Embedded Agents for District Heating Management

Paul Davidsson and Fredrik Wernstedt
School of Engineering, Blekinge Institute of Technology
372 25 Ronneby, Sweden
{paul.davidsson, fredrik.wernstedt}@bth.se

Abstract

We investigate the applicability of multi-agent systems as a control approach for district heating systems. The consumers, i.e., the heat exchange systems, in current district heating systems are purely reactive devices without communication capabilities. In this work, the possibilities of a new type of heat exchanger system that has an open software environment and communication capabilities are explored. Operators of district heating systems have several, often conflicting, goals, e.g., to satisfy the demand of the customers and to minimize production costs. Our approach is to embed a cooperative agent in each consumer. Results from a simulation study indicate that the approach makes it possible to reduce production while maintaining the quality of service. In another experiment in a controlled physical environment, two agent-based approaches are evaluated and compared to existing technologies. The experiment shows that it is possible to automatically load balance a small district heating network using agent technology.

1. Introduction

District heating systems are by nature distributed both spatially and with respect to control. In current systems each substation (i.e., Heat Exchanger System at the consumer side) can be viewed as a "black-box" making local decisions without taking into account the global situation. Thus, today a district heating network is basically a collection of autonomous entities trying to optimize operations locally, which typically results in behavior that is not globally optimal.

ABSINTHE (Agent-based monitoring and control of district heating systems) [1] is a collaboration project between Blekinge Institute of Technology and Cetetherm AB, one of the world-leading producers of district heating substations. The goal of the project is to improve the monitoring and control of district heating, e.g., by increasing the knowledge about the current and future state of the network at the producer side and by performing automatic load balancing at the customer side. In the project, a new type of substation has been developed that has computing and communication hardware. We use this to equip each substation with a software agent that will enable it to cooperate with other substations to perform automatic load control, as well as with the producers to optimize production.

In the next chapter the district heating domain will be described. This is followed by a presentation of the multi-agent system used. Two sets of experiments are then described; a series of simulation experiments and a series of experiments on a real district heating system. Some reflections and pointers to future work conclude the paper.

2. District heating systems

The basic idea behind district heating is to use local heat production plants to produce hot water (or steam). The water is then distributed at approximately 1-3 m/s through pipes to the customers where it may be used for production of hot tap water and household heating. The cooled water then returns to the production plant forming a closed system (see Figure 1).

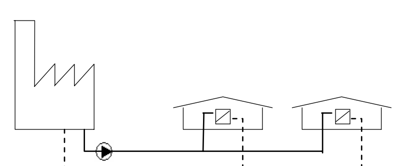

Figure 1. A simple district heating network.

At the customer side, there is a *substation*. As illustrated in Figure 2 a substation is normally composed of two (sometimes three) heat exchangers and a control unit, which receives hot water from the district heating network. The substation heats both cold tap water and the water in the household heating circuit by exchanging the required heat indirectly from the primary flow of the distribution network. The hot network water is returned to the network at a somewhat lower temperature. Both the temperature of the returning water and the flow rate in the network are dependent on the consumption of substations. When the water returns to the heat production plant it is heated and again pumped into the distribution network.

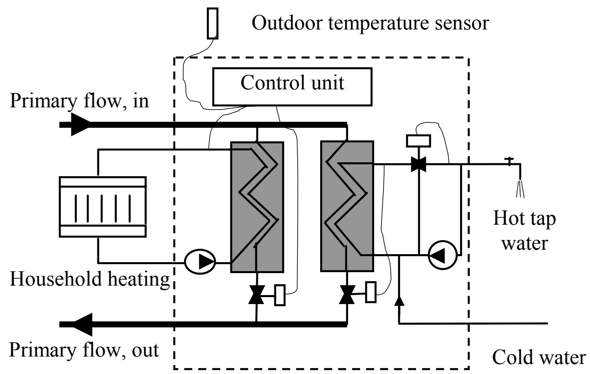

Figure 2. A substation consisting of heat exchangers (the shaded boxes), control valves, pumps and a control unit. The household heating system (radiators) is controlled by the control unit, using information about actual outdoor temperature.

