Distributed Thermal Storage Using Multi-Agent Systems[1]

C. Johansson[1], F. Wernstedt[2] and P. Davidsson[3]

[1]Blekinge Institute of Technology, PO Box 214, SE-374 24 Karlshamn, Sweden
christian.johansson@bth.se
[2]NODA Intelligent Systems AB, Drottninggatan 5, SE-374 35 Karlshamn, Sweden
fredrik.wernstedt@noda.se
[3]Malmö University, SE-205 06 Malmö, Sweden
paul.davidsson@mah.se

Abstract. Thermal storage is an essential concept within many energy systems. Such storage is generally used in order to smooth out the time lag between the acquisition and the use of energy, for example by using heat water tanks within heating systems. In this work we use a multi-agent system in order to maintain and operate distributed thermal storage among a large group of buildings in a district heating system. There are several financial and environmental benefits of using such a system, such as avoiding peak load production, optimizing combined heat and power strategies and achieving general energy efficiency within the network. Normally a district heating system is purely demand driven, resulting in poor operational characteristics on a system wide scale. However, by using the thermal inertia of buildings it is possible to manage and coordinate the heat load among a large group of buildings in order to implement supply driven operational strategies. This results in increased possibilities to optimize the production mix from financial and environmental aspects. We present a multi-agent system that combines the thermal storage capacities of buildings with production optimization strategies. The system consist of producer agents responsible for valuing the heat load management, consumer agents managing the quality of service in individual buildings while consenting to participate in heat load management, and a market agent acting as a mediating layer between the producer and consumer agents. The market agent uses an auction-like process in order to coordinate the heat load management among the consumer agents, while the producer agents use load forecasting in order to evaluate the need for heat load management at any given point in time. A consumer agent uses continues feedback regarding indoor climate in order to uphold quality of service while participating in heat load management. Real-time data from a district heating system in Sweden is used in order to evaluate the agent system in relation to operational peak load management. The results show clear financial and environmental gains for the producer as well as participating consumers.

Keywords: coordination, distributed decision making, real-time agreements

---

[1] AT2012, 15-16 October 2012, Dubrovnik, Croatia.

1 Introduction

This paper describes a multi-agent system used in order to improve the operational management of a district heating system in relation to financial and environmental aspects. District heating and cooling is an integrated part of the energy infrastructure in many countries. This is especially true in Northern and Eastern Europe, although district heating and cooling is growing rapidly in popularity around the world [1]. A district heating system consists of one or more production units which distribute heated water or steam throughout at pipe network. Buildings are then connected to this network, usually through the means of heat exchangers which transfer the heat to the heating- and tap water system in the building. District heating systems are often considered environmentally and financially sound, since the centralized set-up facilitates large scale energy efficiency schemes. More or less any type of production units can be used in a district heating system, and the fuels used range from biomass and industrial waste heat to coal and oil. Many district heating systems also make use of combined heat and power production (CHP) which have a very high level of utilization of the primary fuel used [2]. A CHP plant produces electrical power at the same time as it heats the water for the district heating system. First water is boiled into steam which is used in order to run a turbine, which in turn is connected to a generator which converts the mechanical energy to electrical energy. Afterwards the steam is passed through a condensing unit which transfers the heat to the district heating system. As the steam cools off it can be re-heated before being passed into the turbine again. A traditional power plant has an energy efficiency of about 30-50%, while a CHP plant utilizes about 80-90% of the primary fuel. The electricity produced in the CHP plant is then sold on the power market while the heat is being sold within the connected district heating system.

Normally a district heating system is purely demand driven, in the sense that the production units can only react to the current heat demand within the network. The bulk of the heat demand consist of demand for building heating, although substantial heat load peaks are also generated by social patterns in relation to tap water usage. The heating demand in a building is correlated to the outdoor temperature, while tap water usage is more correlated to time of day. This combination of outdoor temperature dependencies and social behaviour over the day causes fluctuations in the heat demand. Such fluctuations or heat load peaks are very undesirable for a number of reasons. Most district heating systems have a production set-up using some sort of base load boilers which are complemented with a range of peak load boilers. The base load boilers are often fuelled by cheap and environmentally friendly fuels such as biomass or industrial waste heat, while the peak load boilers normally use fossil fuels such as oil or gas which are not only expensive but also generate undesirable emissions.

