperature (in K) as $T^t_{ext} \in \mathbb{R}^+$.

We assume that the home is heated by an electric heat pump, whose thermal output (in kW) is given by r_h. We denote $o_h \subseteq T$ as the set of time slots at which the heat pump is actively producing heat, and also define the variable $\eta^t_{on} \in \{0,1\}$ for every $t \in T$ such that $\eta^t_{on} = 1$ if $t \in o_h$ and 0 otherwise. Given this, the amount of energy delivered (or lost) from the home, η^t, in any time slot is given by:

$$\eta^t = \eta^t_{on} r_h - \phi\left(T^t_{in} - T^t_{ext}\right) + \epsilon^t \tag{1}$$

where ϵ^t is Gaussian noise reflecting additional effects due to the householder's activities (e.g. opening or closing windows, cooking, etc.). The internal temperature of the home after this energy transfer is then given by:

$$T^{t+1}_{in} = T^t_{in} + \frac{\eta^t \Delta t}{c_{air} m_{air}} \tag{2}$$

where Δt is the duration of the time slot (in seconds), and the heat capacity and the total mass of air in the home are $c_{air} \in \mathbb{R}^+$ (in J/kg/K) and $m_{air} \in \mathbb{R}^+$ (in kg) respectively.

The heat pump is controlled by a timer and a thermostat. We denote $o_t \subseteq T$ as the set of time slots at which the heat pump is enabled (but not necessarily actively producing heat), and we define the variable $\eta^t_{timer} \in \{0,1\}$ for every $t \in T$ such that $\eta^t_{timer} = 1$ if $t \in o_t$ and 0 otherwise. The thermostat acts to keep the internal temperature of the home at the thermostat set point, T_{set}, by applying the rule:

$$\eta^t_{on} = \begin{cases} 0 & T^{t-1}_{int} > T_{set} + \Delta T, \\ 1 & T^{t-1}_{int} < T_{set} - \Delta T, \\ \eta^{t-1}_{on} & \text{otherwise.} \end{cases} \tag{3}$$

for all $t \in o_t$. Note that ΔT induces hysteresis such that the thermostat does not continually cycle at the set point.

Finally, we consider the cost and carbon emissions of the electricity that supplies the heater. We assume that the electrical power of the heat pump is given by r_e (in kW). Note that the heat energy produced by a heat pump is greater than the electrical power used, such that, $r_h = COP \times r_e$, where COP, the *coefficient of performance*, is typically between 2 and 4. Thus the total cost of electricity consumed over T is given by $\sum_{t\in o_h} r_e p^t \Delta t$, and the total carbon emissions are given by $\sum_{t\in o_h} r_e c^t \Delta t$, where p^t is the price (pence per kWh), and c^t is the carbon intensity (in gCO_2/kWh), of the electricity at time t.

3. REAL-TIME ENERGY FEEDBACK

Now, to provide real-time feedback to householders on the cost and carbon emissions of their heating preferences, it is not sufficient to simply provide instantaneous measurements of the heater's use, since this requires that the householder extrapolate these figures to the end of the current day. Rather, we require that our smart controller perform this integration automatically, and to do so, it must understand the thermal properties of the home (specifically, the thermal output of the heater and the leakage rate of the home), and also the environmental factors that effect the cost and emissions (specifically, the external temperature and the carbon intensity of the electricity grid). In the next sections, we show how we can learn the former through a simple regression process, and show how we can make prediction of the later, over a 24 hour period, using Gaussian processes.

3.1 Learning Thermal Properties

As described in Section 2, the thermal characteristics of the home is dependent on the thermal output of the heater, the leakage rate of the home, and a number of parameters such as the thermal capacity and the mass of air inside the home. Rather than attempting to estimate all these parameters, the smart controller uses a thermal model, is just as expressive, but which is denoted in terms of the air temperature increase over a time slot, $\mathcal{R} \in \mathbb{R}^+$ (in K), due to the heater and the decrease, $\Phi \in \mathbb{R}^+$ (unitless), due to leakage:

$$\overline{T}_{in}^{t+1} = \overline{T}_{in}^{t} + \mathcal{R}\eta_{on}^{t} - \Phi\left(\overline{T}_{in}^{t} - T_{ext}^{t}\right) \quad (4)$$

Given historical observations (typically over the preceding 24 hours) of internal, T_{in}^t, external, T_{ext}^t, temperatures and the times when the heater was providing heat, η_{on}^t, the smart controller can predict the evolution of the internal temperature over the same period (initialising at $\overline{T}_{in}^1 = T_{in}^1$). The error in this prediction is given by:

$$\sum_{t\in T}\left(\overline{T}_{in}^{t} - T_{in}^{t}\right)^2 \quad (5)$$

Thus, the best estimates of the heater output, $\mathcal{R}$, and leakage rate, Φ, are then those that minimise this error.

3.2 Predicting Environmental Factors

We now turn to predicting the future values of the environmental factors that affect the cost and carbon emissions of the home (specifically, the external temperature and the carbon intensity of the electricity grid). To do so, we exploit the fact that these two time series are likely to be correlated with other time series for which we are able to retrieve predictions over the internet. For example, we expect the external temperature to be correlated to that of the local weather forecast (but not necessarily identical to it due to additional local factors in the immediate environment of the home). Similarly, we observe that carbon intensity is often closely correlated with the total demand for electricity within the grid, for which the grid operators provide accurate predictions in real-time for the next 24 hours. Thus, to predict these correlated time series, we present two novel variations of the Gaussian process (GP): a multi-output GP which explicitly parameterises the correlations between the time series, and a *difference* GP where we model the difference between the normalised time series as a single-output GP.

3.2.1 Gaussian Process Prediction:

A Gaussian process (GP) represents a function as a multivariate Gaussian distribution [9]. It is specified by a covariance function $k(t,t')$, describing correlations between observations of the function at different times, and a mean function, $\mu(t)$, typically taken to be zero, that is the expected value of function prior to seeing any training data. Within the energy domain, Leith et. al used GP for electrical demand forecasting, and shown that it consistently generates significantly more accurate forecasts than comparative time series models [5]. However their work addresses a single time series only, unlike our case where we consider multiple correlated time series.

Now, as discussed above, we have one main time series, for which we have local historical observations, T_M, (i.e. external temperature and carbon intensity), and another correlated time series, for which we again have historical observations, T_C, and also a future prediction over the next 24 hours supplied by an outside agency, T_C^*, (i.e. historical temperature measurements and a 24 hour forecast at the local weather centre, and historical observations of demand and a 24 hour prediction from the grid operator). Our aim is to predict the values of the main time series over the same interval as T_C^*, and we denote these prediction T_M^*.

3.2.2 Multi-Output Gaussian Process:

We first apply a multi-output GP in which the cross-correlation between the time series is explicitly represented as a hyperparameter [8]. To this end, we represent the correlation between two observations in either time series as the Hadamard product of a covariance function over time alone and a covariance function over the time series labels:

$$k([l,t],[l',t']) \triangleq k_L(l,l')k_T(t,t') \quad (6)$$

The covariance function over labels is simply given by:

$$k_L(l,l') \triangleq \begin{cases} 1 & \text{if } l = l', \\ \rho & \text{otherwise,} \end{cases}$$

where $k_L(l,l')$ is unity if both inputs are from the same time series and where ρ is a hyperparameter that determined the cross-correlation between the time series.

The temporal covariance function, k_T, provides more flexibility and is composed of a number of more basic covariance functions. Our first choice is to use the standard *squared exponential* function, given by

$$k_{SE}(t,t') = \sigma_f^2 \exp\left(-\frac{(t-t')^2}{2\ell^2}\right) \quad (7)$$

where σ_f^2 is the amplitude of the process and ℓ is the characteristic length-scale that determines how rapidly the correlation between outputs should decay as t and t' diverge. As the squared exponential implies a strong smoothness assumption regarding the process, we also use the Matérn class

Table 1: Gaussian Process Covariance Functions

Quantity	Method	w	Covariance Function, k_T
Carbon Intensity	Single-output	1	$k_{SE} + k_P + k_N$
	Multiple-output	2	$k_L \cdot (k_M + k_{SE} \cdot k_P) + k_N$
	Difference	2	$k_{SE} + k_{SE} \cdot k_P + k_N$
External Temperature	Single-output	2	$k_{SE} + k_{SE} \cdot k_P + k_N$
	Multiple-output	2	$k_L \cdot (k_M + k_{SE} \cdot k_P) + k_N$
	Difference	1	$k_{SE} + k_{SE} \cdot k_P + k_N$

to represent processes with more erratic changes:

$$k_M(t,t') = \sigma_f^2 \left(1 + \frac{\sqrt{3}(t-t')}{\ell}\right) \exp\left(-\frac{\sqrt{3}(t-t')}{\ell}\right) \quad (8)$$

Furthermore, since both the time series considered here exhibit strong daily periodic patterns, we also use a squared exponential periodic covariance function with unit periodicity given by:

$$k_P(t,t') = \sigma_f^2 exp\left(-\frac{sin^2\pi(t-t')}{2\ell^2}\right) \quad (9)$$

Finally, we also consider the noise in measurements by applying an additive gaussian noise function given by:

$$k_N(t,t') = \sigma_n^2 \delta_{t,t'} \quad (10)$$

where δ is the Kronecker delta and σ^2 represents the noise variance.

3.2.3 Difference Gaussian Process:

In addition to the multi-output formalism presented above, we also propose an alternative approach, whereby we use a conventional single-output GP to predict the point-wise difference between the normalised outputs of T_M and T_C, over the range of T_C^* to recover T_M^*, and then use the values of T_C^* to recover predictions of T_M^*.

3.2.4 Empirical Evaluation:

We evaluate our GP formalisms by comparing them to a number of benchmarks. For carbon intensity, we hypothesise that the value of carbon intensity during any time slot in the future will be equal to the value in that time slot, exactly 1 day, and also 7 days, earlier (reflecting the 1 day and 7 day periodicity observed in the data). We also consider a conventional single-output GP that does not take account of the demand time series, and an approach that determines the correlation between carbon intensity and demand by performing a linear regression between the two, and uses this linear relationship to directly convert a demand prediction into a carbon intensity prediction. In the case of the external temperature, we again consider a single-output GP, the case any time slot is identical to the time slot 1 day earlier, and when we use the online forecast directly.