Several different energy sources may be used for heat production, e.g., waste energy, byproduct from industrial processes, geothermal reservoirs, otherwise combustion of fuels as oil, natural gas etc. is used. If the demand from the customers is high several heat producing units must be used. A district heating system in a large city can be very complex, containing thousands of substations and hundreds of kilometers of distribution pipes resulting in distribution times up to 24 hours. In addition, they are dynamical as new substations may be added or old substations may be replaced by new ones with different characteristics.

2.1. District heating management

Current approaches to district heating management are centralized and reactive. The operator basically has two variables to balance, supply temperature and flow, when assuring sufficient resources for the demand of customers. However, a high supply temperature causes increased heat losses during the distribution as well as reduced production capacity in many production facilities. On the other hand, a large flow implies high pump costs as well as problems with the control of production facilities [2]. The load of a district heating system is mainly a consequence of the customers' demand for household heating. As a result, operations of most district heating systems are based on a simple mapping between the current ambient temperature and the supply temperature. Note that this approach is purely reactive and does not take into account future events that may be possible to anticipate. Furthermore, in order to ensure sufficient heat supply, the tendency has been to produce more heat than necessary and hence a waste of energy [3, 4, 5].

A more advanced approach to decide the supply temperature would be to use an optimization model. In a general optimization model the network appears as a set of constraints (where consumers have fixed and given demand), and the objective function is composed of cost for production. However, an optimization model of a large district heating system with many loops and more than one heat production plant is extremely computationally demanding [6, 7, 8]. In fact, it is argued that when the complexity of a district heating system reaches approximately 100 components and restrictions, the present computer and software technology is insufficient to find an optimum operational strategy [7].

A general argument against centralized approaches for problems as complex as the management of district heating systems is that when the problems are too extensive to be analyzed as a whole, solutions based on local approaches often allow them to be solved more quickly [9]. Furthermore, it would not be easy to collect and use the sensor information from each entity in the network in a centralized fashion. Substations, pumps and valves, etc., are typically manufactured by different organizations, i.e., different entities in the system often have different design of interfaces to the system. It would be a complex task to keep track of these aspects centrally. To develop local monitoring and control software adapted for each type of substation, but with the same interface to the rest of the software system, seems as a more fruitful approach.

Demand Side Management, i.e., actions taken on the customer's side to change the amount or the timing of energy consumption, has primarily been used by electric utilities to reduce peak loads [10]. For district heating systems, on the other hand, the most advanced approach is to use a local flow regulator which makes simple prioritizing, i.e., reducing the primary flow to the household heating system during hot water tapping. However, the effects of such reduction strategies are limited since the flow needed for the hot water tapping typically is larger than the flow caused by the household heating.

The consumption in a district heating network is mainly composed of two parts [11]:

- The heating of buildings, which mainly is a linear function of the outdoor temperature.
- The consumption of hot tap water, which mainly is dependent on consumption patterns, e.g., social factors.

The hot tap water consumption of a substation is very "bursty" even in large buildings, and therefore very difficult to predict, whereas the household heating is "smoother" and therefore easier to predict assuming that reliable weather predictions are available. Also, since the heating of buildings is a slow process, we should be able to make time-limited reductions to the flow caused by household heating without significantly affecting the comfort of the end users [10, 12].

2.2. IQHeat optimal 100

Due to the rising demand of automation of building services (heating, ventilation, and air-conditioning etc.) Siemens have developed the Saphir, an extendable I/O platform with an expansion slot for a communication card, suitable for equipment control. A Rainbow communication card can be used to get easy and fast access to the Saphir (see Figure 3).

