Deployment of Agent Based Load Control in District Heating Systems

Christian Johansson
Blekinge Institute of Technology
PO Box 520, SE-372 25, Ronneby, Sweden
chj@bth.se

Fredrik Wernstedt
NODA Intelligent Systems AB
Drottninggatan 5
SE-374 35, Karlshamn, Sweden
fw@noda.se

Paul Davidsson
Blekinge Institute of Technology
PO Box 520, SE-372 25, Ronneby, Sweden
pdv@bth.se

ABSTRACT

This paper describes results and experiences from an industrial proof-of-concept installation of a multi-agent based load control system in three major district heating systems in Sweden. A district heating system is a demand-driven system, i.e. the consumption controls the level of energy input which the district heating producer needs to deliver into the system. The basic idea of load control is that the individual consumers can be utilized as heat load buffers which, when coordinated on a system-wide scale, can be used to adjust the total consumption demand instead of having to change the production scheme. Load control leads to several important benefits such as giving the district heating producer the capability to avoid using expensive and environmentally unsound peak load boilers, while at the same time lowering the overall energy consumption at the consumer side. In order for load control to work the system needs to be able to coordinate the behaviour of a large amount of consumer substations in relation to the dynamic status among a range of production units, while continuously maintaining a sufficient level of quality of service among the consumers. The results show that the multi-agent based system was capable of reducing the peak loads with up to 20% of the total load, and to lower the average energy consumption with about 7,5% without any deterioration of the experienced indoor climate. Different theoretical aspects of load control have long been studied, but it is not until recently that technical advances in hardware and communication infrastructure has made it possible to implement these schemes in real-world settings.

Categories and Subject Descriptors

J.7 [**Computer Applications**]: Computers in other systems—*industrial control, process control*

General Terms

Management, Performance

Keywords

Agent-based industrial applications

1. INTRODUCTION

A district heating system consists of one or more production plants, a distribution pipe network and a collection of consumers. The medium used to transport heat is normally water, although steam is also sometimes used. At a production plant boilers are used to heat the water which is then pumped through the distribution network. Each consumer node consists of a heat exchanger system which transfers the heat from the primary distribution network into the secondary radiator network within the building. The heat exchanger system is also used to heat tap water, although the tap water system is separated from the radiator system. District heating networks are consumer driven systems, i.e. it is the consumption which controls the amount of energy that needs to be produced.

The basic idea with load control is that the consumer side can constitute an heat load buffer which can strategically be used in order to perform consumption reductions instead of having to produce more energy[6]. Substantial environmental and financial benefits can be found by smoothing out variations in the heat load in such a way. This is due to the fact that a production unit consists of a range of boilers using differently priced fuels.

The results discussed in this article are based on a multi-agent based system which has been deployed in parts of the district heating network in three different cities in Sweden; Stockholm, Västerås and Linköping. All Swedish cities with more than 10 000 inhabitants have district heating networks, and about fifty percent of all heating in Sweden is based on district heating. The Swedish district market is worth about €2.5 billion ($3.5 billion) annually, with the combined total European and Russian market being worth about €100 billion ($140 billion)[3].

1.1 Load Control

As consumption rises the producer has to engage increasingly expensive fuels. Such peak load boilers are usually fuelled by expensive and environmentally unsound fossil fuels. The cost for producing heat using such peak load boilers is usually not covered by the price paid by the costumer. Using load control it is possible to cut such peaks in the heat load by using the consumer buildings as the equivalence of a large storage tank. If the need for peak load production is lowered then district heating companies will be able to forestall large investments in peak load capabilities. The ability to perform load control also means that new customers can be added to the district heating system without having to invest in more production capabilities. That there are periods where district heating companies would rather lower consumption need that to sell that power is clearly shown by the fact that they are increasingly using pricing rates based on the level of momentary effect usage and not only on the total amount of energy used. It is not uncommon for customers to install a local effect guard which cuts energy use above a certain threshold of momentary effect usage. This, however, is often not desirable from a system wide point of view, as there is no connection between the locally lowered energy usage and the actual system wide status. Basically it can be said that this is a distributed information problem, in the sense that the consumer systems have no system wide perspective and thus are not able to decide whether it is appropriate or not to perform load control. In order to do this the system needs information about the global status of the district heating network, i.e. the total heat load in connection to the present production state. In essence a system wide perspective is need in order to successfully perform load control[4].