In this paper we describe a multi-agent system which uses the thermal inertia of buildings in order to manage the heat load demand in relation to the operational status at the production units. By managing the demand in such a way it is possible to make the district heating system more supply driven, i.e. it is possible to optimize the heat load usage in relation to the production status and not only the other way around. This is similar to using large storage tanks, although here the thermal inertia of the connected buildings is used as a distributed storage instead. By using such a set-up it is possible to shed heat load peaks in order to avoid using peak load boilers, or to move the heat load peaks in time in order to coordinate the heat load peaks with high spot prices for electricity when using CHP. The use of large storage tanks is well understood and is a mature technique within district heating systems, and similar strategies are possible to utilize with the proposed distributed thermal storage [3].

The multi-agent system consists of three different agent entities; the producer agent, the consumer agent and the market agent. This paper focuses on the market agent component which acts as a mediating layer between the consumer and producer agents. The market agent uses as an auction-like process in order to coordinate heat load demand among the consumer agents in relation to the operational status of the producer agents.

The paper starts off by describing related and previous work concerning multi-agent based heat load management within district heating systems. In section three the physical process of managing the heat load usage in relation to thermal inertia of a building structure is discussed. Section four presents the functionality of the coordination process within the multi-agent system. Section five describes the experimental set-up used in evaluating the system, and the results are presented and discussed in sections six and seven. Finally the paper is concluded in section eight and nine with conclusions and future work.

2 Related work

The framework for the multi-agent systems is built around three different agent types; the consumer agents, the producer agents and the market agent used to coordinate the interaction between the two former agent groups. This basic set-up was first introduced in a previous paper [4], and has since been expanded to cover the practical issues present in a real-time industrial setting [5]. The last few years the system has matured through a series of industrial installations showing the potential of the concept in operational settings [6]. Lately the theoretical foundation for the system has been refined through work considering the producer and consumer agent entities. An optimization algorithm for coordinating the desired heat load in relation to operational constraints was presented in a recent paper [7]. Together with heat load forecasting algorithms such optimization algorithms are used as decision support tools by the producer agents when calculating their desired future state. Heat load forecasts are produced by relating historical operational data of the district heating system with forecasts of the outdoor temperature, in combination with estimations of social behaviour affecting tap water usage and ventilation. Such systems have been studied extensively in previous papers [8] and [9]. More complex solutions use seasonal autoregressive integrated moving averages [10], or basic Box-Jenkins autoregressive integrated moving averages [11] and [12]. Other approaches for heat load forecasting include a grey-box process explicitly using climate data to forecast heat consumption in a large geographical area [13].

The consumer agent tries to balance the ability to partake in load control while simultaneously upholding required levels of quality of service (QoS). The concept of quality filters was proposed in [14]. This concept is based on the idea that the consumer agent should be responsible for the QoS in the building by means of continues measuring and evaluation, thus being able to take informed decisions regarding participation in any requested load control. Maintaining a sufficient level of QoS is at the core of any heating or cooling strategy, and a consumer agent must only act within the constraints set by hereby. An energy balance model for large-scale QoS evaluation within a multi-agent system was presented in [15]. The focus of [15] and [7] was on the specific application of combined heat and power generation while this paper combines the conceptual foundation for producer and consumer agents with a framework for a general application of distributed thermal storage by extending the market agent process.

The basic process of using heat storage facilities within buildings in order to improve the operational functionality of a district heating system has been previously studied in [16] and [17]. The ability to do this on a large scale throughout an entire district heating network was evaluated and quantified in [18], and further studies on the process of storing heat in building structures are presented in [19]. The process of pre-cooling and demand limiting in relation to this has been studied in several studies such as [20] and [21]. Basic heating control strategies are presented in [22], while more complex adaptive strategies are presented in [23] and [24]. All these strategies make use of the thermal inertia within the building structures in order to improve the heating and/or cooling demand.

An early attempt on demand side management was presented during the 1980s. And although the current hardware and communication capabilities limited the operational functionality of the proposed system, the future potential of the system was apparent [25]. It should be noted that the importance of two-way communication in order to enable feedback was acknowledged even in these early works [26].

3 Multi-agent system overview

The multi-agent system uses three different main types of agents in order to distribute and coordinate the heat load among the participating buildings; producer agents, consumer agents and market agents. A producer agent will work to optimize its operational production in regards to fuel prices in the general case and power market spot prices in the case of combined heat and power generation. Such operational management is done by distributed heat storage among buildings connected to the district heating network. Whenever the producer agent deems it necessary to perform heat load management it will request load control among the consumer agents. The producer agent will have to pay the participating consumer agents so it is important for the producer agent to set the price constraints correctly. The coordination of this load control between the consumer and production agents is done by the market agent.