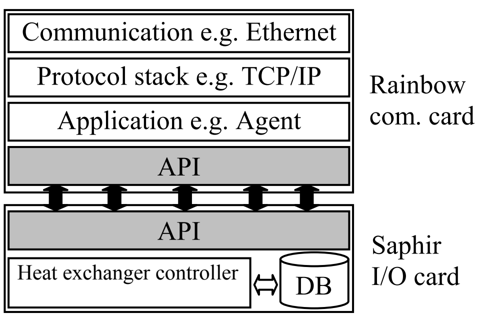

Figure 3. The Rainbow communication and computation card is here shown on top of the Saphir hardware interface card.

The Saphir contains a database that continuously is updated with sensor data from the I/O channels by a small real-time operating system, which is directly accessible from the Rainbow card. On the Rainbow card a small computational platform (a handheld PC) makes it possible to easily deploy software and by that providing the possibility to host an agent. Hence, the deployment of an agent on such a platform enables it to potentially read all connected sensor input as well as send commands over the I/O channel to actuators on the hardware, e.g., the valves of a heat exchanger. It has previously not been possible to develop, or make commercially available, such an advanced platform due to high costs.

The Saphir platform and the Rainbow communication card has been integrated into a new type of substation (IQHeat Optimal 100) developed by Cetetherm AB during the ABSINTHE project.

3. MAS architecture

We used an agent-based approach to implement a more sophisticated district heating management strategy. In our approach, each substation not only informs the producer about the current consumption, but also makes predictions of future consumption that are reported to the producer. Since local predictions typically are more informed than global predictions, this approach should give better results. The MAS architecture we suggest below also introduces a means for automatic redistribution of resources. In order to solve the problem of producing the right amount of resources at the right time, each consumer is equipped with an agent that makes predictions of future needs that it sends to a production agent through a redistribution agent.

The other problem, to distribute the produced resources to the right consumer, is approached by forming clusters of consumers within which it is possible to redistribute resources fast and robustly. This usually means that the consumers within a cluster are closely located to each other. In this way it is possible to cope with the discrepancies between predicted and actual consumption. An important requirement for a redistribution mechanism of this kind is that it can deal with shortages in a way that is fair to the consumers.

We have chosen a semi-distributed approach since a completely centralized approach may result in communication bottlenecks without achieving greater fairness. Also, each message would need to travel a longer route, which would increase the total network load. A completely distributed approach, on the other hand, would result in a larger network load than the semi-distributed approach without achieving greater fairness.

Based on the above insights, we used the GAIA methodology [13] to design the MAS. This led us to an architecture that has the following three types of agents:

- *Consumer agents*: (one for each consumer) which continuously (i) make predictions of future consumption by the corresponding consumer and (ii) monitor the actual consumption, and send information about this to their redistribution agent.

- *Redistribution agents*: (one for each cluster of consumers) which continuously (i) make predictions for the cluster and send these to the producer agent, and (ii) monitor the consumption of resources of the consumers in the cluster. If some consumer(s) use more resources than predicted, it redistributes resources within the cluster. If this is not possible, i.e., the total consumption in the cluster is larger than predicted, it will redistribute the resources available within the cluster according to some criteria, such as, fairness or priority, or it may take some other action, depending on the application.
- *Producer agents*: (one for each producer, however, we will here only regard systems with one producer) receives predictions of consumption and monitors the actual consumption of consumers through the information it receives from the redistribution agents. If necessary, e.g., if the producer cannot produce the amount of resources demanded by the consumers, the producer agent may notify the consumers about this (via the redistribution agents).

The interval between predictions is larger than between redistributions, i.e., during each prediction interval there may be a number of redistributions performed. Each consumer agent produces one prediction for each prediction interval and sends this to its redistribution agent, who sums the predictions of all consumer agents belonging to the cluster and informs the producer agent about this. The predictions made by a consumer agent must reach the producer at least by a time corresponding to the production plus the distribution before the resources are actually consumed (the distribution time is individual for each consumer). Typically, there is also a production planning time that also should be taken into account.