To perform load control basically means that the system will cut the energy usage from time to time. In order to do this without jeopardizing the indoor climate some sort of intelligence is needed in the load control system. There must be some kind of feedback between the load control and the indoor climate. A certain level of energy must at all times be supplied to the radiator systems in order to avoid sudden temperature drops and to ensure that the indoor temperature is always kept within the acceptable range, even during longer periods of load control. It is clear that lowered energy usage will result in a lowered indoor temperature, but it is also equally clear that the process of heat loss from a building is very slow and that this makes it possible to utilize the building as a heat load buffer which can be used to move the heat load without affecting the perceived indoor climate.

Earlier experiments with a distributed multi-agent system in a smaller building area in the south of Sweden have shown good system results by having the heat exchanger stations in a number of buildings communicate and cooperate in regards to performing distributed load control without any perceived deterioration of the indoor climate[7].

The question faced in this project was that if this behaviour could be replicated over a larger number of buildings in a geographically spread area while constrained by commercially viable terms. Our previous work has shown a theoretical and small scale experimental feasibility of the load control system[9]. This specific project was about investigating the possibility to actually perform load control in an industrial setting using a multi-agent system.

2. SYSTEM DESCRIPTION

There are a range of financial and environmental benefits to be had from a system that facilitates real-time control of consumer energy usage. By having the system coordinate consumer behaviour in a large group of buildings it is possible to achieve substantial system wide benefits. This system is based on the idea that the district heating company is not interested in the behaviour of individual buildings but rather in the total merged heat load. In theory load control gives rise to long line of advantages:

- It is possible to better utilize base load boilers, instead of having to use peak load. Using base load is financially as well as environmentally desirable since peak load is normally more expensive and more likely than not to emit large amounts of carbon-dioxide.
- In connection to shorter peaks it is possible to perform load control in order to entirely avoid starting a peak load boiler. During start-up of a boiler the emissions are usually higher, since it takes a few hours for the boiler to reach its operational temperature.
- By using load control it is possible to not only cut the peaks, but also to move them in time. This provides benefits for combined heat and power generation, since it is possible to better match the demand on the power market.
- When adding new customers to the district heating network it is possible to forestall investments in new production capacity.
- In certain circumstances it is possible to use load control in order to bridge narrow segments of the district heating network without having to lay new piping.
- Load control alleviates the need for expensive storage tanks in the district heating network.
- By implementing load control techniques it is possible to prioritize between different customers during periods of shortage or extreme cold.
- Lowering the return temperature in the primary district heating network. This favours environmentally sound production, since these are normally less energy rich than fossil fuels.

2.1 Multi-agent System

One of the fundamental aspects of district heating is its reliability and high quality of service in regards to the end customer. An adequate ability to uphold this fact must be considered one of the more important requirements to any energy efficiency measure performed in a district heating network. In regards to load control this is a question of coordinating a connection between two functions in conflict, i.e. to uphold an acceptable indoor climate while achieving the needed system wide consumption profile. It is the ability to simultaneously fulfil these two requirement that differentiates so called intelligent load control from simpler forms of load control. Such simpler forms of load control might for example be local systems which implement load control based on predefined lists, or which uses tap-water prioritization. The basic problem with these type of solutions are not coordinated globally which, along with the fact that they do not incorporate any feedback from the indoor climate in the individual buildings, means that have no ability to achieve global production oriented goals while sustaining a desired indoor climate. In order to solve these issues we have designed the system based on a multi-agent approach, where each consumer and producer node are assigned to an agent. The goal of a consumer agent is to uphold the desired indoor climate while trying to participate in the overall system wide load control as much as its local heat load buffer allows. A producer agent is responsible for recognizing the need for load control, i.e. there is a need to manipulate the energy consumption among the buildings. The producer agent will then try distribute this load control among the participating consumer agents.

Obviously the indoor climate is connected with the local energy consumption, but there is a certain delay in this physical process. During shorter periods of time it is possible to manipulate the local consumption without any noticeable influence on the indoor climate. Normally, changes within a single degree Celsius will not be noticed by the inhabitants of a building. This value can be changed by the system operator and can be set individually for each consumer agent.

An important aspect of quality of service in regards to load control is to only ever try to control the radiator circuit within a building, and never the tap-water circuit. Besides directly affecting the comfort of the inhabitants, lowering the energy usage on the tap water circuit might also give rise to other problems such as growth of Legionella bacteria.