To perform the evaluation, we use real external temperature sensor data collected from the University of Southampton campus, and also observations and 24 forecasts of Southampton temperature available online through The Weather Channel (`http://uk.weather.com/`). We also make use of carbon intensity and demand observations from the UK grid (for which, again, 24 hour predictions are available - see `http://www.bmreports.com`). All data is at 30 minute intervals, and we perform 31 sequential 24 hour predictions over the entire month of January 2010.

Table 2: Prediction Accuracy

Quantity	Method	RMSE
Carbon Intensity (gCO_2/kWh)	Single-output GP	32.9
	Multi-output GP	22.3
	Difference GP	20.1
	Linear Regression	27.1
	1-Day Repeat	37.7
	7-Day Repeat	78.4
External Temperature (K)	Single-output GP	2.41
	Multi-output GP	2.13
	Difference GP	2.73
	Weather Forecast	2.36
	1-Day Repeat	3.04

We used Bayesian Monte Carlo to marginalise over the hyperparameters of the GP, and we note that the specification of sensible initialisation values for this has not been actively investigated. We employed a principled approach whereby we investigated the autocorrelation properties of the time series under consideration. The knowledge of such correlations also helped guide the choice of appropriate covariance structure. We set the history length for best performance based on empirical observations, and Table 1 shows the covariance functions, and the number of days of historical data, w, used to generate the predictions in this section.

Our results, in terms of the root mean square error of these predictions, are shown in Table 2. We note that the multi-output GP significantly outperform our benchmark approaches on both prediction tasks. This approach is able to effectively exploit both the periodicities within each single time series, and also correlations between them. We also note the while the difference GP works very well in the case of carbon intensity (providing the best predictions overall), it performs less well on external temperature. When the time series are strongly correlated, both approaches provide good predictions. However, when correlations are not present (as is sometimes the case in the external temperature time series) the difference GP has no mechanism to compensate. In contrast, the multi-output GP can adjust the hyperparameter that describes this correlation and perform well. Thus, the difference GP can be an effective approach, but must be used with care since it is not adaptive to changing correlation.

3.3 Estimating Cost and Carbon Emissions

Now, given the external temperature and carbon intensity time series predicted above, the thermal properties of the home learned earlier, and the description of the thermostat, the smart controller can predict the operation of the heater over the course of the day, and hence predict the total cost and carbon emissions that will result from any particular thermostat setting. To evaluate these predictions we have simulated a home and the smart controller installed within it (see Figure 2). In doing so, we use leakage rate $\phi = 90$ W/K and $m_{air} = 1205$ Kg (corresponding to a small well-

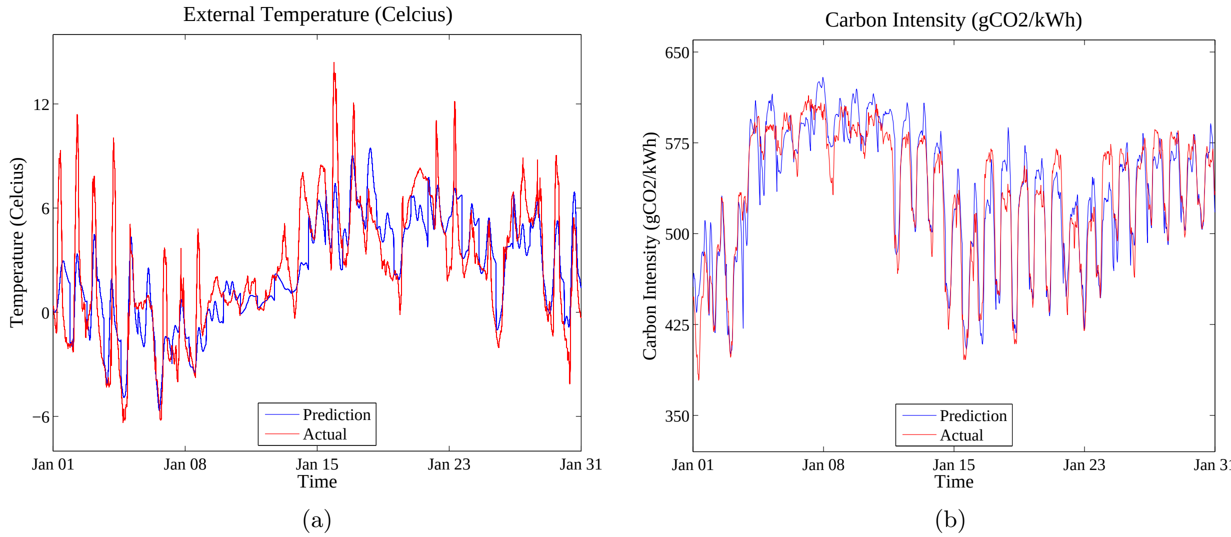

Figure 1: Comparison of Gaussian process predictions for (a) external temperature and (b) carbon intensity using real data for January 2010.

insulated home), an electric heat pump with a thermal heat output of 2.5kW and COP of 2.5, and used a standard value of $c_{air} = 1000$ J/Kg/K in the thermal model. We consider a thermostat set point, $T_{set} = 20$ C, and we assume that the heater is enabled (but not necessarily actively producing heat) between 07:00 and 23:00 each day. Furthermore, we consider a setting in which the price of electricity exhibits a critical pricing period between 15:00 to 20:00 when the price of electricity is £0.24 per kWh, a cheap overnight rate between 20:00 and 06:00 of £0.08 per kWh, and a standard rate of £0.12 per kWh at all other times.

Using the same data for January 2010, we predict the cost and carbon emission for each day, comparing these to the actual values as the true external temperature and carbon intensity is revealed. Over this period, the actual mean daily heating cost is £1.59 and the root mean squared error in our predictions is £0.14, while the actual mean daily carbon emissions are 5.99 kg and the root mean squared error in our predictions is 0.54 kg (an error of 9% in both cases). We observe that in both cases, the prediction errors are less those achieved by simply using the current day as a prediction for the next (in this case, £0.16 and 0.63 kg, respectively). Crucially, our approach allows the effects of changes to the thermostat set-point, or heater timing, to be calculated since it explicitly models these variables. Furthermore, we note that the maximum error occurs at the very start of the day; updating the predictions over the course of the day will cause the predicted cost and carbon to converge to the actual cost and carbon by the end of the day. Figure 2 shows this cost information being displayed on the smart controller's graphical user interface.

4. HOME ENERGY OPTMISATION

Having shown how our smart controller can predict the cost and carbon emissions of the home's heating system, we now show how it can adapt to real-time cost and carbon intensity signals, adjusting the timing of heater use in order to satisfy the home owner's preferences for comfort minimising cost and carbon emissions. Our approach uses mathematical programming to ensure that the smart controller delivers the same level of comfort as a standard thermostat whilst also minimising either cost or carbon. This is attractive since it does not require the householder to explicitly trade-off between cost and comfort (or carbon and comfort), nor indeed, even be aware of the underlying comfort model.

4.1 Evaluating Comfort

We use a standard model of comfort, based on the ANSI / ASHRAE Standard 55-2010, that has previously been used within the smart home environment [3]. To this end, we denote $o_c \subseteq T$ as the set of time slots at which the householder requires comfort, and we define the *instantaneous* discomfort, $\Delta d^t \in \mathbb{R}^+$, such that:

$$\Delta d^t = \begin{cases} \omega_1(\tau_{in}^t - T_{opt})^{\phi_1} & T_{in}^t \geq T_{opt}, \\ \omega_2(T_{in}^t - T_{opt})^{\phi_2} & T_{in}^t < T_{opt} \end{cases} \tag{11}$$

Then, the actual discomfort at time t is a combination of the instantaneous discomforts at t and at $t-1$, such that:

$$d^t = \Delta d^t + \gamma \Delta d^{t-1} \tag{12}$$

where $\gamma \in [0,1]$ scales the effect of the previous time slot on the current one (capturing the psychological persistence of discomfort). The total discomfort over is the sum of the discomfort at every $t \in o_c$ given by given by $\sum_{t\in o_c} d^t$.

4.2 Optimising Heating

Given this comfort model, the smart controller can predict the total discomfort that the thermostat will to deliver over the day. We denote this as the target discomfort, D_{target}. The aim of the smart controller is then to determine when the heater should on in order that the total discomfort is less than the target discomfort with the minimum cost or carbon emissions. More formally, specifying $\eta_{on}^t \in \{0,1\}, \forall t \in T$ as the decision variables, the objective function for cost and

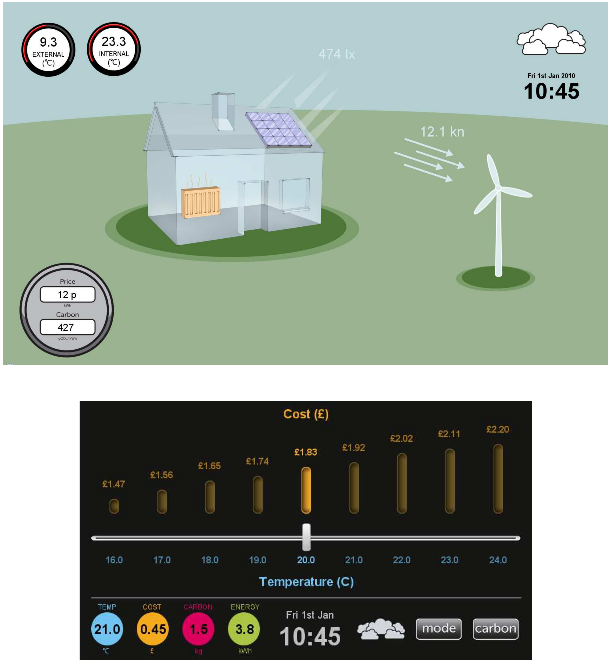

Figure 2: Simulation of the home, the smart controller and its graphical user interface.

carbon respectively, is given by:

$$\arg\min_{\eta^t_{on}} \sum_{t\in T} \eta^t_{on} r_e p^t \Delta t \qquad \text{and} \qquad \arg\min_{\eta^t_{on}} \sum_{t\in T} \eta^t_{on} r_e c^t \Delta t$$

where the evolution of the internal temperature is described by Equation 4, that of discomfort by Equations 11 and 12, and the constraint that $\sum_{t\in o_c} d^t \leq D_{target}$.