The agents are controlled by an implicit norm system in which agents are assumed not to cheat or lie. Specifically, in order to avoid collusive behaviour the agents are not permitted to form coalitions. A consumer agent is assumed to be truthful regarding its internal energy balance status, i.e. it is assumed that a building owner doesn't want the residents to complain about the indoor temperature. Furthermore, in multi-producer agent settings it is assumed that the producer agents will not impose on load management instigated by other producer agents, although they have the information to do this based on the on-line auction status.

An energy company and all its production units are represented by a producer agent. Most district heating systems are run by one single energy company, but there could also be several energy companies competing within a network. In the case of a single energy company the producer agent will only need to consider its own operational constraints, while in a multi-energy company setting it will also need to perform real-time evaluation in regards to the status of its competitors.

Each consumer sub-station participating in the system will be represented by a consumer agent. Normally each building has its own sub-station, but several buildings can also be connected to a single sub-station. A consumer sub-station consists of heating control systems, heat exchangers, valves and pumps, and this is the point of control for the consumer agent. Normally such a sub-station will be connected to an outdoor temperature sensor which acts as the input signal for the control system, i.e. the colder it gets, the more the heat demand will increase. In order to control the heat load demand of the building a Linux-based computer I/O platform is used. This platform is connected between the outdoor temperature sensor and the existing control system. By manipulating the outdoor temperature signal the consumer agent is thus able to influence the control behaviour of the sub-station. Figure 1 shows an example of how this works in practice.

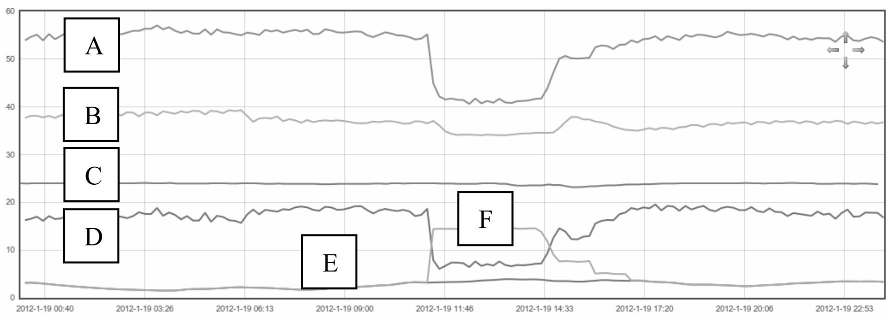

Fig. 1. Load control by consumer agent

The data in Figure 1 is collected from an online consumer agent operating on the I/O platform. The data shown covers twenty-four hours. The forward temperature into the heating system is shown by the line A, while the return temperature from the heating system is shown at line B. Together with the mass-flow (not shown in the picture), the amount of energy used can be calculated. Line C shows the average indoor temperature of the building in question. Line D shows the difference between the forward and return temperatures in the heating system, i.e. line A minus line B. Line E shows the actual outdoor temperature, while line F shows the manipulated outdoor temperature. In the example it is shown that E and F are identical until the consumer agent performs a short load control, when E and F diverge in the middle of the figure. When this happens the forward temperature into the heating system (A) is lowered, which leads to a decrease in heat load instantaneously. Conversely the opposite can be done if the consumer agent wants to buffer the thermal inertia of the building by inserting extra heat.

Unlike the producer and consumer agents there is no physical object corresponding to the market agent. However, there is a good reason to separate this functionality since it will make for a more transparent coordination process. Since actual money is involved in the process it is important that the market process is deemed trustworthy by all participating entities.

3.1 Heat block value

In a previous paper we presented an optimization algorithm for synchronizing heat load usage in relation to power market spot prices. It was shown that such optimization resulted in substantial financial benefits when using combined heat and power generation on a day-ahead and intraday spot price market. In this paper we extend this algorithm to cover the general case, i.e. adapting it for distributed heat storage in general. The generalized optimization model is shown in equations 1 to 7.

$$\max \sum_{i=0}^{m} \sum_{j=0}^{n} f(x_{i,j}) * (price_j - fuel_{i,j} - op_{i,j})$$

(eq. 2)

$$subject\ to$$

$$f(x) = \begin{cases} 1\ if\ x_{i,j} \neq 0 \\ 0\ if\ x_{i,j} = 0 \end{cases}$$

(eq. 3)

$$if\ f\left(x_{i+1,j}\right) = 0\ then\ f\left(x_{i,j}\right) = 0$$

(eq. 4)

$$b_{upper} \geq \sum_{i=0}^{m} f(x_{i,j})$$

(eq. 5)

$$b_{lower} \leq \sum_{i=0}^{m} f(x_{i,j})$$

(eq. 6)

$$\sum_{i=0}^{m} f(x_{i,j}) \leq \sum_{i=0}^{m} f(x_{i,j-1}) + b_{dynamic}$$

(eq. 7)

$$\sum_{i=0}^{m} f(x_{i,j}) \geq \sum_{i=0}^{m} f(x_{i,j-1}) - b_{dynamic}$$

(eq. 8)

