



Distributed monitoring and control of office buildings by embedded agents

Paul Davidsson ^{a,*}, Magnus Boman ^b

^a *Department of Systems and Software Engineering, Blekinge Institute of
Technology, Soft Center, 372 25 Ronneby, Sweden*

^b *Swedish Institute of Computer Science (SICS), P.O. Box 1263, 164 29, Kista,
and the Royal Institute of Technology, Forum 100, 164 40 Kista, Sweden*

Accepted 27 September 2004

Abstract

We describe a decentralized system consisting of a collection of software agents that monitor and control an office building. It uses the existing power lines for communication between the agents and the electrical devices of the building, such as sensors and actuators for lights and heating. The objectives are both energy saving and increasing customer satisfaction through value added services. Results of qualitative simulations and quantitative analysis based on thermodynamical modeling of an office building and its staff using four different approaches for controlling the building indicate that significant energy savings can result from using the agent-based approach. The evaluation also shows that customer satisfaction can be increased in most situations. The approach here presented makes it possible to control the trade-off between energy saving and customer satisfaction (and actually increase both, in comparison with current approaches). © 2004 Elsevier Inc. All rights reserved.

Keywords: Software agents; Intelligent buildings; Distributed control

* Corresponding author.

E-mail addresses: paul.davidsson@bth.se (P. Davidsson), mab@sics.se (M. Boman).

1. Introduction

In a de-regulated market, the distribution utilities will compete with added value for the customer in addition to the delivery of energy. We will here describe a system consisting of a Multi-Agent System (MAS) that monitors and controls an office building in order to provide services of this kind. It is a further development of a system that was implemented as a part of the ISES (Information/Society/Energy/System) project [12,4]. A major improvement compared to the previous version is that the system now is decentralized and therefore scalable. The goal of the ISES project was to assess and demonstrate new business opportunities for future service-centric utilities.¹

The system uses the existing power lines for communication between the agents and the electrical devices of the building, i.e., sensors and actuators for lights, heating, ventilation, etc. The objectives are both energy saving and increasing customer satisfaction through value added services. Energy saving is realized, e.g., by lights being automatically switched off, and room temperature being lowered in empty rooms. Increased customer satisfaction is realized, e.g., by adapting temperature and light intensity according to each person's personal preferences.

In the MAS, which will be described in detail below, different agents control different parts of the building, as well as different aspects of the environmental conditions of the building. Other agents represent the persons in the building in order to maintain their preferences concerning temperature, light intensity, etc. The goal is to make the system transparent to the people in the building in the sense that they do not have to interact with the system in any laborious manner. To make it possible for the MAS to automatically detect in which room each person is at any moment and adapt the conditions in the room according to that person's preferences, we assume a Bluetooth-based indoor positioning system [5,14] (including Bluetooth-equipped PDAs and Bluetooth access points).

In order to evaluate the MAS approach to control environmental parameters such as temperature and light in office buildings, we have run a number of qualitative simulations as well as made quantitative calculations comparing two versions of the approach to the two currently most used methods for this type of control. In addition, fielded experiments at our test site—the Villa Wega building in Ronneby, Sweden—have been made to assure that the performance of power line communication is sufficient for controlling, e.g., radiators [12].

¹ The ISES project was a collaboration between a number of Swedish universities and some of the leading players in the European energy market, such as, EnerSearch AB (owned by Sydskraft and IBM Utility and Energy Services), ABB Network Partner AB, Electricité de France, and PreussenElektra.

1.1. Related work

The research efforts on intelligent buildings and environments have increased rapidly during the last couple of years. However, much work has been spent on either developing infrastructures supporting such applications or finding solutions to particular sub-problems, rather than on general control mechanisms on the system level. Also, most current work does not make use of the flexibility that agent technology offers and therefore we believe that its potential in this domain has not been sufficiently explored.

For instance, Hasha [6] describes a platform based on distributed active objects and has many characteristics in common with a normal multi-agent system platform. However, in the existing fielded implementation of this platform (the Gates Estate in Medina, Washington, USA), both hardware and installation costs were very high. Since it is based on a large number of computers (more than 120) connected via a dedicated network, rather than on (potentially cheap) smart sensors and actuators equipped with minimal processing capability and communicating with each other via the existing power lines, we fear that this approach will be too expensive to be widely used also in the future.