4. Simulation experiments

The MAS described above was first evaluated in a series of simulation experiments [14], which will be described in this section. In the next section, some experiments on an actual district heating system will be described.

4.1. The simulation software

The MAS as well as the simulated environment was implemented in JADE (Java Agent DEvelopment framework) [15]. Thus, we used an agent-based approach that was time-driven, i.e., were the simulated time is advanced in constant time steps, to simulate the environment. Each simulated entity was implemented as a separate agent.

A model of hot tap water usage, implemented by Arvastson and Wollerstrand [16] and based on detailed field measurements performed in Sweden by Holmberg [17], was used to generate hot tap water consumption values. A common resistance/capacitance model [18] was used for simulation of the energy needed for household heating. Figure 4 shows an example of the total consumption during a day using these models.

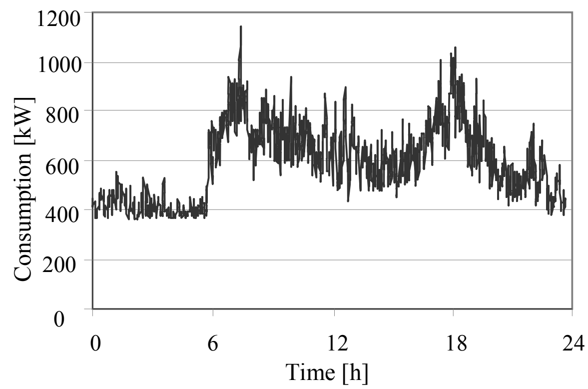

Figure 4. Example of the total consumption of 500 apartments during one day.

The predictions of consumption for each substation were calculated by averaging the consumption for that particular ten minutes interval for that particular substation of the last five days.

4.2. Experiment design

Two series of simulation experiments have been made. In each experiment there were one producer agent and one redistribution agent (since a cluster is independent of other clusters, we do not need to simulate more than one cluster). In each experiment the district heating system was simulated for 24 hours. Each experiment was run over 5 different days (different series of consumption values) and the averages are shown in the diagrams below.

In the first series of experiments, the redistribution agent was managing a cluster consisting of 10 consumer agents (substations), where five of the consumer agents were serving a building with 40 apartments and five consumer agents serving a building with 60 apartments. It was assumed that the distribution time from the producer to the consumers is 1 hour and that there is a single production source. We ran experiments on different degrees of surplus production (from 0% surplus production, in steps of 1%, to 5%), where surplus production is defined with respect to the predicted consumption. For example, if the predicted total consumption is 200 kW and the surplus production is 2%, 204 kW is produced. The quality of service was measured in terms of the number of restrictions issued by the redistribution agent.

We also compared the suggested approach to a reference control scheme, which, we believe is a very optimistic approximation of the current production strategy for district heating. Also here, different degrees of surplus production were tested. The production of the reference control scheme is shown in Figure 5.

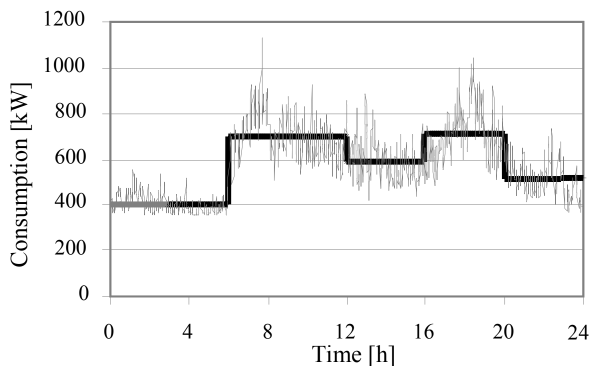

Figure 5. The amount of hot tap water produced by the reference control scheme (indicated by the thick line, 0% surplus production), compared to the 5-day average consumption.