Equation 2 maximises the total earnings. The value of sold heat can only change from hour to hour, while the cost for producing the heat can change both in time and in relation to the level of heat load demand. The costs are separated in fuel costs and operational costs. Such operational costs include costs for starting and shutting down boiler and costs related to managing the boilers. An m by n matrix represents the heat load demand during a time period, where m is the theoretical maximum amount of heat load in the system. The value of n is normally 24, i.e. the head load demand is evaluated one day ahead and is discretizised into hours. Equations 3-8 define the operational boundary conditions for the optimization model.

An optimization model for production is dependent on heat load forecasts ranging from hours to a few days ahead. Such forecasts are calculated based on a combination of previous operational data and weather forecasts. The main parameter from the weather forecast to consider is the outdoor temperature, although general weather conditions such as precipitation and wind affect the situation to some extent. Using the heat load forecasts a producer agent can arrange the wanted heat load demand hour by hour in relation to fuel prices and other related costs. By doing this it is possible to minimize the use of expensive and environmentally unsound peak load fuels. In order to make the model useful in practice the heat load demand represented by a matrix which is discretizised into blocks. Each such block is called a heat block, and it is these heat blocks that the producer agent will sell to the consumer agents through the market process.

A heat block represents a certain monetary value for the producer agent, since if the consumer agents successfully implement the load control specified by the heat block this will translate into an improved operational status at the physical production units. However, some of these added earnings will have to be used to pay the consumer agents participating in managing the heat load demand within each heat block. Based on the optimization model the producer agent will be able to evaluate the financial value for heat block, and will thus be able to set a maximum price it is willing to pay for the implementation of the associated load control. The producer agent is also able to set the size for the heat blocks based on the boundary conditions in the optimization model, i.e. the size of the matrix is based on the operational conditions within the production units.

3.2 Consumer agent assets

The consumer agents get paid by the producer agents in order to perform load control. The payments are accumulated during the payment period, usually a month, and then subtracted from the heating bill. Obviously a consumer agent would want to maximize this revenue, by performing as much load control as possible. At the same time the consumer agent is responsible for maintaining an acceptable level of quality of service, i.e. a building owner wants to save money on the heating, but only to the point where the residents of the building are not complaining about it. Therefore the consumer agent assets are modelled based on the indoor temperature, and thus the agent's ability to accumulate heat load control through the auction process is inversely proportional to any deviation from the wanted indoor temperature. The consumer agents also has an upper and lower limit of maximum allowed deviation from the wanted indoor temperature, outside which the agent is not allowed to perform any load control at all. Furthermore the consumer agent also has time constraints in regards to how long the temperature is allowed to deviate from the wanted temperature, i.e. the agent cannot just set the temperature at the upper or lower limit and let it stay there indefinitely. All these operational constraints are set by the building owner depending on characteristics of the building in question.

In order to fulfil the task of maintaining quality of service the consumer agent must have some way of evaluating the indoor temperature. The I/O platform used in the installed system is equipped with wireless communication for such sensors which make it easy to collect indoor sensory data. In addition to this the consumer agent continuously calculates an internal energy balance model of the building based on sensory data from the consumer sub-station in relation to the energy signature and time constant of the building in question. The energy signature is a characteristic value for each individual building indicating the amount of heat load needed to uphold one degree of difference between the outdoor temperature and the indoor temperature. The time constant is another value specific for each building indicating how long time it takes for the indoor temperature to reach equilibrium with the outdoor temperature if no external heating is supplied. Using these values is convenient since they can normally be estimated within acceptable levels without any costly energy analysis of each building. It is also convenient in relation to the auction process since a change in the indoor temperature in relation to the specific energy signature of a building translates to a specific amount of heat load, while the time constant translates into an estimate of how long a specific heat load change can be upheld without jeopardizing the quality constraints. This makes it possible for the consumer agent to relate the indoor temperature status directly to an amount of available currency, expressed in the possible heat load change, to be used in the auction process.

3.3 Auction process

When a producer agent estimates the need for load control it will issue one or more heat blocks to the market agent. Each such heat block is the size of the heat load times one hour, i.e. the length of time for a heat block is static while the size in heat load may vary. The market agent will then divide each heat block into heat slots which it auctions off among the consumer agents. This basic set-up is visualized in Fig 2.