Another interesting piece of work is the Intelligent Room project at the MIT AI lab [2]. Its main focus is on the interaction between the users and the system, in particular on how to integrate different sensor modalities, such as vision, gestures, and speech. In contrast, our approach is to make this interaction as simple and transparent as possible for the users (i.e., by just using a PDA).

Colley et al. [3] have performed a series of interesting experiments in the area of intelligent buildings. Their work complements ours by focusing on issues that we have not studied in-depth, such as mobile robots, wearables, and learning.

One of the approaches most similar to ours is the ACHE system [9], which also aims at energy saving and increased personal comfort. While we have assumed that the persons working in the building enter their preferences manually, ACHE learns these automatically by observing the behavior of the persons of the building, e.g., when they manually adjust the settings of lighting or thermostats. An interesting idea would be to use this adaptability to learn, or at least fine-tune, the preference settings of our systems. However, ACHE does not have any system for locating and identifying individual persons and is thus unable to deal with personal preferences.

2. The building infrastructure

A typical office building contains an electrical network and a number of electrical devices that constitute an important part of its infrastructure. At

the Villa Wega test site, communication with the devices at the hardware level is facilitated by LonWorks technology (cf. <http://www.echelon.com>). Each device in the system is connected via hardware nodes to the LonWorks system, allowing the exchange of information over the electrical network.

Some of the devices are sensory and some are actuator devices. The sensory devices we use in the work presented here are temperature and light intensity. It is quite possible to include other types of sensors, such as presence sensors and fire detectors. The Bluetooth-based system makes it possible to know which persons are in each room at any moment. Note that even if establishing a connection via Bluetooth is too slow for many applications involving moving humans, its use here is less sensitive to such delays since positioning is typically used in Villa Wega for triggering agent behavior which is not time-critical, at least not on a second-basis. The other oft-cited drawback of the large number of base stations typically required for adequate indoor positioning via Bluetooth can with a similar argument be defended in our application. As we have not performed actual Bluetooth live tests in Villa Wega, we will not enter into details on the topic.

The actuator devices differ from the sensory devices in that it is possible, besides reading the state of the device, to change the state of the device (in order to change the state of the building). The actuator devices in the current application are lamps and radiators, but there are generic devices (such as LonPoint modules) that can be used to integrate any sensor or actuator with the system, for instance, a coffee machine, or a personal computer. The sensory devices provide input to the MAS and the actuator devices occasionally receive instructions from the MAS.

3. The multi-agent system

Each agent is intuitively linked to a particular entity in the building, such as an office, a meeting room, a corridor, a person, or an electrical device. The behavior of each agent is determined by a number of rules that express the desired control policies of the building conditions. The occurrence of certain events inside the building (like a person moving from one room to another) will generate messages to some of the agents that will trigger some appropriate rule(s). The agents execute the rule(s), with the purpose of adjusting the environmental conditions to some preferred set of values. The rule will cause a sequence of actions to be executed, which will involve communication between the agents of the system and eventually with an actuator device. The programming language used to implement the MAS was April [8] together with its extension April++. More information about the April Agent Platform (AAP) and the platform itself are available at <http://www.nar.fujitsu-labs.com/aap/>.

The agent-based approach provides an open architecture in the given context, i.e., agents can be easily configured and dynamically re-configured. It is possible to add new agents at run-time without the need of interrupting the normal operation of the system. Such changes reflect changes to the infrastructure of the building, or among the staff.

There are four main categories of agents in the MAS:

- *Personal Comfort (PC) agents*, which each corresponds to a particular person. It contains personal preferences and acts on that person's behalf in the MAS trying to maximize customer value. Thus, the agent does not model the behavior of a person, but tries to act in that person's interest. The PC agents are located in PDAs and communicate with other agents via Bluetooth.
- *Room agents*, which each corresponds to and controls a particular room with the goal of saving as much energy as possible. Taking into account the preferences of the persons currently in the room, it decides what values of the environmental parameters, e.g., temperature and light, are appropriate. Each room agent is embedded in a Bluetooth access point.
- *Environmental Parameter (EP) agents*, which each monitors and controls a particular environmental parameter in a particular room. They have access to sensor and actuator devices for reading and changing the parameter. For instance, a temperature agent can read the temperature sensor and control the radiators in a room. The goal of an EP agent is to achieve and then maintain the value of the parameter decided by the Room agent.