2.1.1 Allocating Resources by Auction

When the producer agent wants the system to perform load control it will start by analysing the size and time needed for the total load control, and based on this, the producer agent tries to distribute this change in energy usage among the consumer agents. This is done continuously as long the producer agent deems it necessary to perform load control. As the available heat load buffer in the individual buildings is drained, the system will try to re-allocate the load control among the consumer agents. This allocation is based on a first-price sealed-bid auction process [2]. The reason for using this type of auction is that all the agents within the system are programmed to be completely cooperative, so there would be no gain in using other, possibly more complex, auction processes. Each consumer agent will continuously calculate the amount of load control that the building can afford without jeopardizing the indoor climate. The parameter for this can be set for each individual agent, i.e. while one building can handle a temporary deviation of 1°C from the desired indoor temperature another building might only be able to handle a temporary deviation of up to 0.5°C. The calculation is done by a energy balance equation based on the geometry and characteristics of each individual building in combination with continuous sensor data from the heat exchanger station in that building. This value is then used by the consumer agents as currency when participating in the auction process. This basically means that the more suitable a building is to perform load control, i.e. a building with low or no deviation from the desired indoor temperature, the more currency, i.e. size of heat load buffer, it will have to spend on the auction, and thereby be more likely to win. Each such bid consists of three parts:

- The shifted outdoor temperature. The consumer agent uses shifted outdoor temperatures to make the heat exchanger station perform load control, e.g. by telling it the outdoor temperature is five degrees Celsius when in reality it is two degrees Celsius. This makes it easy to manipulate the behaviour of the heat exchanger station without extensive work and expensive equipment.
- How much the heat load usage will change.
- The amount of time this change can be implemented while staying within acceptable quality of service constraints.

These bids are then used by the system as a basis for computing how much of the heat load that can be shed and how large the total available heat load buffer is. This data can then be used in order to display real-time information about the system-wide status through a graphical user interface. When developing industrial implementations of research based systems it is important not to forget the importance of providing ways to offer direct interaction with human operators. After each completed auction the winner or, more often than not, winners will implement the load control according to their local prerequisites. The auction process basically consists of the following steps:

1. The producer agent identifies a need to perform load control.
2. An action is started were all participating consumer agents will submit bids
3. The producer agent distributes the load control among the consumer agents according to the bids. If the total need for load control is larger than the winning bid the producer agent will give the winning agent load control corresponding to the full bid, then to the second agent and so on until either the total need for load control is fulfilled or all the consumer agents are given load control corresponding to their full bids. In the latter case the total need of load control cannot be fulfilled using the available buildings.

In order to calculate how the dynamics of the indoor climate in the individual buildings each consumer agent uses a mathematical energy balance model. The agent continuously uses this in order to calculate its bids. The parameters for this model are unique for each building and is based on the geometry of the building and building materials in combination with continuous input and output of energy through the building. There is an obvious connection between the indoor climate and the energy input into the building, but due to the physical process of energy loss there is a substantial delay between energy input change and indoor climate change. This delay creates a time frame which the consumer agent can utilize in order to participate in load control without jeopardizing the perceived indoor climate. There is, in other words, a heat load buffer in each building and the size of this decides how much the building can participate in the system wide load control strategy. The energy balance model is used continuously by the consumer agent in order to calculate the future ability of the building to participate in load control as well as to calculate current status within the building and to perform controlled returns in energy input after a load control instance is finished. Such controlled return in energy input is needed in order to prevent the underlying control system to overshoot the desired indoor temperature by trying too hard to compensate for the energy input drop during the load control. The energy model uses outdoor temperature and radiator system temperatures as input, and is modelled as a system of differential equations which are numerically solved over and over again based on changes in the input.

2.2 Additional Hardware and Software

In order for the deployed system to work there was a need to develop additional hardware. The complete system is based on the system-wide multi-agent system software, the computing and communicating capabilities in the individual buildings and server systems for data handling and user interface. In particular the hardware for the individual buildings had to be developed. There were no systems available on the market that could provide the needed functionality while not being overly expensive. In order to handle the multitude of sensors present in different buildings a new hardware platform had to be developed. This took the form of an I/O card which could handle 10 inputs and 4 outputs, both analogue and digital, which was connected to a small computer that could handle the consumer agent software. This computer then used either normal Internet access or GPRS modems for communication.

A lot of companies and organizations use different types of network solutions in order to secure their networks, and it is not uncommon for communication to be a problem even if there exists the physical infrastructure for Internet access. In order to overcome such problems a Virtual Private Network (VPN) was set up, within which all the agents could communicate freely without disturbing underlying network structure. The VPN used within the project is based on OpenVPN, which is a full scale SSL VPN solution based on open source software.

3. DEPLOYED SYSTEM

The system described in this paper is by its very nature distributed and uses the combined heat load buffer within a range of buildings in order to achieve system wide advantages. In order to study this behaviour in a operational system it is thus important to have access to a large enough collection of buildings. During this project fifty-eight district heating consumer stations were connected through three separate multi-agent systems, Stockholm (27), Västerås (21) and Linköping (10). Each consumer station can serve several buildings, so the number of actual buildings participating in the system was greater than fifty-eight. The buildings in Stockholm and Linköping were mostly multi-apartment buildings of different sizes and types, while the buildings in Västerås were schools and other public buildings. This provided a good diversity among the buildings which made it easier to draw general conclusions based on the results.