4.2.1 *Mixed-Integer Quadratic Programming:*

Now, it is typical to consider that $\phi_1 = \phi_2 = 2$ within Equation 11, and thus the objective function is quadratic. In this case, the optimisation described above can be directly implemented as a mixed-integer quadratic program, using a standard solver, such as IBM ILOG CPLEX.

4.2.2 *Greedy Heuristic Optimisation:*

More generally, if $\phi_1 \neq \phi_2 \neq 2$, or if the computational resources to implement a full mixed-integer quadratic solver are not available on the hardware installed within the home, we also present a greedy heuristic (see Algorithm 1) that incrementally finds the individual time slot where switching the heater on results in the largest reduction in discomfort, for the minimum additional cost or carbon. This process repeats until the total discomfort in equal to, or below, that of the target discomfort.

4.3 Empirical Evaluation

We now compare the performance of the mixed-integer quadratic solver and the greedy heuristic against thermostatic control within the simulation setting previously described. For all experiments, comfort is required between 8am and 11pm, and we use values of $\omega_1 = \omega_2 = 1$, $\phi_1 = \phi_2 = 2$, $\gamma = 0.8$ and $T_{opt} = 23.5$ °C in the comfort model. We consider time

Algorithm 1 Greedy Algorithm to Optimise Home Heating

while $D_{op} > D_{target}$ **do**
  $index \leftarrow 0; bestratio \leftarrow 0$
  for $j = 1$ to $|T|$ **do**
    $\eta_{\mathbf{test}} = \eta_{\mathbf{op}}; \eta^{\mathbf{j}}_{\mathbf{test}} = \mathbf{1}$
    for $t = 1$ to $|T|$ **do**
      Compute T^t_i, Δd^t_{test} and d^t_{test} as per (4), (11) and (12).
    end for
    $D_{test} = \sum_{t\in o_c} d^t_{test}$
    $\Delta D = D_{op} - D_{test}$;
    if $\frac{\Delta d_{test}}{p_t} \leq bestratio$ **then**
      $bestratio = \frac{\Delta d_{test}}{\Delta p}; index \leftarrow j$
    end if
  end for
  if $index > 0$ **then**
    $\eta^{index}_{op} \leftarrow 1; D_{op} \leftarrow D_{test}$
  end if
end while

slots that are 5 minute long, such that $|T| = 288$, we perform 31 sequential optimisations over the month of January 2010. We consider 5 minutes time slots (i.e. $|T| = 288$), and thus, we limit the mixed-integer quadratic solver to 5 minutes of running time on a standard desktop PC (the greedy heuristic runs in milliseconds).

Figure 3 show examples of the optimisation process. In this case, using the MIQP optimisation routine to minimise both cost and carbon. In both cases, we present the internal temperature of the home when the heating is controlled by both the standard thermostat and the smart controller. The green shaded area represents the time interval over which comfort is required, and the red shaded area represents when the heating system is actually producing heat.

Consider Figure 3(a) where the smart controller is minimising cost. Note that the smart controller applies heat before the peak pricing period, allowing the temperature to increase, and then allows this heat to leak away over this period. In contrast, the standard thermostat applies heat uniformly across this period. Similarly, note that the smart controller also exploits the low price of electricity before 06:00 and supplies heat even though it is not immediately required. In both cases, it is effectively storing cheap electricity in the form of hot air, so that this stored energy can be used when electricity is more expensive, and this approach provides an alternative to the use of more costly electrical storage batteries proposed by [11].

Figure 3(b) where the smart controller is minimising carbon. As before, we compare the internal temperature of the home when the heating is controlled by both the standard thermostat and the smart controller. Again the smart controller applies more heat before the peak period when carbon emissions would be the highest, allowing the temperature to increase, and then allows this heat to leak away over this period. In contrast, the standard thermostat again applies heat uniformly across this period (indeed, the operation of the thermostat is identical to the previous case). In addition, note that the smart controller also exploits the lower carbon estimates before 06:00 and supplies heat even though it is not immediately required.

Finally, in Table 3 we compare the total cost and car-

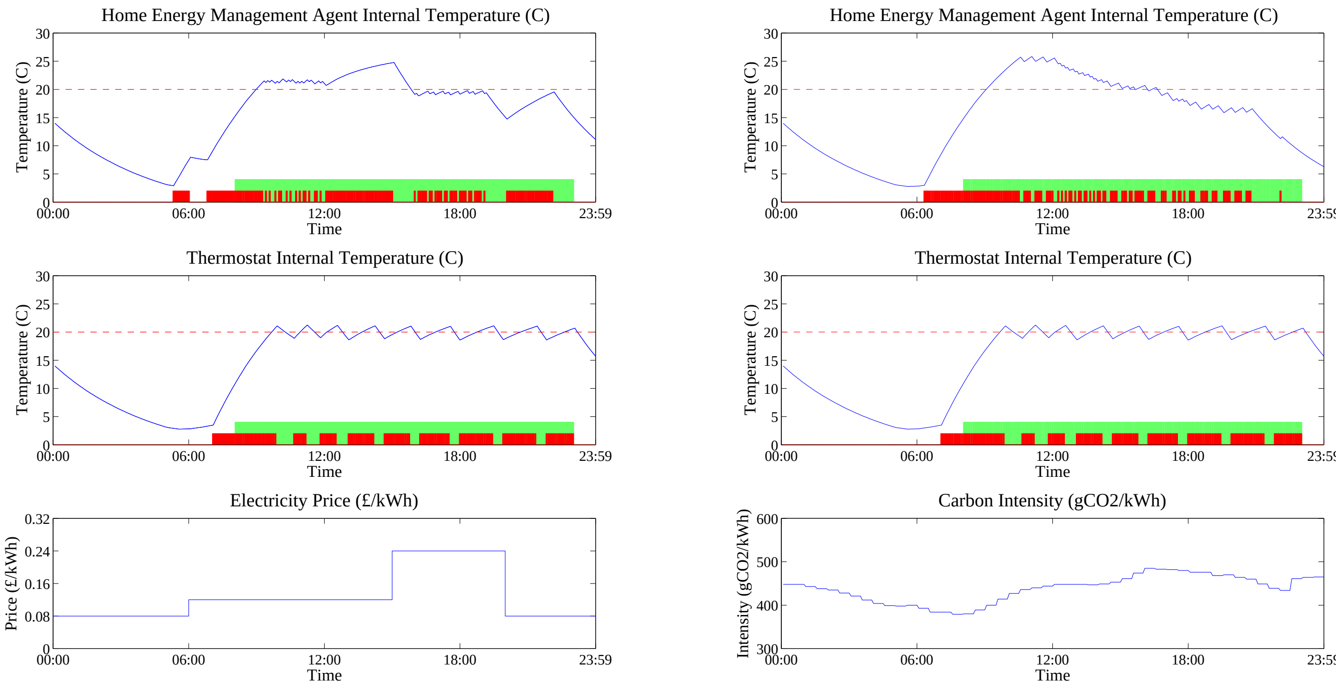

(a) (b)

Figure 3: Example showing the MIQP formulation optimising heating use to maintain comfort whilst (a) minimising cost in a setting with a critical price period between 15:00 and 20:00, and (b) minimising carbon emissions given the variable carbon intensity of the electricity grid.

.

bon emissions of both the MIQP and greedy heuristic over the month of January 2010. Overall, the smart controller reduces cost and carbon emissions by 15%, and 9%, respectively, compared to a standard thermostat, and that the greedy heuristic performs well, being within 1% of the mixed-integer quadratic program (MIQP) solutions (at a fraction of the computational cost).

5. CONCLUSIONS

In this paper, we addressed the challenge of adaptively controlling a home heating system in order to minimise cost and carbon emissions in response to real-time price and carbon intensity signals. We proposed a home energy management agent, or smart controller, that operated on behalf of the householder, and learned the thermal properties of the home, used a Gaussian process model to predict the local environmental parameters, in order to both provide real-time feedback concerning the cost and carbon emissions of their heating preferences, and optimised their energy use. We showed, using real-world data and a simulated home, that our approach could reduce cost and carbon emissions by up to 15%, and 9%, respectively compared to using a standard thermostat.

Our future work is focused on deploying the home energy management agent described here within a number of real homes in collaboration with our industrial partners. The key challenges that we expect to face in doing so, are ensuring that the relatively compuationally expensive prediction and optimisation processes can be performed effectively on the low power devices to be deployed, and to ensure that the thermal model that we have developed is sufficiently representative of a real house, such that it is able to accurate costs and carbon emissions with the necessary accuracy.

Table 3: Comparison of smart and thermostat control when optimising for cost and carbon emissions for January 2010.

	MIQP	Greedy	Thermostat
Cost (£)	42.44	42.81	50.37
Carbon (kg CO_2)	174.60	174.90	191.48

6. ACKNOWLEDGEMENTS

The work presented in this paper was joint funded the iDEaS project (`www.ideasproject.info`), and through the Intelligent Agents for Home Energy Management project (EPSRC EP/I000143/1).

7. REFERENCES

[1] DECC. Smarter grids: The opportunity. Technical report, Department of Energy and Climate Change (DECC), 2009.

[2] DOE. *Grid 2030: A National Vision For Electricity's Second 100 Years*. United States of America Department of Energy, 2003.

[3] Y. Guo, R. Li, G. Poulton, and A. Zeman. A Simulator for Self-Adaptive Energy Demand Management. In *Proceedings of the Second IEEE International Conference on Self-Adaptive and Self-Organizing Systems*, pages 64–73, 2008.