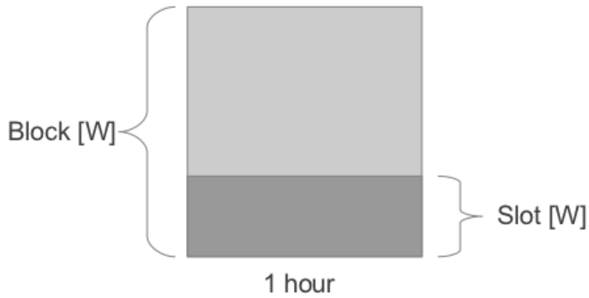

Fig. 2. Heat block structure

The reason to divide the heat block into smaller slots is that normally no individual building is large enough to handle a whole block by itself. The market agent performs the sizing of these slots based on knowledge regarding the outcome of previous auctions, i.e. the market agent will have a rough idea about how large chunks the consumer agents can handle at a time. If the market agent does not have any such prior knowledge it will start by trying to auction off the whole block as one slot, then dividing it in half until consumer agents start to respond to the auction calls.

Each slot is sold through a reverse Dutch auction process. A Dutch auction starts at a high price which is then decreased in increments by the auctioneer until a buyer accepts the price. The purpose of using such auctions is mainly that they come to a conclusion fast, thus minimizing communication overhead which is imperative in an industrial setting such as this. A reverse Dutch auction is like a normal Dutch auction except that it starts from zero and works itself up until either a bid is accepted by some consumer agent or the maximum price set by the producer agent is reached. Once a bid is accepted the market agent will issue a contract to the winning consumer agent. The contract specifies the size of the heat slot and also the time when to start implementing the load control. A consumer agent can participate in several successive auctions, as long as it is able to maintain the quality of service.

A producer agent can issue a heat block many hours, or indeed days, before it is to be implemented. However, if the block is issued to early the consumer agents might be reluctant to commit themselves since they cannot predict their indoor temperature state too far into the future. On the other hand, if the heat block is issued to late a large portion of the consumer agents might already be under contract, potentially preventing them from participating further.

If the market agent isn't able to allocate the whole heat block, then the producer agent can choose to either accept the partially filled block or to reject it. In this case the consumer agents involved are bound by contract until the producer agent either accepts or rejects the block. On the other hand, if the market agent does succeed in allocating the entire heat block, the producer agent is obliged to accept it.

The market agent will keep track of the outcome of all auction activity. This data is publicly available to all entities participating in the agent system, which serves two purposes. First of all it ensures a transparency in the process, which is important since both consumer and producer agents want to be able to ensure the functionality of the debiting system. Secondly if such data is freely available all agents will share the same knowledge, thus decreasing the risk for unfairness in the market process.

4 Experimental setup

The experiments in this project were performed using data from a district heating system at Gothenburg Landvetter Airport in Sweden. This system has around thirty buildings connected, and the size of these buildings range from large arrival halls and air plane hangars to smaller buildings of various kinds. The data studied spans the month of January 2011. This district heating system is installed with the previously mentioned I/O platform, although the overall agent system was not yet operational at the time of this study. Hence the agent coordination evaluation was done using the DHEMOS simulation environment, although operational data from the district heating system was used. DHEMOS is a simulation framework for district heating systems, combining simulation models for production, consumption and distribution [27]. The basic function of DHEMOS is described in [28]. The optimization calculations for the producer agents are done using Octave 3.2.4. The DHEMOS simulations have been run on a Linux/Ubuntu 11.10 computer with Intel Core i5 CPU with 4GB of RAM.

The scenario studied involves the management of peak load. The production units use a base load burner using biomass fuel and an oil based peak load burner. The oil burner is necessary to use at heat load demand levels above 4.5 MW. The goal is to minimize the use of the oil burner since it is both costly and environmentally unfriendly. The prices involved are based on averages during the month of January 2011. The fuel costs are valued at €32.29/MWh for biomass and €75,84/MWh for oil, while the income is €54,17/MWh. Biomass emit zero CO_2 while oil results in emissions at 271 kg/MWh.

5 Results

Fig 3 shows a comparison between DHEMOS and the measured operational data at the production site. The DHEMOS data has distributed heat storage (DHS) inactivated in order to compare to the operational data. During the first half of the month the operational data displays somewhat of a volatile behaviour in comparison with the simulated data. These violent fluctuations are caused my measurement errors. However, the general correlating trend between the measured data and the simulation is apparent. If the simulation data is used as a measurement of actual heat load demand it can been seen that short periods of heat load shortage has most likely resulted during the first and last peaks, at around 50 and 630 hours. The fact that no-one has complained during these times might be considered an indication in regards to the potential of thermal storage within buildings.