3.1 Results

Normally a building heating system uses the outdoor temperature as a control signal. The agent system uses this outdoor temperature sensor as an interface towards the existing control systems in the building. By manipulating this sensor a consumer agent can force the control system to act according to the agents desire, without actually having to implement any changes on the existing control system. When the agent decides to implement load control it does this by faking the outdoor temperature signal, thus causing the building control system to lower or raise its heat load accordingly. Figure 1 shows an example of how this works in practice.

The difference between the supply and return temperature of the radiator system (Trad, supply and Trad, return) shows the energy usage. This value clearly drops when the faked outdoor temperature (Tout, LC) deviates from the actual outdoor temperature (Tout).

That the existing control system in a building will increase or decrease its heat load when it registers that the outdoor temperature changes is hardly controversial. The complexity of the process instead arises when the system tries to coordinate this behaviour among a group of buildings in order to achieve system-wide goals. In order to evaluate the system we have to analyse both the momentary heat load

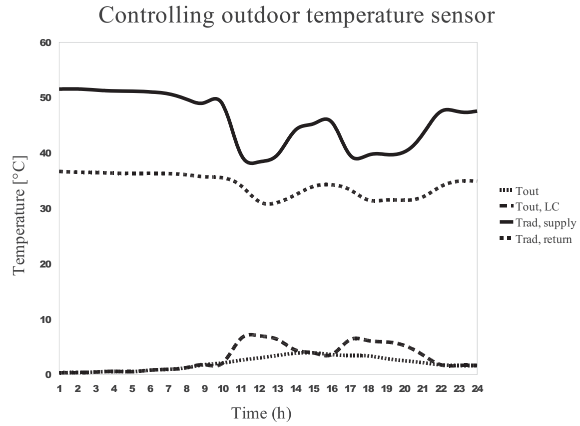

Figure 1: Controlling the energy usage behaviour through the outdoor temperature sensor

and the energy usage (heat load over time). Figure 2 compares the heat load on a day with two load control instances being shared among a group of consumer agents with a day without any load control. The first load control starts at about six o'clock in the morning and the other starts about six o'clock in the evening. Both load control instances take about two hours to complete.

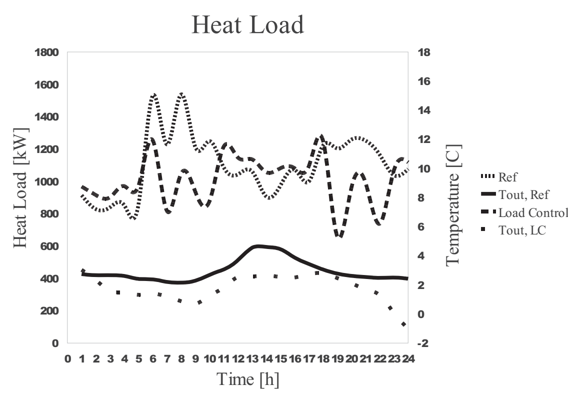

Figure 2: Heat Load during Load Control

Figure 3 shows the momentary heat loads from Figure 2 sorted according to size.

Figure 4 shows the energy usage in relation to the outdoor temperature in the same group of buildings during the same two days. The energy usage follows the same pattern as the heat load.

Figure 5 shows the same energy usage sorted according to size.

The outdoor temperature is slightly lower during the day with load control, so normally this day would have a higher energy usage. However, by using load control the total energy usage during the day without load control is 26578 kWh while the energy use during the day with load control is 24727 kWh. Despite the difference in outdoor temperature the saving in energy usage is still about seven percent.