[4] Kempton. Two theories of home heat control. *Cognitive Science*, 10(1):75–90, 1986.

[5] D. Leith, M. Heidl, and J. Ringwood. Gaussian process prior models for electrical load forecasting. In *Proceedings of the International Conference on Probabilistic Methods Applied to Power Systems*, pages 112 –117, 2004.

[6] Low Carbon. The UK low carbon transition plan: National strategy for climate and energy. Technical report, Department of Energy and Climate Change (DECC), 2009.

[7] A. Molderink, V. Bakker, M. Bosman, J. Hurink, and G. Smit. Management and control of domestic smart grid technology. *IEEE Transactions on Smart Grid*, 1(2):109–119, 2010.

[8] M. A. Osborne, S. J. Roberts, A. Rogers, S. D. Ramchurn, and N. R. Jennings. Towards real-time information processing of sensor network data using computationally efficient multi-output gaussian processes. In *Proceedings of the Seventh International Conference on Information Processing in Sensor Networks*, pages 109–120, Washington, DC, USA, 2008.

[9] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2005.

[10] F. Schweppe, B. Daryanian, and R. Tabors. Algorithms for a spot price responding residential load controller. *IEEE Power Engineering Review*, 9(5):49 – 50, 1989.

[11] P. Vytelingum, T. D. Voice, S. D. Ramchurn, A. Rogers, and N. R. Jennings. Agent-based micro-strorage management for the smart grid. In *Proceedings of the Nineth International Conference on Autonomous Agents And Multiagent Systems*, pages 39–46, Toronto, Canada, 2010.

Short Papers

Information Transfer Analysis for Ongoing Commissioning of Buildings

X. Rosalind Wang
CSIRO ICT Centre
PO Box 76, Epping
NSW 1710, Australia
Rosalind.Wang@csiro.au

Joseph T. Lizier
Max Planck Institute for Mathematics in the Sciences
Inselstraße 22
04103 Leipzig, Germany
lizier@mis.mpg.de

Joshua Wall
CSIRO Energy Technology
PO Box 330, Newcastle
NSW 2300, Australia
Josh.Wall@csiro.au

ABSTRACT

Commissioning ensures buildings deliver the performance and energy savings promised by their design. Ongoing commissioning aims to fully automate the commissioning process, providing periodic or even continuous monitoring, identification and possibly remedial actions ensuring persistence in intended operation and performance . One of the major research areas inongoing commissioning at the moment is that of fault detection and diagnosis (FDD) of heating, ventilating, and air conditioning (HVAC) systems. First generation technology mainly focuses on detection of faults in sensors and actuators with manual diagnosis and remedial actions performed by system experts. In this paper, we propose an information theory method to analyse the information transfer between time series of the sensors and actuators (or the agents) of a HVAC system. The information transfer can help in the diagnosis of faults in the system by automatically identifying the related agents that caused the particular faults in another agent. We present some preliminary results from the proposed method using real data gathered from the a commercial office building in Newcastle, Australia.

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: Miscellaneous

General Terms

Algorithms, Experimentation

Keywords

information dynamics, transfer entropy, energy efficiency, ongoing commissioning

1. INTRODUCTION

Literature on related building case studies suggest that virtually all buildings have some sort of *heating, ventilation, and air conditioning* (HVAC) operational problems, and the vast majority of buildings are not carefully commissioned [2, 11]. With considerable potential to achieve persistent energy savings at 'negative' cost, effective commissioning, post-occupancy evaluation, and *fault detection and diagnosis* (FDD) procedures could provide the key to immediate and significant emission cuts from the built environment.

We are interested in the problem of automating the FDD process for office buildings. Fault detection and diagnostics tools allows us to address many of the barriers to good commissioning and maintenance practice and consequential energy savings. Techniques from scientific disciplines such as artificial intelligence and machine learning [5, 6] are emerging as viable approaches where rules-based techniques may be less suitable.

Many current approaches focus on individual components, while this gives good fault detection algorithms, it does not allow for automatic fault diagnosis. In this paper, we examine the information transfer [8] between sensor and actuators (agents) within a building HVAC system. Information transfer is a recently published framework that characterises the information dynamics of distributed computation in terms of the elements of Turing universal computation. While information transfer is not a direct measure of causal effect [1, 7], it can be used for simple inference about causal relationships between two sets of measured data within the HVAC system. By knowing these relationships, we can then better diagnose faults by looking at those in related systems.

We begin the paper with an overview of building commissioning, HVAC systems and the continuous challenges in the field in Section 2. Section 3 introduces the framework for the information dynamics of distributed computation that was used. The experimental setup and data is shown in Section 4. We then present and discuss the results of information dynamics within the building energy sensors in Section 5. Finally, we discuss some future direction of this work and a conclusion.

2. ONGOING COMMISSIONING

The aim of commissioning new buildings is to ensure that they deliver, and sometimes exceed, the performance and energy savings promised by their design. When applied to existing buildings, commissioning identifies the almost inevitable "drift" from where things should be and puts the building back on course. Commissioning is a 'stealth' strategy for saving energy and money and reducing greenhouse gas emissions. Commissioning is also a risk-management strategy that ensures the building owners get what they paid for. [10]

Since each building is unique, the commissioning process would also be different for each new building. Commissioning is performed at widely different levels of effort and applied to buildings as a whole or a specific sub-system/end-use. Commissioning complements the energy efficiency practices by ensuring and maintaining building energy performance (and other benefits, such as indoor environmental quality). On the same token, it can simply focus on saving energy by improving conventional building systems, irrespective of whether or not the building is equipped to be particularly energy efficient.

In this paper, we focus on automated ongoing commissioning of the HVAC systems. Ongoing commissioning is a commissioning process that takes place in the building periodically. This is necessary since commissioning and retro-commissioning fixes do not persist over time [4].

The process of ongoing commissioning has two major steps: detection-diagnosis and repair. An automatic FDD system for HVAC will identify uncalibrated or malfunctioning sensors, valves, dampers and actuators; etc [4]. The repair step can be automated for tasks such as recalibration, new control strategies, and new set-points; or manual for tasks such as repair and re-installation of equipments. We are interested in the problem of automating the FDD process.

Fault detection and diagnosis are two separate processes. Detection aims to find data that deviates from the 'normal' behaviour of a system, while diagnosis is required to deduce the cause of the deviation. Automated FDD has become increasingly important in science and technology, however, it is a relatively recent discipline in building operation. Only within the last decade or so has the field been researched. A good review of the subject has been published by Katipamula and Brambley [5, 6].

FDD in building can be applied at various levels: from individual components such as heating coils, to specific systems such as the heating systems with all the components, to to the entire building, taking into account all systems such as heating, cooling, lighting, etc. Ideally, applications would complement each other, however, the current situation is far from achieving this goal.

Most FDD methods available now are designed and developed for individual components whose design and thermodynamic behaviour are already know. Special FDD methods are developed on the basis of individual faults. This detailed knowledge of the system is typical of FDD at component level and distinguishes it in principle from FDD operating at system or building level [3].

One reason for most of the research up until now to focus on individual components is that at the macro level of the entire building, most of the parameters are often not known. An example of these parameters is the effects of the occupants' behaviour on energy consumption in a building. At the moment, an accurate representation of the system or of the individual faults is only possible to a limited extent. Therefore, in more complex systems, fault analysis is often carried out manually by an expert. In this paper, we aim to use information theory to discover the relationships between various sensors and actuators within the HVAC system, thus allowing the FDD to operate at a system level, and ultimately, the building level. Knowing the relationship between various sensor and actuator's time series data will also allow the FDD to diagnose faults quickly and autonomously.

3. INFORMATION TRANSFER

Information theory [9] is an increasingly popular framework for the study of complex systems. In part, this is because complex systems can be viewed as distributed computing systems, and information theory is a natural way to study computation, e.g. [8]. Information theory is applicable to any system, provided that one can define probability distribution functions for its states. This is a particularly important characteristic since it means that information-theoretic insights can be directly compared across different system types.

The fundamental quantity for information theory is the (Shannon) *entropy*, which represents the uncertainty in a sample x of a random variable X. Shannon Information [13] was originally developed for reliable transmission of information from a source X to a receiver Y over noisy communication channels. When dealing with outcomes of imperfect probabilistic processes, it is useful to define the information content of an outcome x, which has the probability $P(x)$, as $-\log_2 P(x)$. Crucially, improbable outcomes convey more information than probable outcomes. Given a probability distribution P over the outcomes $x \in \mathcal{X}$, the average Shannon information content of an outcome is determined by

$$H(X) = -\sum_{x\in\mathcal{X}} P(x)\log_2 P(x), \tag{1}$$

Intuitively, Shannon information measures the amount of freedom of choice (or the degree of randomness) contained in the process — a process with many possible outcomes has high entropy.

Transfer entropy (TE) [12], a measure to formulate the information transfer based on information theory, can be useful in ongoing commissioning work. The information transfer between a source and a destination agent is defined as the information provided by the source about the destination's next state that was not contained in the past of the destination. While information transfer is not a direct measure of causal effect [1, 7], it can be used for simple inference about causal relationships between two sets of measured data within the HVAC system. For example, if the temperature within a zone is directly dependent on the speed of the fan, then there should be a large amount of information transferred from the fan speed to the temperature data.

Transfer entropy does not require a model of the interaction and is inherently non-linear. Transfer entropy addresses concerns that the mutual information, $I(X;Y) = H(X) - H(X|Y) = H(X) + H(Y) - H(X,Y)$, as a de facto measure of information transfer was a symmetric measure of statically shared information. The transfer entropy from a source agent Y to a destination agent X is the mutual information between the previous state of the source y_n and the next state of the destination x_{n+1}, *conditioned* on the semi-infinite past of the destination $x_n^{(k)}$ (as $k \to \infty$ [8]):

$$T_{Y\to X} = \lim_{k\to\infty} \sum_{\mathbf{u}_n} p(\mathbf{u}_n) \log_2 \frac{p(x_{n+1}|x_n^{(k)}, y_n)}{p(x_{n+1}|x_n^{(k)})}, \tag{2}$$

where $\mathbf{u}_n$ is the state transition tuple $(x_{n+1}, x^{(k)}, y_n)$, $T_{Y\to X}(n,k)$ represents finite-k approximation.