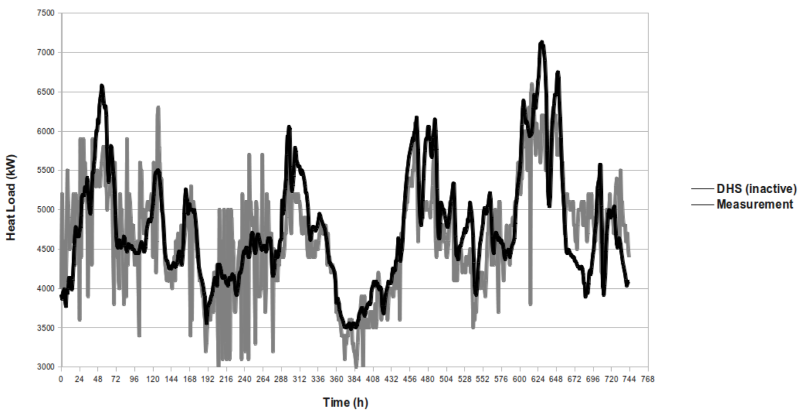

Fig. 3. Comparison between DHEMOS and measured data

Fig 4 shows a simulation run with DHS activated in order to avoid peak load usage above the limit of 4.5 MW. The producer agents wants the heat load demand to be as close as possible to this limit, since they want to sell as much as possible of the energy produced by biomass.

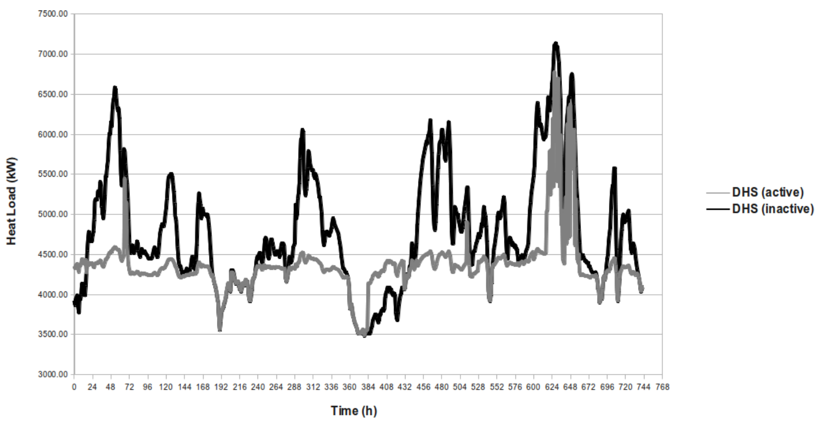

Fig. 4. Comparison between active and inactive distributed heat storage

It is obvious that the agent system isn't able to fully avoid all peak load usage. During the latest peak load, at around 600 hours, the consumer agents are not capable of responding to the heat blocks being offered by the market agent. However during the major part of the month the system is able to successfully manage the heat load demand. At around 400 hours it can be seen how the consumer agents buffer heat by increasing their heat load demand in order to prepare for the coming peak at around 450 hours.

Fig 5 shows an example of the indoor temperature variations in a building during the simulation period. The consumer agent used the energy balance model in order to calculate this value.

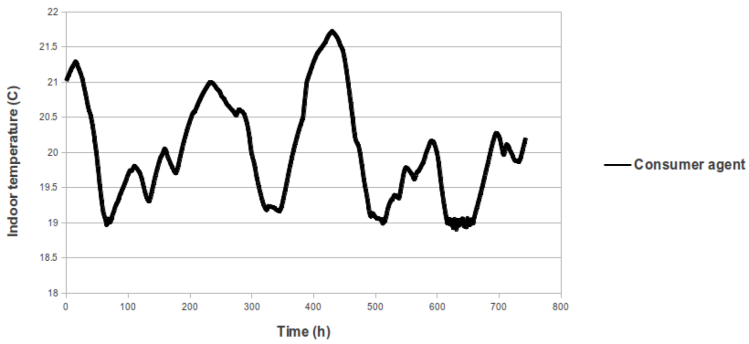

Fig. 5. Indoor temperature variations in a single building

The lowest temperature value accepted by the consumer agent in this case is nineteen degrees, while the wanted temperature is twenty-one degrees. The consumer agent approaches the lower limit on several occasions, most notably during the latest peak at around 650 hours. When this happens the consumer agent will start to lose auctions, thus rendering it unable to perform further load management. It should be noted the calculated indoor value shown above should be viewed as a control variable for the consumer agent, and not necessarily a representation of the actual indoor temperature. This is due to the fact that the energy balance model doesn't take into account heat contributing factors such as solar radiation, excess heat from electronic equipment or social activities.