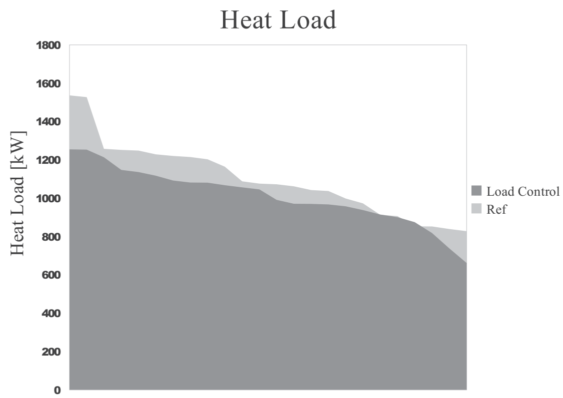

Figure 3: Heat Load sorted according to size

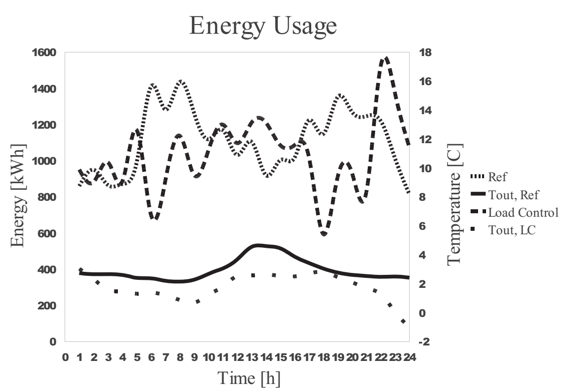

Figure 4: Energy usage during Load Control

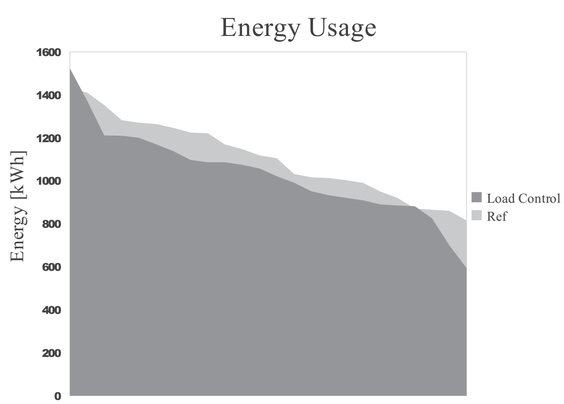

Figure 5: Energy usage sorted according to size

Figure 6 shows the total energy usage in a group of coordinated buildings during periods of different weeks, with one data point for each weekday of the week. The data for load control comes from a single week, and the data without load control comes from two reference weeks. The figure shows the energy use in relation to the outdoor temperature.

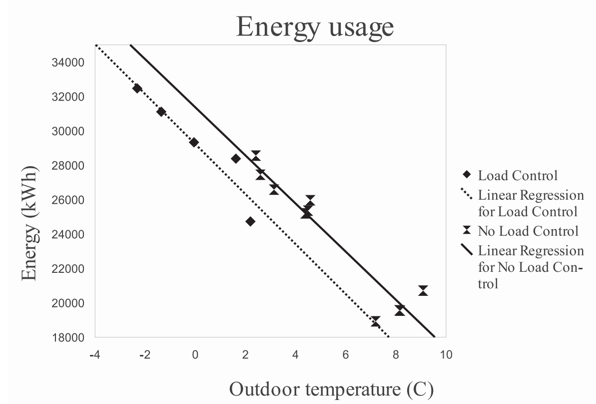

Figure 6: Total energy usage in an entire system of buildings

The values shown in figure 6 are actual measurements from the buildings. For the week with load control it is also possible to calculate what the energy usage should have been without load control, based on the historical energy performance of the building in relation to the outdoor temperature. The actual amount of energy used during the week with load control is 197 215 kWh and the calculated value for the same week but without load control is 213 352 kWh. This gives an energy saving of about 7,5%. This value is an average value from a large group of buildings with energy usage behaviour coordinated by a multi-agent system. The value of 7,5% is about the same as was shown in figure 5.

3.2 Indoor Climate

An important aspect of the system is to secure an adequate indoor climate at all time in connection with load control. The most basic measurement of this if the building occupiers have complained during time intervals with load control active. Another way to measure this is to install indoor temperature sensors and see what happens with the indoor temperature during load control. This was done during the project in selected buildings within the multi-agent system.

In practice the system will reduce the heat load in the building, which will inevitable lead to a reduced indoor temperature if allowed to continue uncontrolled. Each consumer agent is responsible for making sure that this reduction in indoor temperature is small enough for the building occupiers not to notice. There were no complaints during the project that were derived from load control.

The indoor temperature sensors showed no sign of being influenced by load control. The indoor temperature could vary wildly in individual apartments. This happens all the time and is due to social behaviour such as using electrical appliances, opening windows, or having ten kids over for a birthday party. It was not possible to determine when load control was active or not by studying the data from the indoor temperature sensors.