In the second series of experiments, we varied the size of the clusters in order to study how the cluster size affects the quality of service. The number of consumer agents in the cluster that we studied were 2, 4, 8, and 16, where each consumer agent served a building with 40 apartments.

4.3. Simulation results

We found that the multi-agent system described above performs well, coping with faulty predictions (even though the discrepancy between the predicted and the actual consumption sometimes is quite large). Figure 6 shows the total number of restrictions to hot tap water and the number of restrictions to household heating (radiators) during one day for different degrees of surplus production. We see that there is a clear trade-off between the quality of service (number of restrictions) and the amount of surplus production and that there are almost no restrictions of any kind when the production is 4% larger than the predicted consumption.

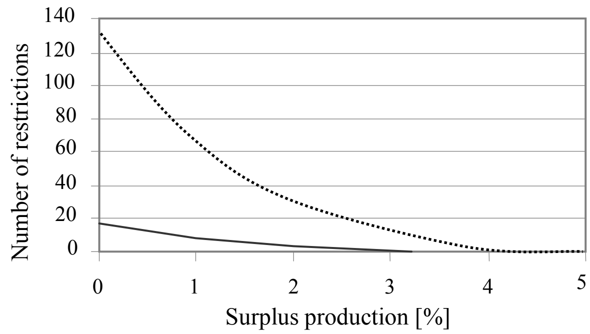

Figure 6. Trade-off between quality of service and surplus production. The dashed line corresponds the number of restrictions to household heating and the other line to the number of restrictions for the hot tap water.

Using the reference control scheme, 7% surplus production was needed to achieve this level of quality of service. It should be noted that it is in current practice very difficult to find this minimum degree of surplus prediction. This is because the operators of district heating networks typically do not have exact information on neither the quality of service delivered to the customer, nor the amount of the actual degree of surplus production.

The second experiment showed that there is an almost linear reduction in the number of restrictions when the size of the cluster increases. However, it should be noticed that the upper limit of the cluster size is determined by the structure of the district heating network. If the actual distance between the substations within a cluster is too large, the assumption of robust and fast redistribution does not hold.

5. Small scale in situ experiment

In the experiments we used two heat exchange systems (600 kW and 400 kW respectively) and one heat producing unit. As described earlier the substations (IQHeat Optimal 100) that were used are developed by Cetetherm AB and contain Saphir hardware and a small real-time operating system. The production and the consumption were controlled by a computer, which also was responsible for making all measurements during the experiments.

The consumer agents are able to reduce the consumption at a substation by changing the temperature set value for the household heating water circuit. However, they may maximally reduce it by

15% from the initial value used by the control circuit. A reduction by 15% may seem small but will actually often result in a temporary shut-down of the household heating system (since the water returning from the radiators is sufficiently warm). As a result, the consumption will start to decrease immediately, but only at a relatively slow rate.

5.1. Experiment design

In each experiment the scenario illustrated in Figure 7 was utilized.

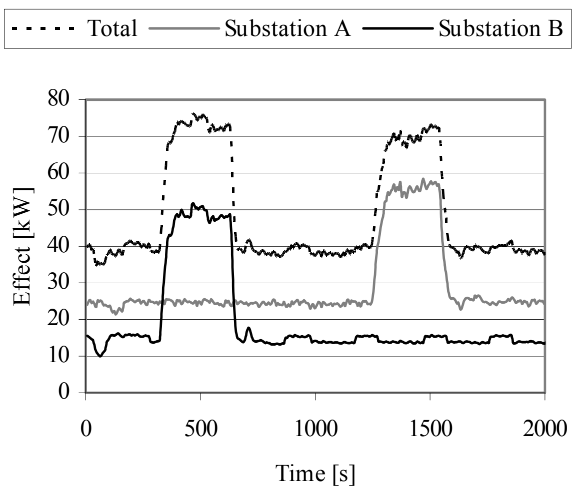

Figure 7. Household heating and hot water tapping from two substations.