As the producer agent manages the heat load demand the total amount of energy sold will be reduced, but since the majority of this energy is produced by expensive oil the net profit will increase. Table 1 shows a comparison between active and non-active distributed heat storage.

	Energy	**Gross**	**Net**	**Agent cost**	**Profit**	**CO_2**
DHS off	3548,24	192212	63905	0%	63905	85540
DHS on	3245,62	175819	69157	30%	65480	11660

Table 1. Energy in MWh, gross and net income and profit in euro, CO_2 in kg

Table 1 show that the profit will increase by €3675 during the period and that the CO_2 emissions will decrease by more than 86%. The agent cost represents the amount of money the producer agent uses to pay the consumer agents. In our simulations this value was between 10-30% although this is dependent on the ratio of competing consumer agents in relation to the amount of heat blocks available on the market, i.e. a direct relation between supply and demand. The cost of the installed system is about €20000 which in relation to a profit of about €3500 per month will lead to the system being amortized within one heating season.

6 Discussion

The presented system uses an implicit norm system, which assumes that the agents are benevolent and acting towards globally known goals using universally accepted strategies. However, in a fully operational system these agents would be acting on a financial market including energy companies and building owners who want to maximize their own financial gains. In the current system norms are present in such a way that agents are expected to behave in a certain fashion, e.g. not to cheat or lie about their status. In a practical setting it is possible to maintain this system since the implementation of all agent parameters where set within this project. However, in a real-life system the participating building owners or energy companies would be able to set any parameters they like, and could even re-program the agents as long as they adhere to the communication protocol.

The strategy of the consumer agent is to bid as a combined function of its valuation of the load management and its estimation concerning other agents' valuation. This obviously opens up for coalition forming among the consumer agents, although the current implementation doesn't allow for this. Consumer agents might want to form coalitions in order to achieve high prices by agreeing to withhold auction bidding and then splitting the proceeds. If a group of consumer agents were successful in rigging the market like this they could achieve more than the 10-30% profit share they got during the presented experiments.

In regards to the producer agents there is no incentive for them to lower the maximum price below its true valuation, since this would only harm the producer agent itself by depriving it of financially profitable load management.

7 Conclusions

In this paper we have presented a multi-agent system for distributed heat storage. It was shown that such a set-up can be used to reduce, and in many instances remove, the need for financially and environmentally unsound peak load fuel usage. Since the multi-agent system is operationally adaptable it is able to adjust the heat load demand in relation to actual operational constraints among the production units. This makes it possible to increase the net profit for the energy company even though the total amount of energy being sold is less.

It was shown that the consumer agents will cooperate in load management as long as they are able to simultaneously uphold their desired quality of service.

The net profit in district heating system in question was increased by about 2.5% which translates to a return of investment in less than one heating season considering the installation costs involved. At the same time the CO_2 emissions were reduced by more than 86% due to the shift in fuel composition.

8 Future work

In the future a framework for explicit management of norms will have to be added. This is imperative since the physical entities represented by the agents need to maintain trust in the system, even when those entities are allowed to set their own agent parameters. This system should be based on a set of explicit norms or ground-rules, which are then managed through a layer of individual trust among agents coupled with a globally visible reputation system. The market agent will act as the main manager for this system, and will use punishment by exclusion in order to maintain order. Exclusion translates into financial loss, which will act as deterrence for the agents to deviate from the norms.

9 Acknowlegements

The authors would like Swedavia for letting us use operational data from the district heating system at Gothenburg Landvetter Airport. Also, we would like to thank NODA Intelligent System AB for the use of the agent I/O-platform.