3.3 Available Load Control Ability

In order to evaluate the financial aspects of this type of system it is important for the energy company to be able to estimate the number of buildings need within the multi-agent system. The system-wide operational benefits of the system is directly proportional to the number and size of buildings available in relation to the total size of the district heating system. A large building has a greater available heat load buffer than a smaller building. Normally it is possible to estimate an average size of buildings within the district heating network, which can then be used as a basis for a first estimation of the potential of the system. It is the possible to make a rough estimate of the amount of buildings needed for a given district heating system:

$$Amount = \frac{Heatload}{eSig * (T_b - T_{out}) * LC_{max}} \quad (1)$$

Where $Heatload$ is the amount of heat load [W] that the system should be able to handle, $eSig$ [W/°C] is the average energy signature in the area, T_b [°C]is the outdoor temperature above which a building needs no heating, T_{out} [°C] is the outdoor temperature at time of using load control and LC_{max} is the maximal share of the total heat load that the system should be allowed to control. This last value, which is set to 70% for this study, is a general limit that is used in order to ensure that the system doesn't completely shut of the heating, not even during shorter periods of time. This is a social, rather than physical, consideration, e.g. if someone puts their hand on the radiator it should not be completely cold. This is the maximum value that the system can control of the heat load, although normally the system will control significantly less than this. In order to calculate how long this amount of buildings can uphold load control at each time the following equation can be used:

$$t = \frac{T_{diff}}{T_{in} - T_{out}} * (1 - e^{-1}) * Timeconstant \quad (2)$$

Where t is the time [h] that the system can uphold load control at a LC_{max} of 1,0, T_{diff} [°C] is the acceptable temperature drop in indoor temperature during load control, T_{in} [°C] is the indoor temperature at the beginning of activated load control, T_{out} [°C] is the outdoor temperature during the load control and $Timeconstant$ [h] is the time constant of the average building within the installation area. The value of $(1 - e^{-1})$ is derived from the defintion of the time constant for a building [1]. The time constant tells you how fast the indoor temperature of a building will drop when the heat load is totally shut off, given a nominal outdoor temperature, normally -20°C. The time constant is measured in hours, e.g in a building with a time constant of 100 it will take about 100 hours before the indoor temperature will have fallen $(1 - e^{-1})$, or about 63%, of the difference between the initial indoor temperature, e.g 21°C, and the nominal outdoor temperature. Finding the exact time constant of a building can be time consuming, but normally it is possible to use approximate values. During this project the following templates were used:

- Light building: 80h (light construction)
- Semi-light building: 150h (light/semi-light construction, concrete grounding)
- Heavy building: 300h (heavy construction)

Calculating the amount of time the system can uphold each turn of load control basically means to estimate how long it takes for the indoor temperature to drop below the acceptable limit. In practice the system will have a certain capability to enforce load control, and this capability will diminish as the individual buildings exhaust their heat load buffers. If a larger group of buildings are available through the agent system it is possible to uphold load control during longer periods of time since the buildings can relieve each other as they drain their heat load buffer. By using the auction process, this behaviour automatically manifests itself when a large enough group of buildings is connected.

The ability to perform load control is dependant on the total level of heat load in the buildings. If there is a high level of load control, then the ability to perform load control is equally high. This means that the ability to perform load control increases as the outdoor temperature drops, which in turn means that the ability to perform load control is at its highest when the need for it is the greatest.

3.4 Cost-Benefit Analysis

The largest financial value of lowered heat load is probably when load control ability will help to avoid the need for more peak load production. The amount of heat load that can be moved or lowered is dependant on the characteristics of both the production units and the building types available in the district heating system. Normally the need to perform load control only occurs during shorter periods of time. Normally 30-40% of the heat load has a length of less than 500 hours. In most district heating system these loads are produced by oil, electricity or coal. An oil boiler costs about €200/kW ($270/kW) to install, which is the financial value of the load control system if this installation capability can be avoided. During this project the load control system has lowered heat loads during peak hours between 15-20% on average. Avoiding 15% oil based peak load in a 100 MW system is then worth about €3M ($4.1M) in saved installation costs.

Transferring peak load to base load also means lowered variable production costs. As an example the ability to move from tree oil (cost €0,04/kWh ($0.055/kWh) including taxes) to bio-fuel (cost about €0.02/kWh ($0.027/kWh) including taxes) gives a saving of €0.02/kWh ($0.027/kWh). Fortum, who owns the district heating network in Stockholm, produces slightly less than 10 000 GWh annually. The value of reducing fossil fuel in production is not only a financial one, but also based on environmental considerations. For security reasons it is unlikely that peak load production can be avoided all together, but the benefits of reducing it as much as possible are numerous.