Two substations, A and B, are set to have a constant household heating demand. Substation A has a set value for household heating circuit temperature of 48C° (approximately 25kW) and Substation B has a temperature set value of 51C° (approximately 15kW). The system is first allowed to reach a steady state during five minutes. After five minutes Substation B initiates a hot water tapping of 0,2 kg/s for a duration of five minutes. The system is then given ten minutes to stabilize. After the stabilization period Substation A initiates a hot water tapping of 0,2 kg/s, also with a duration of five minutes. After the second hot water tapping the system is given time to stabilize before shutdown.

In this experiment four different approaches where compared. In addition to the hierarchical MAS evaluated in the simulation experiment, a decentralized MAS without a redistribution agent was considered, together with two non-cooperative approaches. The four approaches are described in table 1. For more details on the in situ experiments see [19].

Table 1: Control strategies

No restrictions	Substations are free to consume the amount requested.
Local restrictions	Consumer agents individually enforce reductions when the consumption reaches a specific limit. The limit for substation A is 30 kW and the limit for substation B is 20 kW.
Hierarchical MAS	The redistribution agent uses the aggregated limit of the two substations, i.e., 50 kW, as a global limit. The redistribution agent requests reductions of both substations when system consumption reaches the global limit.
Decentralized MAS	Consumer agents individually enforce reductions when the consumption reaches a specific limit. The limit for substation A is 30 kW and the limit for substation B is 20 kW. When enforcing a local reduction the consumer agent also requests reduction assistance from the other substation.

5.2. Experimental results

Figure 8 shows the total energy consumption for the four different control strategies.

We see that the currently most advanced strategy, to use local restrictions, clearly reduces the consumption peaks and that the strategies to use hierarchical as well as decentralized multi-agent based approaches reduces the peaks even further. However, the MAS approaches require slightly more time to reach the desired level after the restriction period. Note also that the strategy with local restrictions is unable to reach the desired level (50 kW) during the 5 minutes of high consumption, whereas the MAS approaches reach this level after approximately 3 minutes.

A factor that slows down the control is the way that the consumption is measured. Ideally, the agent should be able to measure the (accumulated) consumption at any time. However, when using the existing sensors of the substation, it is only possible to discover changes in consumption at the same rate as the rate of incoming pulses from the flow gauges (which could be less than one per minute when the flow is low). The agents are instead using extrapolated consumption values over an interval of one minute. Unfortunately, the result is that agents have a limited view of their environment, making the reaction time for the reductions longer.