4. EXPERIMENTAL SETUP

The data used for this paper was gathered at the Australian Commonwealth Scientific and Research Organisation (CSIRO)'s Energy Centre in Newcastle, New South Wales. The building was completed in 2003, with a total air-conditioned area of 3320 m^2 divided into 15+ zones on 4 levels. A HVAC zone is a group of adjacent offices and/or

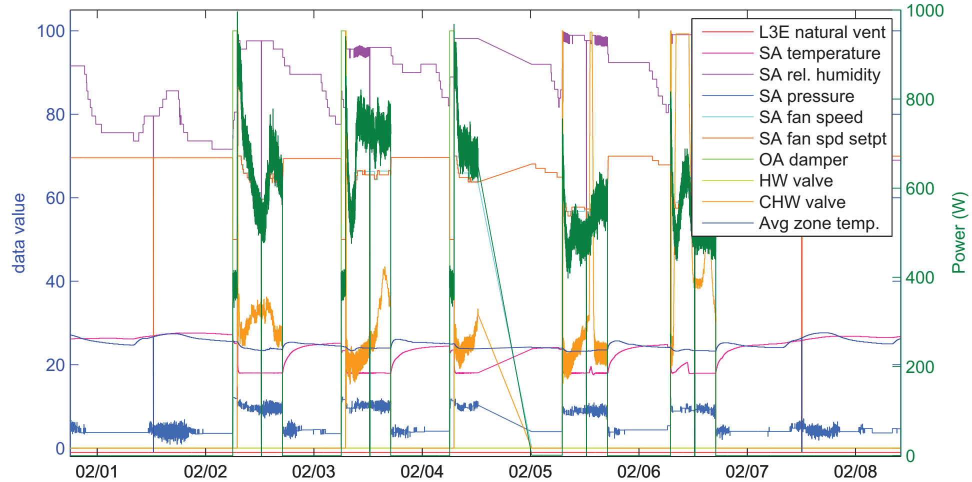

Figure 1: A example of the data used for the experiment, the days shown here represents a week from Sunday 1st to Sunday 8th.

spaces serviced by a common air-handling unit (AHU) or air-terminal device. Air is delivered to each zone via underfloor supply air plenums and floor air terminals. Return air is circulated back to the plant rooms via high level grilles. Space temperature control is provided through modulation of supply air temperature and and fan speed. The HVAC system accounts for approximately 57% of the total building electrical load.

A number of measurements have been conducted to gather historical data on Level 3 (East) office wing, in which three AHUs (AHU 08, 09 and 10) service three adjacent HVAC zones consisting of 18 offices and associated corridors. For this paper, we used the data gathered for AHU-10 (which services 8 adjacent offices) along with one set of data for level 3. The sensors and actuators used were:

1. Level 3 East natural ventilation on/off,
2. AHU-10 fan power (W),
3. AHU-10 supply air temperature (°C),
4. AHU-10 supply air relative humidity (%),
5. AHU-10 supply air pressure (Pa),
6. AHU-10 supply air fan speed (%),
7. AHU-10 supply air fan speed set point (%),
8. AHU-10 outside air damper (%),
9. AHU-10 hot water valve (%),
10. AHU-10 chilled water valve (%),
11. AHU-10 average zone temperature (°C).

We took the data collected between 0900hr 8th Jan and 1150hr 9th Feb 2009, the data was sampled at 5 seconds interval. It should be pointed out that this takes place in the middle of Australian summer, thus the HVAC system is mainly engaged to cool the building. Figure 1 shows a portion of the data used, the dates here represents a week in February.

5. RESULTS AND DISCUSSION

To calculate the transfer entropy between pairs of data sets, we set the history size of $k = 2$. There are a total of approximated 336000 data points in each set of data, therefore we used a kernel width of 1 for kernel estimation with a normalised kernel.

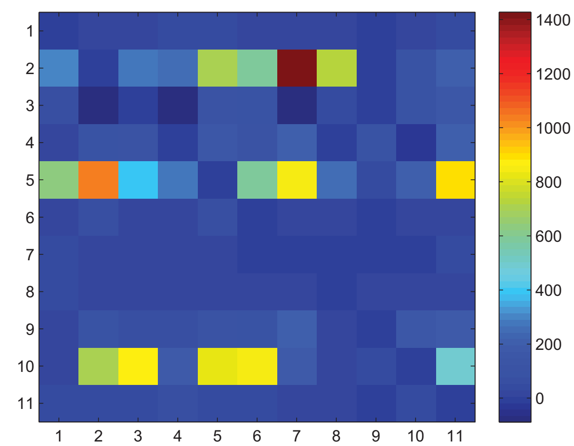

Figure 2: Transfer entropy calculated for the data sets described in Section 4. The rows represent the *destination* agent, and the columns represent the *source*.

Figure 2 shows the resulting total transfer entropy between pairs of the data sets as described in Section 4. Each row represents the destination agent of the transfer and each column represents the source agent. For example, cell (2, 7) represents the transfer entropy from data set 7 (which is AHU-10 supply air (SA) fan speed set point) to data set 2 (which is AHU-10 fan power). The units of shown in the image is bits.

One of the first feature to notice on Figure 2 is that it is not symmetrical. This is to be expected, since the causal relationship should be directional. So, while we expect the supply air's fan speed to affect the supply air temperature, we don't expect the reverse. This was the reason that mutual information isn't used in determining the causal relationship as it is unidirectional.

Another expected result that we see is that the diagonal is zero, since we wouldn't expect any transfer within one data source. This does not mean that an agent's past state doesn't affect it's future, rather transfer entropy calculates the information provided by a source about a destination's next state that was not contained in the past of the destination. Therefore, if a source is also the destination, then the terms $p(x_{n+1}|x_n^{(k)}, y_n)$ and $p(x_{n+1}|x_n^{(k)})$ in Equation 2 will be equal, thus TE is zero.

The highest transfer entropy shown in Figure 2 is that in cell (2, 7), indicating a positive transfer of information from the SA fan speed set point to the fan power. This might seem surprising at first since set points are usually constant for a long time and do not change often. However, looking closely at the data we have, we found that the set point in our case does get altered regularly, presumably to cool the building as outside temperature rose.

Of the other systems' relationship with fan power, we see that the water valves have very little transfer entropy, which is reasonable. The average temperature of the zone also doesn't have much effect, due to the fact we are only measuring one past step of the source agent, as it will take a lot more than 5 seconds (which is the time interval used in the data) for the average zone temperature to affect the fan power.

The other time series that was affected by many others was the SA pressure (row 5 in Figure 2. We can see from the results that it was most notably affected by (in order of transfer entropy value) fan power, average temperature of the zone, SA fan speed set point, and whether the natural ventilation of the floor was on or off which means whether the air conditioning system was on or off.

There are some relationships in Figure 2 that were expected to appear but didn't. For example we expect that the position of the chilled water valve will affect some of the time series such as the average zone temperature. This is due to the fact that the data was gathered at 5 second interval and we only used one step of the source agent's data, which would not be enough time for effects to be noticeable. Another reason for these none-appearance is that the data was not pre-processed before the TE calculations were made. This is especially a problem for systems that are turned off overnight, for example the fan power, as we see in Figure 1. Therefore, when we calculated the total TE using all the data, a lot of noise was introduced to the results, thus covering up some of the features that could showed up.

6. CONCLUSIONS AND FUTURE WORK

In this paper, we measured the information transfer between time series data of various sensors and actuators in an HVAC system. The amount of information transfer can be used for simple inference about the causal relationship between two sets of data. This relationship can then be used to allow us better diagnose faults detected within the system.

We found with our preliminary investigation that for most pairs of time series, the measured information transfer provides a good indication of the causal relationship between the two. However, we also found with other sets, the expected causal relationship were not present in the results. There were two main reasons for those: firstly the data was recorded at a 5 second interval, which is too little time for expected effect to appear; secondly, as the data was not preprocessed, some time series still has a lot of noise in the data, which obscured the real features in the data. In future works, we would like to look into having a longer history length for the source agent's time series; and pre-process the data to eliminate the noise before calculating the transfer entropy. We also aim to apply this method to similar data from other buildings.

7. ACKNOWLEDGMENTS

We like to thank the High Performance Computing and Communications Centre (http://www.hpccc.gov.au/) for the use of their supercomputer clusters in performing the experiments for this paper.

8. REFERENCES

[1] N. Ay and D. Polani. Information flows in causal networks. *Adv. Complex Syst.*, 11(1):17–41, 2008.

[2] D. E. Claridge, J. Haberl, M. Liu, J. Houcek, and A. Athar. Can you achieve 150% of predicted retrofit savings? is it time for recommissioning? In *Proc. ACEEE 1994 Summer Study on Energy Efficiency in Buildings*, volume 5, Aug. 1994.

[3] L. Jagemar and D. Olsson. The EPBD and continuous commisioning. Technical report, CIT Energy Mangement AB, Göteborg, Sweden, Oct. 2007.

[4] S. Katipamula. Automated continuous commissioning of AHUs. In *ASHRAE Winter Annual Meeting*, Chicago, USA, Jan. 2003.

[5] S. Katipamula and M. R. Brambley. Methods for fault detection, diagnostics, and prognostics for building systems - a review, part i. *HVAC&R Research*, 11(1):3–25, Jan. 2005.

[6] S. Katipamula and M. R. Brambley. Methods for fault detection, diagnostics, and prognostics for building systems - a review, part ii. *HVAC&R Research*, 11(2):169–187, Apr. 2005.

[7] J. T. Lizier and M. Prokopenko. Differentiating information transfer and causal effect. *EPJ B*, 73(4):605–615, 2010.