As illustrated in Fig. 1, the control of the actuators is performed locally which makes the system scalable. However, it is of course possible to monitor the state of the complete building at a central node.

To illustrate agent control, we describe what happens when a person moves from one room to another. The Bluetooth-based system continually informs the relevant Room agents (those residing at the access points that detect the PDA) on the location of PDAs in the neighborhood. When a person movement is detected by a room agent, the room agent of the new location informs the appropriate PC agent about this. The PC agent then provides the room agent with the personal preferences. The room agent decides, based on these preferences and on energy saving considerations, the new desired environmental conditions and pass them on to the EP agents. The EP agents then try to achieve and keep the values decided by the room agent by monitoring the relevant sensors and sending commands to the relevant actuators.

Using the approach described in [13], each agent contains a number of components contributing to the overall functionality of the agent. The architecture of the room and PC agents is depicted in Fig. 2. The EP agents have simpler structure and are implemented as April processes. Some of the components

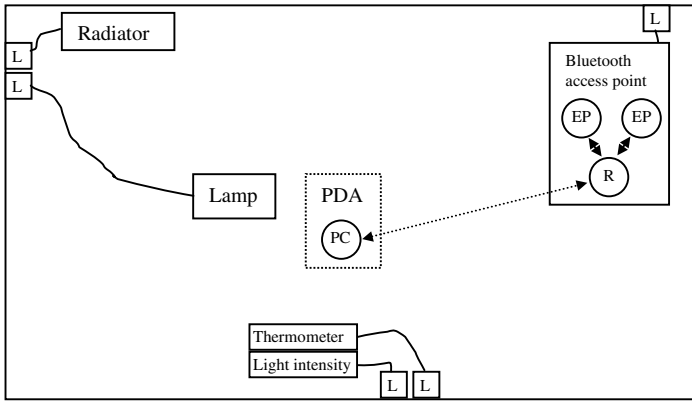


Fig. 1. An example hardware and software configuration of a room (including one person, i.e., one PDA). All stationary equipment is connected to electrical outlets via LonWorks adapters (*L*). The circular objects correspond to agents (*R* = room agent, *EP* = environmental parameter agent, and *PC* = personal comfort agent) and the arrows to communication between agents.

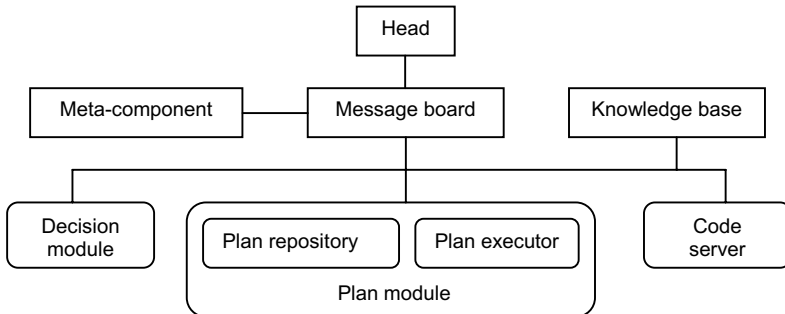


Fig. 2. The architecture of the room and PC agents.

(rectangular boxes) are included in each agent by default, but it is possible to add other domain-dependent components (rounded boxes).

One generic module is the *head* that plays the role of the communication interface of the agent. All messages directed to the agent are sent to the head and are subsequently forwarded to the internal modules. In this way external entities do not need direct access to the agent’s modules. Also, a shared *knowledge base* is included that can be used to store shared information. The *meta-component* is used for administrative purposes during the addition and deletion of components. Finally, the communication between the components is facilitated by the *message board*.