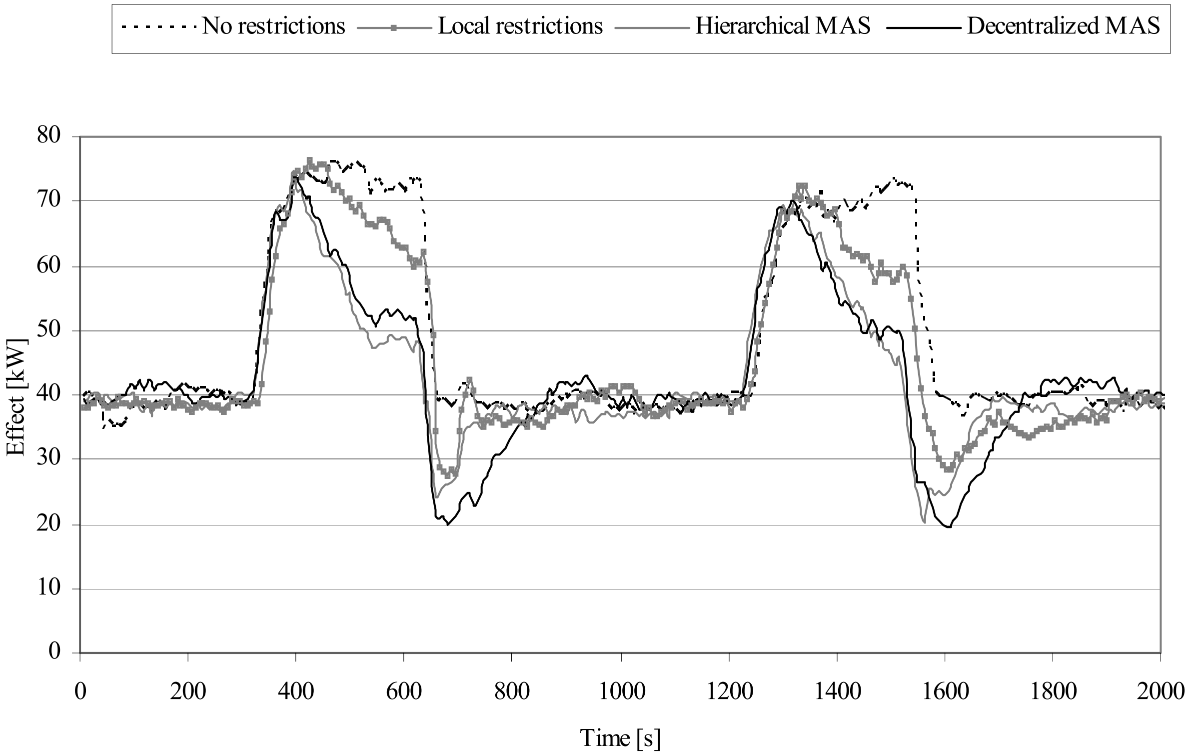

Figure 8. Total system energy consumption for the four control strategies. The desired global consumption is 50 kW.

6. Conclusions

We have suggested a MAS approach to district heating management. The approach has been evaluated in both simulation experiments and in situ experiments in an actual (but small) district heating system.

The simulation experiments showed that the suggested approach to district heating management makes it possible to control the trade-off between quality-of-service (measured in terms of the number of restrictions) and degree of surplus production of resources. Compared to a reference control scheme, it is possible to reduce the amount of resources produced, from approximately 7% to 4% surplus production, while at the same time maintaining the quality of service. The simulation results also indicate that the larger the cluster size, the higher is the quality of service that can be achieved. However, the number of restrictions decreases only logarithmically as the cluster size increases. Also, there are often other factors, such as distance between consumers that limits the number of consumers that can belong to the same cluster.

The small-scale in situ experiment confirmed that it is possible to automatically load balance district heating systems. To our knowledge, agent technology has never been used for monitoring and control of district heating networks. There have been experiments performed with centralized control of substations [8], however we show that we can achieve distributed concurrent automatic load balancing by the use of agent technology.

A substation (IQHeat Optimal 100) containing the Saphir and the Rainbow communication card is currently in production and available on the market. The agent-based control system (IQHeat Optimal 200), founded on the work described here is planned to be released by Cetetherm AB during the autumn of 2005.

Future work includes:

- Trying out other restriction policies than fairness, for example, based on priority.
- Improve the prediction mechanism, for instance by using neural networks [4, 20].
- Trying out strategies where producers are able to impose pro-active restrictions to cover for future temporary heating needs (peak-shaving).

- Finalize on-going experiments with the agent-based Optimal 200 control system in the district heating system of the city of Karlshamn in Sweden.

Acknowledgements. This work has been financially supported by VINNOVA (The Swedish Agency for Innovation Systems). We also wish to acknowledge the valuable input from the ABSINTHE project members employed at Cetetherm AB. Finally, we would also like to thank Markus Bergkvist for assistance during the in situ experiments.