Load control not only gives the ability to lower the heat load, but also to move it in time. In this case the buildings can be viewed as the equivalence of the large storage tanks found in many district heating systems. When using combined heat and power production (CHP) it is normal to optimize the production in relation to the price variation on the spot-price market for power. The spot-price varies during the day but normally there are price peaks in the morning and evening. The physical process of producing heat and power cannot be separated, but by using load control it is possible to move the heat load consumption a few hours in time by controlled pre-heating of the buildings.

4. DISCUSSION

Performing uninformed local load controls by manipulating the outdoor temperature sensor has been shown to be easy enough, although due to aspects of dynamic (time dependant) indoor climate feedback a certain level of intelligent behaviour is required when actually implementing energy saving process using such techniques. More interesting and complex problems arise when trying to coordinate this behaviour from the system wide aspect of the producer, while still considering the indoor climate in all the individual buildings. When sorting the data from figure 2 according to size like in figure 3 it becomes apparent that the largest reductions in heat load have been done during the higher peaks. It should be noted that the outdoor temperature during the reference day is +2,9 °C and during the day with load control +1,7 °C. No adjustment for this temperature difference have been made so the effect of the load control is in reality even more significant than what is shown in this example.

In order to satisfy the heating need an energy company would normally use a whole range of different production units, with different characteristics and fuels. The production units are started in order by their production costs, which normally means that the boilers using oil will only be started during periods of peak load with short duration. The absolutely highest peak loads normally have durations of a few hours. During this project it was noted that in order to lower the heat load with 10% the system needed to control the load during no more than eight hours at most. This is well within the capability of the system studied.

The theoretical optimization normally used in order to find the financially best operational strategy for the production units is based on solutions of the Economic Dispatch Problem (EDP) and the Unit Commitment Problem (UCP) [5]. By solving the EDP and the UDP an energy company can find a desired consumption level for each hour during the next few days. This value can then be used as decision data by the multi-agent system in order to implement load control throughout the day in order to uphold the desired consumption level.

5. CONCLUSIONS

This system consists of distributed units with communication ability that can measure sensor data and control the heating system in individual buildings. These units are then controlled by consumer agents which enables the system to coordinate the heating behaviour on a system wide scale. By using this system a energy company can, within the constraints set by indoor climate limits, actively plan, optimize and control the heat load within a district heating network.

The results show that the system has lowered the total heat load with about 20% in connection with the highest peak loads. This value is, however, still far from the maximum set limit of LC_{max} in equation 1, i.e. the limit of 70% of the heat load which is used in this study. In figure 3 it is shown that the largest heat load reductions are done during time intervals with the highest heat load (measured in kW). A district heating network utilizes several differently priced fuels in order to satisfy the demand, ranging from cheap waste heat and bio fuels to more expensive alternatives such as oil and electricity. The normal situation in Sweden is that the highest 5-15% of the heat load needs to be satisfied with some peak load boiler (usually oil)[3]. Compare this with figure 3 where the highest heat loads are sorted left to right. It is clear that the agent system manages to shed these peak loads in the left-most part of the figure.

By using correct pricing strategies within the district heating network the energy company wants to achieve transparency so that the costumers understand and act based on the costs of producing the heat. This is a very complex problem and it is not always easy to make the pricing strategy understandable as well as transparent in regards to actual production costs. Also, it is not easy to set a price according to different conditions since if the customers actually react to this price it would change the conditions. This leads to a situation with a variable pricing goal which in practise demands that it is possible to dynamically adjust prices in order for the right production situation to arise. By a correct pricing strategy it should be simple to give the customers the incentive to invest in load control equipment. It is, however, important that this equipment is capable of interacting on a system-wide scale, since local uninformed load control can do more harm the good from a production perspective. If there is no method for system-wide control it is very hard, if not impossible, to use the behaviour of the buildings in order to optimize the production. It is also not enough to use static predefined lists of buildings to cut heat load in during times of peak loads. Using such predefined lists does not take the dynamics of the indoor climate into consideration. It is essential to pay attention to the dynamics of the quality of service as well as the production status.

In relation to this discussion it should be noted that it is not suitable to let human consumers themselves make the dynamic decisions to implement load control, e.g. by having a wall mounted display showing graphs of energy consumption and then expecting people to actually do something when the energy price is high or the heat/power load is above some threshold. Such schemes have been tried several times and the results are normally the same, i.e a consumer will use the system for a month or two, but will then eventually grow tired of the whole thing and stop using it. We believe that the only way to implement such systems in the long term is to exchange the human decision maker with an automated system of some sort.

Besides reducing peak loads the system has also shown a clear ability to lower the energy usage in the participating buildings, without any noticeable difference in the perceived indoor climate. During the project the system showed a energy reduction of about 7,5% during week-long periods. From the perspective of the energy company and the national economy it is preferable if this reduction in energy consumption in connection to peak load production instead of during base load production. Only by using system-wide load control can such a consumer behaviour be enforced.