[8] J. T. Lizier, M. Prokopenko, and A. Y. Zomaya. Local information transfer as a spatiotemporal filter for complex systems. *Phys. Rev. E*, 77(2):026110, 2008.

[9] D. MacKay. *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press, 2003.

[10] E. Mills. Building commisioning: A golden opportunity for reducing energy costs and greenhouse gas emissions. Technical report, Lawrence Berkeley National Laboratory, July 2009.

[11] M. A. Piette and B. Nordman. Costs and benefits from utility-funded commissioning of energy-efficiency measures in 16 buildings. *ASHRAE Transactions*, 102(1):LBL–36448, 1996.

[12] T. Schreiber. Measuring information transfer. *Phys. Rev. Lett.*, 85(2):461–464, 2000.

[13] C. E. Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27:379–423, 623–656, July, October, 1948.

Dynamic coalition formation and adaptation for virtual power stations in smart grids

Radu-Casian Mihailescu
University Rey Juan Carlos
Madrid, Spain
radu.mihailescu@urjc.es

Matteo Vasirani
University Rey Juan Carlos
Madrid, Spain
matteo.vasirani@urjc.es

Sascha Ossowski
University Rey Juan Carlos
Madrid, Spain
sascha.ossowski@urjc.es

ABSTRACT

An agent-based organizational model for a smart energy system is introduced relying on a dynamic coalition formation mechanism for virtual power station (VPS) creation. Following, for VPS maintenance a solution concept is proposed that stems from the existent stability solutions in coalitional games and is introduced in conjunction with a corresponding algorithm, deployed in distributed environments populated by negotiating agents. The algorithm is intended as an open-ended organizational adaptation, concerned with achieving stable configurations that meet the desired functionalities within stochastic scenarios.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence *Coherence and coordination, intelligent agents, multiagent systems*

General Terms

Algorithms, Design

Keywords

virtual power stations, distributed generation, multi-agent systems, coalition formation, emergent organization

1. INTRODUCTION

In recent years, there is an increasing interest in the integration of distributed, small-scale, renewable generation into the power system. An efficient use of distributed energy resources (DER) may increase the flexibility and the resilience of the power system at distribution level. Furthermore, it is possible to reduce the dependence from large-scale, non-renewable, power plants and therefore to contribute to a sensible reduction of CO_2 emissions. According to the US department of Energy, a 5% increase in grid's efficiency is equivalent to the fuel and CO_2 emission of 53 million cars.

The potential allure of the multiagent system paradigm (MAS) to the power industry has been extensively documented so far [1]. To this respect, several management systems have been proposed for the organization of micro-grids [3]. A micro-grids is defined as a subsystem of the distribution grid, formed by generation, storage and load device, interconnected at the electrical and the informational level. Micro-grids can be intended as a systemic approach to realize the emerging potential of distributed generation.

Setting aside from this approach that aims at imposing an architectural control, whether centralized or not, on already predefined micro-grids, our vision is intended at proposing a method for congregating the smart-grid actors to dynamically approximate optimal micro-grid configurations. Thus, the main concern is developing techniques aimed at establishing and adapting a suitable virtual organizational structure. The procedure is designed such that it develops on the concept of integrating DERs in the form of virtual power stations [5]. A virtual power station (VPS) is conceived as a bundle of DERs that are connected through an informational infrastructure and act in a coordinated way as a single entity. The challenging problem related to the implementation of the VPS concept is the distributed control of the DERs, mainly due to the stochastic behaviour of the system and the heterogeneity of the devices involved.

The aim of this work is modelling the coordination of virtual power plants in the sense of coalitional games. Instead of considering centralized architectures [4, 5], we claim that a dynamic, bottom-up, approximation of optimal VPS configurations is more effective to ensure flexibility and robustness to the system.

2. PROBLEM REPRESENTATION

The coalition games we aim to address in our approach are the ones where the coalition formation problem is projected on an underlying network topology, given that the cost for cooperation also plays a major role. This class of games is primarily characterized by non-superadditivity, in the sense that gains resultant from forming coalitions are limited by the actual cost of coalition formation[1] and coordination, thus the grand coalition is seldom the optimal structure. Furthermore, the coalitional game is subject to the dynamism of the environment. The challenge is to develop mechanisms that permit large number of autonomous agents to collectively achieve a desired functionality by permanent adaptive dynamics.

In our scenario, we set to investigate the integration of renewable energy resources to the grid in the form of virtual power stations by means of aggregating the power generating potential of various devices in a novel way in the context of MAS. As system designers, we choose to enable the autonomous agents with the basic coordination primitives, and leave to the agents to self-organize and coordinate as the situation may demand, in a fully distributed manner.

[1]The cost of forming a coalition can be perceived through the negotiation process and information exchange which incur costs.

We model the problem as a dynamic coalition formation game with the following formalization:

Let $M = \langle \mathcal{A}, \beta_i, \mathcal{S}, \Phi, v \rangle$ be a multi-agent system where:

- $\mathcal{A} = \{a_1, a_2, ..., a_n\}$ represents the set of agents of a given portion of the distribution grid. We assume that each stakeholder that is connected to the grid is represented by a software agent that manages the corresponding device (e.g., generators, storage devices, intelligent loads).
- β_i is the forecast amount of electricity for the following day associated with agent a_i. If $\beta_i > 0$, then agent a_i is a *provider*, whilst if $\beta_i < 0$ then agent a_i is a *consumer* (or load). Let $\mathcal{P} \subseteq \mathcal{A}$ denote the set of providers, and $\mathcal{L} \subseteq \mathcal{A}$ the set of consumers. In this work we assume that an agent is either a provider or a load, and therefore $\mathcal{P} \cup \mathcal{L} = \mathcal{A}$, $\mathcal{P} \cap \mathcal{L} = \emptyset$. We will refer onwards generically, to an agent belonging to set $\mathcal{P}$ as PA, and to an agent belonging to set $\mathcal{L}$ as LA.
- $S = \{S_1, S_2, ..., S_m\}$ is the set of coalitions that partition the set $\mathcal{A}$. We assume that all coalitions are disjoint, and therefore:

$$\bigcup_{j=1}^{m} S_j = \mathcal{A}, \quad S_j \cap S_k = \emptyset, \forall\, j \neq k$$

- $\Phi = \{\phi_1, \phi_2, \phi_3\}$ is a set of constraints that must hold for every coalition. In this work, we enforce that the number of members of each coalition does not exceed a predefined value N, which corresponds to the safety limit imposed by technological constraints (ϕ_1). We also want that each coalition is able to supply electricity to all the loads, so as the energetic balance between generation and consumption must be positive (ϕ_2). Finally, each coalition must realise a desired generation profile of electricity that would qualify them as VPS (ϕ_3). Formally:

$$\phi_1 : |S_j| \leq N \quad \forall j \in \{1, ..., m\}$$

$$\phi_2 : \sum_{a_i \in S_j} \beta_i > 0$$

$$\phi_3 : \sum_{a_i \in \mathcal{P}_j} \beta_i = \psi$$

  where $\mathcal{P}_j = \mathcal{P} \cap \mathcal{S}_j$, and ψ represents the desired energetic profile that each coalition wants to achieve.

- $v : \mathcal{S} \to R$ is a function that for every coalition of the set $\mathcal{S}$ returns its utility value, representing a multi-criterion evaluation of domain specific parameters.

The goal of the coordination problem is obtaining a partitioning of $\mathcal{A}$ into a coalition structure $\mathcal{S}$ that complies with the set of constraints Φ and at the same time maximises the social welfare[2] of the system, without jeopardizing the functionality of any of the coalitions. We leave aside for now what this trade-off implies and further develop on this issue in Sections 5.

[2]For computing the social welfare of the system we mean the typical interpretation of averaging over the utilities of all coalitions.

3. COALITION FORMATION MECHANISM

Our proposed approach is intended as a decentralized and dynamic method. Here coalition formation is achieved by opportunistic aggregation of agents, while maximizing coalitional benefits by means of taking advantage of local resources in the grid. It consists of three phases.

- *Coalition initiation.* To accomplishing this task we have chosen a straightforward approach: the PAs whose energy availability exceeds a predefined threshold value are entitled of establishing themselves as VPS initiators and will do so with a probability p inversly proportional to the number of the agent's PA uncommitted neighbors that are also set to do so.

- *Provider aggregation phase.* This procedure dynamically constructs the organizational structure via a negotiation mechanism, where proposals and requests are handled opportunistically. The initiator PA assumes the role of CA (coordinator agent) for the emergent coalition. The mechanism proceeds in a Contract Net-like fashion with the following steps: CAs send requests to their neighboring PAs indicating the VPS profile that they want to realize, in terms of scale and energetic potential [3]. Based upon these specifications PAs evaluate CAs' offers and select the preferred choice from their candidate set. Eventually, CAs receive offer responses and decide committing PAs.

  The CA's decision is based on calculating an association coefficient that reflects the self-sufficiency of the potential coalition i assuming the joining of actor j. This is denoted as a linear combination of the parameters that apply for this particular case study: security measures (which implies computing contingency analysis - C_j), transmission costs (T_j) and energetic balance (E_j):

$$\varphi_{i,j} = w_1 \cdot C_j + w_2 \cdot T_j + w_3 \cdot E_j \tag{1}$$

  We note that the first two parameters (of negative values) are ought to give an indication of the network's capability of avoiding line overloads and incurred transmission costs. The former gives an indication of the impact that the integration of the LA's load would produce on the system's buses in terms of the percentage of capacity drop. By transmission costs we imply the effect of the power loss in the course of transmission, over the coalition's total power, represented by this ratio. The last parameter represents the percent increase or decrease of the coalition's energetic balance, given the desired state of offer matching demand. Formulating the association coefficient in this manner allows placing more emphasis on the contingency and transmission coefficients at the beginning of the aggregation of actors, while the energetic balance gains more significance proportionally to the number of actors involved. This is obviously a key aspect since one cannot expect attaining a reasonable energetic balance at the very

[3]The constraints imposed on the coalition formation process may vary according to the desired feature of the emerging VPS.

beginning of the coalition formation process. The utility function of the coalition, v, represents the sum of the association coefficients for the coalition members, as shown in Equation 2.

$$v(S_i) = \sum_{j \in S_i} \varphi_{i,j} \qquad (2)$$

- *Consumer aggregation phase.*

  Once the VPS energetic potential has been ensured, the only remaining phase requires the aggregation of LAs. This operates in a similar manner. LAs proceed by submitting their forecasted demand to the CAs in their proximity. For each such request, the CAs calculate the corresponding association coefficient. Based on this information, a proposal response is being awarded to the most suitable LAs for joining the coalition, by enclosing its corresponding association coefficient. LAs will conclude the procedure by selecting the coalition whose utility is mostly increased by their commitment. The decision is unequivocally accepted by the CA since it comes as a response to its precedent proposal and it is exclusively addressed to the CA. The LAs' preference for acting in this sense is justified by the fact that the utility of the coalition would have a direct effect on its reliability.