References

1. ABSINTHE project web site: www.ipd.bth.se/absinthe, 2004-01-21.
2. Larsson, G., Rossing, O., "Flödesutjämnande körstrategi" (In Swedish), *Forskning och Utveckling, FOU 2003:86*, ISSN 1402-5191, 2003.
3. Arvastsson, L., "*Stochastic modeling and operational optimization in district heating systems*", Lund Institute of Technology, Sweden, ISBN 91-628-4855-0, 2001.
4. Canu, S., Duran, M., and Ding, X., "District heating forecast using artificial neural networks", *International Journal of Engineering*, Vol. 2(4), 1994.
5. Petersson, S., Werner, S., "Långtidsegenskaper hos lågflödesinjusterade radiatorsystem" (In Swedish), *Forskning och Utveckling, FOU 2003:88,* ISSN 1402-5191, 2003.
6. Aringhieri, R., Malucelli, F., "Optimal operations management and network planning of a district heating system with a combined heat and power plant", *Annals of Operations Research*, Vol. 120: 173-199, 2003.
7. Bøhm, B., Lucht, M., Park, Y., Sipilä, K., Ha, S., Won-tae, K., Bongkyun, K., Koljonen, T., Larsen, H., Wigbels, M., and Wistbacka, M., "Simple models for operational optimisation", *Report 2002:S1, NOVEM*, ISBN 90-57-48021-2, 2002.
8. Tamminen, E., Wistbacka, M., "Capacity and cost models for thermal power systems with random outages of plants", *VTT Energy, Espoo, Research Report ENE6/44/01*, 2001.
9. Rinaldo, J., Ungar, L., "Auction-driven coordination for plantwide optimization", in *Proceedings of Foundations of Computer-Aided Process Operation, FOCAPO*, 1998.
10. Drysdale, A., Stang, J., "Optimised DH systems using remote heat meter communication and control", in *Proceedings of the Eight International Symposium on District Heating and Cooling*, 2002.
11. Werner, S., "Dynamic heat loads from fictive heat demands" (in Swedish), *Fjärrvärmeföreningen, FOU 1997:10*, ISSN 1402-5191, 1997.
12. Österlind, B., "Effektbegränsning av fjärrvärme" (In Swedish), *Rapport 63:1982, Byggforskningsrådet*, ISBN 91-540-3714, 1982.
13. Wooldridge, M., Jennings, N.J., and Kinny, D., "The Gaia Methodology for Agent-Oriented Analysis and Design", *Journal of Autonomous Agents and Multi-Agent Systems,* Vol. 3(3): 285-312, 2000.
14. Davidsson, P., Wernstedt, F., "A Multi-Agent System Architecture for Coordination of Just-in-time Production and Distribution", *The Knowledge Engineering Review*, Cambridge University Press, Vol. 17(4): 317-329, 2003.
15. Bellifemine, F., Poggi, A., and Rimassa, G., "Developing Multi-agent Systems with a FIPA-compliant Agent Framework", *Software: Practice and Experience*, Vol. 31(2): 103-128, John Wiley & Sons, 2001.
16. Arvastson, L., Wollerstrand. J., "On Sizing of Domestic Hot Water Heaters of Instantaneous Type", in *Proceedings of the 6th International Symposium on Automation of District Heating Systems*, 1997.
17. Holmberg, S., "Norrköpingsprojektet - en detaljerad pilotundersökning av hushållens vattenförbrukning" (in Swedish), *M81:5, Departement of Heating and Ventilation Technology*, Royal Institute of Technology, Stockholm, Sweden, 1991.
18. Incopera, F.P., De Witt, D.P., "*Fundametals of Heat and Mass transfer*", Third Edition, Wiley and Sons, ISBN 0-471-51729, 1990.
19. Wernstedt, F., "*Multi-Agent Systems for District Heating Management*", Licentiate thesis 2003:07, Blekinge Institute of Technology, Sweden, ISBN 91-7295-029-3, 2003.
20. Malmström, B., Ernfors, P., Nilsson, D., and Vallgren, H., "Shortterm forecasts of district heating load and outdoor temperature by use of on-line-connected computers" (in Swedish), *Värmeforsk, 1996:589*, ISSN 0282-3772, 1996.