The system controls the heat load in individual buildings by manipulating the outdoor temperature sensor that the existing control system uses. Each load control is implemented by faking the outdoor temperature for the existing control system, e.g. if the existing control system during a short period of time is led to believe that the outdoor temperature is 10 °C when in reality it is 5 °C, then this will lead to a reduction in energy usage during that time. A lowered heat load will obviously led to a lowered indoor temperature sooner or later, which is why such load control has to be done within a controlled process. Several indoor temperature sensors were placed within the buildings during the project, in order to identify lowered temperatures due to load control. Despite of this it has not been possible to identify where and if any reduction of the indoor temperature was present. It was shown that the indoor temperature within individual apartments varies so much due to social behaviour that the influence of load control is lost in the noise. In all likelihood the lowered energy usage must have led to a temporary reduction in the temperature, either in the building body or in the air contained within the building, but neither the indoor temperature sensors or any building occupiers have been able to perceive this.

In a heating system where the ventilation has a considerable impact on the indoor climate it is complicated to implement load control by manipulating the secondary circuit of the heating system. An alternative in these cases is to only control the heat load in the radiator system and not to include the ventilation hot water in the load control.

During this project we have studied load control within district heating systems, but these system-wide techniques are equally interesting in other energy systems, e.g. power grids or industrial production processes. When, for example, comparing a power grid and a district heating system it is obvious that there are major differences in the physical process of how energy is transported (water versus electricity), but on a system-wide scale we believe that there are still many overlapping themes.

This project has been a step up from previous small scale experiments and simulations. It has been shown that the theory holds in practice, and that coordinated load control within district heating systems as a technique does indeed work.

6. FUTURE WORK

Early small-scale practical experiments have shown a 4% decrease in energy usage when using a multi-agent based system for load control in a district heating network, and has predicted a theoretical energy saving of about 10% [8]. During this study we have shown an energy saving of about 7.5% averaging over groups of buildings. With further tweaking of the behaviour of the system we expect this figure to be able to reach the predicted 10%. This might be done by selecting other techniques for allocating the resources. Using an auction based mechanism for this implies that the consumer wishes to spend as little as possible of its utility, when in reality all agents within the system are cooperating in achieving the desired load control, i.e a consumer wants to spend as much as possible of its available buffer but never more than this. This might be interpreted as a suggestion that other methods could be even more successful at distributing the resources among the agents in such a system, e.g. bargaining schemes or some sort of distributed/centralized optimization technique.

7. ACKNOWLEDGEMENTS

This work as been a collaboration between NODA Intelligent Systems AB and Blekinge Institute of Technology. The results are based on a project financed by NODA Intelligent Systems AB, Fortum Värme AB, Mälareenergi AB, Tekniska Verken i Linköping AB and The Swedish District Heating Association.

8. REFERENCES

[1] C. Barley, P. Torcellini, and O. Van Geet. *The Van Geet Off-Grid Home: An Integrated Approach to Energy Savings*. Buildings Technologies Program. National Renewable Energy Laboratory, 2004.

[2] S. Clearwater. *Market-Based Control: A Paradigm for Distributed Resource Allocation*. World Scientific Publishing Company, 1996.

[3] N. Constinescu. *District Heating and Cooling: Country by Country, 2007 Survey*. Euroheat and Power, 2007.

[4] I. D. Heating and Cooling. Optimised district heating systems using remote heat meter communication and control. *Annex VI report*, 2002.

[5] C. A. Koay, L. L. Lai, K. Lee, H. Lu, J. B. Park, Y. H. Song, D. Srinivasan, J. G. Vlachogiannis, and I. K. Yu. Applications to power system scheduling. In K. Y. Lee and M. A. El-Sharkawi, editors, *Modern Heuristic Optimization Techniques*. The Institute of Electrical and Electronics Engineers, Inc, 2008.

[6] H. Nordvik and P. E. Lund. How to achieve energy efficient actions as an alternative to grid reinforcement. In *Proceedings of ECEEE Summer Study*, 2003.

[7] F. Wernstedt. *Multi-Agent Systems for District Heating Management*. Blekinge Institute of Technology, 2005.

[8] F. Wernstedt, P. Davidsson, and C. Johansson. Demand side management in district heating systems. In *Proceedings of Sixth International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, Honolulu, Hawaii, USA, 2007.

[9] F. Wernstedt and C. Johansson. Intelligent distributed load control. In *Proceedings of The 11th International Symposium on District Heating and Cooling*, Reykjavik, Iceland, 2008.