10 References

1. Constinescu, N. District Heating and Cooling: Country by Country, 2007 Survey. Euroheat and Power (2007)
2. Horlock, J.H.. Combined Heat and Power. 2th ed. Pergamon Books In; 2008
3. Andrepont, J.S. Thermal Energy Storage: Optimizing district heating systems in more way than one. District Energy, First Quarter 2012. International District Heating Association (2012)
4. Wernstedt, F., Davidsson, P. and Johansson, C. Demand Side Management in District Heating Systems. In Proceedings of Sixth International Conference on Autonomous Agents and Multiagent Systems (AAMAS), Honolulu, Hawaii (2007)
5. Johansson, C., Wernstedt, F. and Davidsson, P. Deployment of Agent Based Load Control in District Heating Systems. First International Workshop on Agent Technologies for Energy Systems. Toronto, Canada (2010)
6. Wernstedt, F. and Johansson, C. Demonstrationsprojekt inom effekt- och laststyrning. Report 2009:26, The Swedish District Heating Association (2009) (In Swedish)
7. Johansson, C., Wernstedt, F. and Davidsson, P. Combined Heat and Power Generation using Smart Heat Grid. In Proceedings of the 4th International Conference on Applied Energy. Suzhou, China (2012)
8. Jonsson, G. A model for predicting the yearly load in district heating systems. In Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy 216, pp 277-281 (2002)
9. Dotzauer, E. Simple model for prediction of loads in district heating systems. Applied Energy 73, pp 277-284 (2002)

10. Grosswindhager, S., Voigt, A. and Kozek, M. Online Short-Term Forecast of System Heat Load in District Heating Networks. In Proceedings of the 31st International Symposium on Forecasting. Prag, Czech Republic (2011)
11. Box, G.E.P and Jenkins, G.M. Time Series Analysis, Prentice Hall (1991)
12. Chramcov, B., Dostal, P. and Balate, J. Forecast Model of Heat Demand. In Proceedings of the 28th International Symposium on Forecasting, Hong Kong, China (2009)
13. Nielsen, H.A. and Madsen, H. Modeling the heat consumption in district heating systems using a grey-box approach. Energy and Buildings, vol. 38, pp 63-71 (2006)
14. Wernstedt, F. and Johansson, C. Intelligent Distributed Load Control. The 11th International Symposium on District Heating and Cooling. Reykjavik, Iceland (2008)
15. Johansson, C., Wernstedt, F. and Davidsson, P. Smart Heat Grid on an Intraday Power Market. Third international Workshop on Agent Technologies for Energy Systems. Valencia, Spain (2012)
16. Wigbels, M., Böhm, B. and Sipilae, K. Dynamic Heat Storage Optimization and Demand Side Management. ANNEX VII I 2005:8DCH-05.06. International Energy Agency (2005)
17. Rolfsman, B. Combined Heat and Power Plants and District Heating in a De-Regulated Electricity Market. Applied Energy 78, pp 37-52 (2003)
18. Werner, S. and Olsson Ingvarsson, L. Building mass used as short term heat storage. In Proceedings of the 11th International Symposium on District Heating and Cooling. Reykjavik, Iceland (2008)
19. Hietmaäki, T., Kuoppala, J.M., Kalema, T. and Taivantii, T. Thermal Mass of Buildings – Central researches and their results. Tampere University of Technology, Institute of Energy and Process Engineering. Report 174. Tampere, Finland (2003)
20. Xu, P., Haves, P., Piette, M. and Zagreus, L. Demand Shifting With Thermal Mass in large Commercial Buildings: Field Tests, Simulations and Audits. Pier Final Project Report, California Energy Commission (2006)
21. Braun, J.E. Load Control Using Building Thermal Mass. Journal of Solar Energy Engineering, vol. 125, pp 292-301 (2003)
22. Björsell, M. Control Strategies for Heating Systems. Contract 8323P-95-02687. The National Board for Industrial and Technical Development
23. Rogers, A., Maleki, S., Ghosh, S. and Jennings, N.R. Adaptive Home Heating Control Through Gaussian Process Prediction and Mathematical Programming. Second International Workshop on Agent Technologies for Energy Systems, Taipei, Taiwan (2011)
24. Chahwane, L., Stephan, L., Wurtz, E. and Zuber, B. On a Novel Approach to Control Natural and Mechanical Night Ventilation. In Proceedings of Building Simulations. 12th Conference of International Building Performance Simulation Association, Sydney, Australia (2011)
25. Österlind, B. Effektbegränsning av fjärrvärme. Report 63:1982, Byggforskningsrådet, ISBN 91-540-3714 (1982) (in Swedish)
26. Österlind, B. Avläsnings- och kontrollsystem för fjärrvärmenät. Report 85:1990, Byggforskningsrådet, SIBN 91-540-5254-8 (1990) (in Swedish)
27. DHEMOS – District Heating Modelling and Simulation Tool. Retrieved June 28, 2012 from http://www.dhemos.org
28. Johansson, C. and Wernstedt, F. Dynamic Simulation of District Heating Systems. Third European Simulation and Modeling Conference. Oporto, Portugal (2005)