  To be noted that the decision of selecting the best candidate is carried without a complete representation of the environment, but rather based on local and possibly incomplete knowledge. This is the underlying reason for the evolutionary nature of the algorithm that iteratively approximates a solution through refinement steps.

4. COALITION SELF-ADAPTATION

The second important issue we tackle regards the notion of stability that the system is capable to achieve, given the dynamism of the environment. Thus, the problem we are facing in open organizations requires a modification of the coalition structure due to the variations occurring within the system. All actors submit on a daily basis their forecasted profile, which typically does not differ exceedingly from their previous one. Nevertheless these cumulated variations would entail a reorganization of the coalitions for the following forcasted period in order to assure enhanced coordination at the coalition level. Therefore, consequent to calculating the energetic balance of the coalition, it is to be determined the actors that would qualify to be signed-off, or the profile of the actors that would be eligible to be signed-in to the coalition. The association network generated during the coalition formation process, revealing the existing interdependencies within the coalition, plays a key role at this stage. The weakest links signify the actors the coalition is least dependent on, based on which agents are to be proposed for being signed out of a coalition.

With these considerations in mind we seek a notion of equilibrium that intrinsically provides an argumentation scheme, which allows for a reorganization of the coalition structure. Furthermore, the solution concept should reflect the decentralization outlook of our scenario, while minimizing the structural adaptations by providing a minimum set of interaction rules in order of attaining the desired stability properties amongst negotiating agents.

The solution we propose can be directly referenced to game-theoretic approaches on issues of stability and negotiation. For further considerations on notions of stability, their strength, limitations and interrelations we refer the reader to [2]. In our scenario, of utmost importance is the agents' capability to coordinate and reorganize into groups or coalitions by transforming traditional game-theory criterions of stability towards operating in dynamic environments. Moreover, we advocate for reorienting game-theory to accomodate situations where coordination is a more natural operational descriptor of the game rather than simply self-interested settings. As it is emphasized in [2], an equivalent formulation for solution concepts can be in interpreted in terms of objections and counter-objections. More formally, let x be an imputation in a coalition game with transferable payoff($\mathcal{A}, v$), we define our argumentation scheme as follows:

- (S, y) is an objection of coalition S to x against T if S excludes i and $e\left(S \cup \{i\}, x\right) > e\left(S, y\right)$
- coalition T counteracts to the objection of coalition S against accepting player i if $e\left(T \cup \{i\}, y\right) / e\left(T, x\right) > 1 + \mu$ or $e\left(T \cup \{i\}, y\right) + e\left(S, y\right) > e\left(T, x\right) + e\left(S \cup \{i\}, x\right) - \tau$.

To correlate the game-theoretic terms introduced above to our setting, we give the interpretation of these terms for our scenario. Specifically, by imputation we mean the distribution of utilities over the coalitions' set, whereas the excess e of each coalition represents the difference between its potential maximum (see Section 2) and its current utility v

We reason that the excess criteria applied for solution concepts such as the kernel and the nucleolus appears to be an appropriate measure of the coalition's efficiency, especially in games where the primary concern lies rather in the performance of the coalition itself. This basis further advocates for argumentation settings where objections are raised by coalitions and not by single players, such as the case of the bargaining set or the kernel. The objection (S, y) may be interpreted as an argument of coalition S for excluding i resulting in imputation y where its excess is being decreased. Our solution models situations where such objections cause unstable outcomes only if coalition T to which the objection has been addressed fails to counterobject by asserting that S's demand is not justified since T's excess under y by accepting i would be larger than it was under x. Such a response would have hold if we simply presumed players to be self-interested and not mind the social welfare of the system. If on the contrary, players are concerned with the overall efficiency of the system, they would consider accepting the greater sacrifice of y in comparison to x only if this would account for an improvement of S that exceeds the deterioration of T's performance by at least the margin τ. Thus, τ is the threshhold gain required in order for justifying the deviation, whereas μ represents S 's tolerance to suboptimal gains.

For applying this solution concept to our setting, we additionally need to take into account the underlying topology and thus restrain the inter-coalition argumentation to the given network structure, representing a particularization of the more generic outline presented herein. When multiple objections are being adressed to one coalition, its decision of considering one would be based on the criteria of maximizing parameter τ, while minimizing parameter μ. Also worth

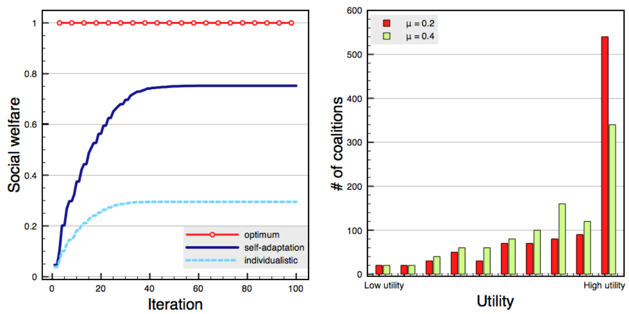

Figure 1: a) Social welfare of the system in number of steps. b) Histogram representation for the utilities of the coalitions

reminding is that the procedure is ought to occur within the domain dependent constraints, that impose maintaining the profile of the coalition within certain limits (see Section 2). Finally, we stress that the aim for our proposed scheme is intended towards an open-ended organizational adaptation concerned with achieving stable configurations in dynamic environments where one-shot optimization procedure are unapplicable.

5. EXPERIMENTAL RESULTS

To begin with, we first evaluate the performance of our algorithm attained through the argumentation scheme introduced. Given the cooperative scenario reflected by our chosen solution concept we have set aside from the Pareto optimal instance[4] where self-interested agents agree to participate in a trade if and only if the contract increases the agent's immediate payoff.

Figure 1(a) points out the average percent increase in social welfare, that the system manages to attain from an initial state to a stable one, achieved during the course of the adaptation phase. A stable configuration of the system is abruptly reached, meaning that the agreements realized earlier improve the social welfare more than the ones performed later. Furthermore, the solution applied is an anytime algorithm that achieves a monotonic improvement of the global (social) welfare of the system

Subsequently we perform a series of experiments to draw more insight to the solution concept introduced. On one hand, our negotiation scheme implies that deviations would only occur if a certain minimum gain τ has been achieved. On the other hand, the extent to which a coalition is willing to decrease its efficiency in detriment of the gain in social welfare is represented by a satisfactory parameter μ. This represents in effect a percentage, which defines what an acceptable performance would be and how tolerant is one coalition towards suboptimal performance. For our simulations we chose an initial value of 0.4 and considered a homogenous population of actors in the system. Although this does not make the objective of our scenario, heterogeneity amongst the actors involved may as well be introduced.

Following, we analyzed the implications of the dependency on this parameter for a better understanding of its functionality. Thus, we have analyzed the stationary states the system falls into as a function of μ. For large values of μ, meaning that coalitions are willing to significantly decrease their utility with respect to the improvement of the global welfare of the system, we encountered an expected global increase in utility, but a considerable variation in the allocation of utilities in the system. Figure 1(b) illustrates a histogram representation of the coalitions' utilities discretised in increasing order of their worth. It can be seen that a 20% increase of μ reduced significantly the number of coalitions operating at high efficiency denoted by the last column of the histogram, while the number of coalitions operating at lower levels of efficiency has been increased. The results emphasized that the best performance of the system was obtained for values of μ in the vicinity of 0.2 . Somewhat surprisingly, what the experiments show is that being willing to accept lower efficiencies in the benefit of the global performance is only advantageous to a certain extent. In actuality, there is a trade-off to be taken into account. Although the overall system utility increases, the ratio between the number of coalitions with low utility and those with high utility is increasing as well. So, for assesing the efficiency of the system not only should we be interested in the global utility, but also in having a uniform distribution of high utilities for the majority of the coalitions.

Future work will further look to a greater extent at the electrical features of the power system and incorporate in a more factual form load-flow computation analyses, that verifies for contingencies and maintains the system within its operational limits.

6. CONCLUSIONS

As a proof of concept our work has introduced a dynamic coalition-based model deployed in distributed environments of negotiating agents. The formation and adaptation mechanism introduced perform an open-ended adaptation of groups of organizational agents, converging towards stable configurations. In particular, we have highlighted the applicability of this approach through the design of a distributed adaptive scheme for the smart electricity grid.

7. REFERENCES

[1] S. D. J. Mcarthur, E. M. Davidson, V. M. Catterson, A. L. Dimeas, N. D. Hatziargyriou, F. Ponci, and T. Funabashi. Multi-agent systems for power engineering applications - part i: Concepts, approaches, and technical challenges. *Power Systems, IEEE Transactions on In Power Systems, IEEE Transactions on, Vol*, 22(4):1743–1752, 2007.

[2] M. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.

[3] J. Oyarzabal, J.; Jimeno, J. Ruela, and C. Engler, A.and Hardt. Agent based micro grid management system. *International Conference on Future Power Systems*, 18(18), Nov 2005.

[4] M. Pielke, M. Troschel, M. Kurrat, and H.-J. Appelrath. Operation strategies to integrate chp micro units in domestic appliances into the public power supply. *Proceedings of the VDE-Kongress*, 2008.

[5] D. Pudjianto, C. Ramsay, and G. Strbac. Virtual power plant and system integration of distributed energy resources. *Renewable Power Generation IET(1)*, pages 10–16, 2007.

[4]Represented in the graphs as the individualistic approach.

A multi-agent system for managing plug-in hybrid vehicles as primary reserve capacity in a Smart Grid

Stijn Vandael[1], Tom Holvoet[1], Geert Deconinck[2]

[1]DistriNet, Department of Computer Science
Katholieke Universiteit Leuven, Belgium
{firstname.lastname}@cs.kuleuven.be

[2]ELECTA, Department of Electrical Engineering
Katholieke Universiteit Leuven, Belgium
{firstname.lastname}@esat.kuleuven.be

ABSTRACT

Driven by environmental concerns, there is an increase in uncontrollable electricity generation from renewable energy sources. To maintain the delicate balance between production and consumption on the electricity grid, the amount of expensive reserve capacities is increasing.

An intelligent electricity grid or Smart Grid will enable the control of appliances at consumers' homes to improve grid balance. In particular, the control of PHEVs (plug-in hybrid vehicles) will offer a large flexible reserve. The goal of our research is to compare several solutions for controlling PHEVs as a reserve capacity. This paper takes a first step by controlling PHEVs with a multi-agent system as a primary reserve capacity.

Simulations show that the multi-agent system is able to control PHEVs as a primary reserve, charge them in time and flatten the transformer load. To guarantee a constant regulation power throughout the day, PHEVs require a grid connection at work and the possibility of V2G (vehicle-to-grid).

1. INTRODUCTION

In an intelligent electricity grid or Smart Grid, it becomes possible to coordinate and control devices at a household level to improve grid balance. One particular interesting application is the coordination of plug-in hybrid electric vehicles (PHEVs), hybrid electric vehicles that can be charged through a simple grid connection. In the context of controlling a large fleet of PHEVs, two important challenges were identified :

- *Large scale and dynamic environment*: research predicts that in 2030, around 30% of all vehicles in Belgium will be plug-in hybrid electric vehicles (PHEVs) [2]. Moreover, PHEVs arrive and depart at varying times, which entails a very dynamic environment.
- *Multiple coordination objectives*: many parties and objectives can be identified in controlling a fleet of PHEVs. E.g PHEVs want to charge their battery, the grid operators want to limit the transformer load etc.

In this paper, the first results of a MAS for controlling PHEVs as a reserve capacity are discussed and evaluated. These results provide a basis for future research.

2. BACKGROUND: PRIMARY RESERVE CAPACITY

The European grid has a constant net frequency of 50 Hz. Deviations from this frequency are caused by imbalances between production and consumption. A frequency above 50 Hz indicates a surplus, while a frequency below 50 Hz indicates a shortage. Reserve capacities are connected to the grid to maintain a stable frequency.

There are three types of reserve capacities available to maintain frequency stability: primary, secondary and tertiary reserves. While primary reserves quickly respond to a frequency deviation with limited power, secondary and tertiary reserves respond slower with more power. The deployment of PHEVs as primary reserve has the largest potential [3].

In case of a frequency deviation, the primary reserve is activated first. The technical requirements of a system deploying a primary reserve are determined in the Operational Handbook of the ENTSO-E [4]. A summary of the most important requirements:

1. A frequency measurement device has to be available.
2. The primary reserve has to be equipped with an automatic control system that reacts to frequency deviations by activating the primary reserve.
3. The reaction time of the primary reserve is within a few seconds of a frequency deviation. 50% of the primary reserve has to be deployed after 15 seconds, while 100% of the reserve has to be deployed after 30 seconds and maintained for at least 15 minutes.
4. The primary reserve has to be available for 24/7.

While (1) and (2) are purely functional requirements, (3) and (4) are quality requirements. The preliminary evaluation in this paper will be based on quality requirement (4).

3. A MULTI-AGENT SYSTEM FOR MANAGING PHEVS AS PRIMARY RESERVE CAPACITY

Based on the identified challenges, a decentralized MAS approach is proposed. First, a schematic overview of the MAS is discussed. Hereafter, the coordination mechanism used by the agents to function as a primary reserve capacity is explained.

3.1 Multi-agent system overview

The multi-agent system (figure 1) is based on our previous work on demand-side management of PHEVs [5]. Each type of agent has its own objectives: a PHEV agent charges its

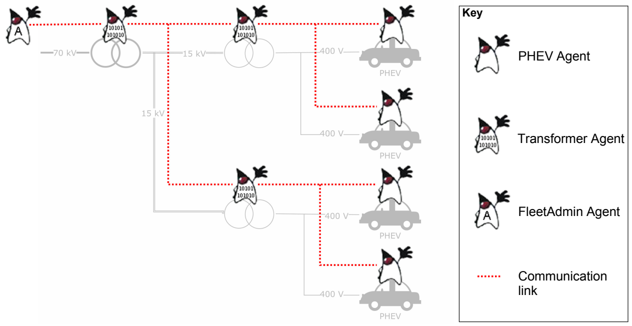

Figure 1: Schematic overview of the MAS.

battery in time and a transformer agent flattens the load of its transformer. A fleet manager agent is added to the original multi-agent system. The goal of this agent is to manage the fleet of PHEVs as a primary reserve capacity.

3.2 Coordination of the multi-agent system

The coordination mechanism for managing a fleet of PHEVs as a primary capacity reserve is an extension of the original coordination mechanism for demand-side management [5]. This DSM coordination mechanism ensures PHEVs charge in a way that flattens the load at each transformer.

The extended coordination mechanism adds two steps. In the first step, each PHEV calculates their individual up and down regulation power E_{phev+} and E_{phev-}. The up (down) regulation power indicates the possible decrease (increase) in charging power of a PHEV. Next, these values are summed up towards the fleet agent to E_{total+} and E_{total-}.

In the second step, the fleet agent calculates the required percentage p_+ and p_- of the total regulation power E_{total+} and E_{total-}, based on the contracted regulation volume (equation 1 and 2). This percentage is broadcasted to all PHEV agents to inform the agents about the percentage of their individual up- and downregulation power that has to be deployed. Consequently, PHEVs are now able to adapt their charging power to frequency deviations.

$$E_{contract+} = p_+ \cdot E_{total+} = p_+ \cdot (E_{phev_1+} + ... + E_{phev_n+}) \quad (1)$$

$$E_{contract-} = p_- \cdot E_{total-} = p_- \cdot (E_{phev_1-} + ... + E_{phev_n-}) \quad (2)$$

with $0 \le p_+ \le 1$ and $0 \le p_- \le 1$

4. EVALUATION

A fleet of PHEVs acting as a primary reserve was evaluated in a large scale simulation of a distribution grid of a typical city in Belgium (17,775 PHEVs).

In the first part of the evaluation, the regulation potential of the fleet was evaluated by observing E_{total+} and E_{total-} of the fleet during an entire day (figure 2). The PHEVs didn't adjust their charging power to the frequency in this scenario, because this would affect the regulation potential. Without V2G (providing electricity from the PHEVs to the grid), down regulation power is very low. During the day, both down and up regulation power are low, because PHEVs are not connected to the grid.

In the second part of the evaluation, the PHEVs use p_+ and p_- for adjusting their charge power to the frequency. As expected from the results of the potential analysis, V2G is essential for down regulation and PHEVs cannot comply with the contracted volume during the day. Consequently, a day-time connection to the grid (e.g at work) is essential.

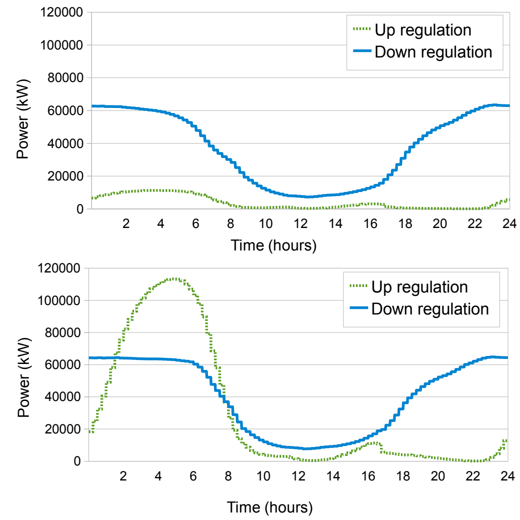

Figure 2: Up and down regulation power of the fleet without V2G (a) and with V2G (b).

5. CONCLUSION AND FUTURE WORK

In the future, the need for reserve capacities will rise due to the increase in electricity generation from renewable energy resources. To identify the best solution for deploying PHEVs as a primary reserve, different techniques have to be compared. In this paper, we proposed a multi-agent solution and preliminary results which provide a basis for future research. This research will consists of, but will not be limited to, the following aspects:

ECONOMIC INTEGRATION. To realistically coordinate PHEVs as reserve capacity, economic objectives have to be integrated. A fleet manager will optimize its profit and mitigate its financial risks, while a PHEV owner will charge at the lowest price.

PREDICTIONS. While the frequency is stochastic, departure and arrival times of PHEVs can be predicted. From this information, a fleet manager can make more informed decisions to control PHEVs.

This research is conducted in coorporation with Linear, the first project in Flanders that researches active demand management with real end-consumers on a large scale.

6. REFERENCES

[1] Martin LaMonica. Cisco: Smart grid will eclipse size of internet. CNET, May 2009.

[2] K. J. Roth, E. L. Smith, and C. W. Anderson. Emissions of road traffic in Belgium. FEBIAC, 2006.

[3] Algemene Raad van de CREG. studie over de mogelijke impact van de elektrische auto op het belgische elektriciteitssysteem. Technical report, CREG, 2010.

[4] European Network of Transmission System Operators for Electricity (ENTSOE). Ucte operation handbook, 2009.

[5] Stijn Vandael, Nelis Boucké, Tom Holvoet and Geert Deconinck. Decentralized demand side management of plug-in hybrid vehicles in a Smart Grid, ATES 2010.