Depending on the situation, an agent needs to execute a sequence of actions, i.e., a plan. The *plan module* is responsible for maintaining such plans and con-

sists of a *plan repository* and a *plan executor*. The plan repository stores the plan descriptions, which include the name of the actions involved in the plan, their temporal relationships, and descriptions of the information they manipulate. Requests for executing single plans are received by the plan executor. The executor fetches the code that implements the specific actions from the *code server* and executes them. The code stored on a code server can be supplied on demand. An alternative approach would be to have the actions hard-wired in the decision module. Although our approach requires some additional communication among the components in order to retrieve the code, it simplifies the configuration of the agents. New actions are simply added to the code server and new plans can be added to the plan repository. Exactly which plan to be executed at a particular moment is decided by the *decision module* based on the agent's current state and any external event.

The MAS conforms to a number of general rules (constraints) that are programmed into the agents. Some examples are listed below:

- When a particular person is in her office, the room agent must adapt temperature, light, etc. to her preferences, otherwise the default conditions are maintained. If an irrelevant person (i.e., another person than the one that normally works in that office) enters, this does not affect the environmental conditions (except for the light being turned on if the room was empty).
- For meeting rooms, the temperature condition is adjusted to the mean value of all the meeting participants, and the light intensity to the highest preference value.
- For other common rooms, like corridors, the temperature remains steady regardless whether there are people in the room or not. The light is turned on only when at least one person is in the room, otherwise it is turned off.
- Every room with no persons in it must maintain some default environmental conditions.
- It must always be possible to over-rule the decisions of the agents in the MAS by physical interaction with the electrical equipment. For instance, even if an EP agent has decided that the light in a room should be on, it must be possible for a person to turn off the light using the switch in the actual room.

These constraints are not hard-wired into the MAS and can be changed easily. With regard to the last constraint, it is worth mentioning that the concept of manual overrides is becoming increasingly important to systems in which human and artificial agents both act [10]. The agents in the MAS must display *adjustable autonomy*. Interestingly, human operation of the hardware in Villa Wega is viewed as a form of interference by the agents controlling the building. At the same time, a difficult object from the design point of view is to make the agent operations transparent to the people in the building, and in doing so prevent persons from viewing agent action on hardware as a form of interference.

Usually, the goals of the room agents and the PC agents are conflicting: the room agents maximizing energy saving and the PC agents maximizing customer value. Another type of a conflicting goal situation is the adjustment of temperature in a meeting room in which people with different preferences regarding temperature will meet.

3.1. Pronouncers

We have made some initial investigations into introducing real-time decision support to the agents in the form of *pronouncers* [1]. A pronouncer can be seen as a decision module located outside the agents themselves, providing normative advice in generic situations, and typically using only information provided by the agents themselves at query time, e.g., in the form of decision trees with probability and utility assessments. We are currently investigating the combination of pronouncers and technical norms. This combination in turn allows for agents to abandon elaborate plans and increases efficiency by freeing agents from the burden of plan revisions (cf. [16]). The intelligent building domain is suitable for pronouncer use, since the size of the agents must be kept reasonably small.

4. Evaluation

Although much of the hardware necessary to evaluate the approach outlined above is actually installed in the Villa Wega building, it would be quite expensive to make the installation complete. Since communication over the electrical network is a new technology and devices currently are produced in small numbers, the required hardware is expensive at the moment, but we expect that this situation will change drastically in the next couple of years. Therefore, we have made the evaluation of the approach (i.e., the MAS) through qualitative simulations [1,4] as well as through quantitative analytical computations.

The total system can be divided into three parts; the hardware, i.e., the building including sensors and effectors, the software, i.e., the MAS, and finally the people working in the building. Thus, we simulate the hardware and the behavior of the people, and let the actual MAS, which would be used in a fielded application, interact with these simulated entities instead of the actual building and people. This simulation of the behavior of the people should be contrasted to the PC agents in the MAS, which serve the persons, and so are agents in the true sense of the word.

In order to monitor the simulations, a graphical user interface (GUI) visualizing the building environment was implemented. Fig. 3 shows a snapshot of the environment visualization GUI that visualizes the state of the building in terms of temperature, light intensity of the rooms, and the persons present in the rooms. In order to verify the behavior of the MAS, a number of scenarios were simulated using this software.

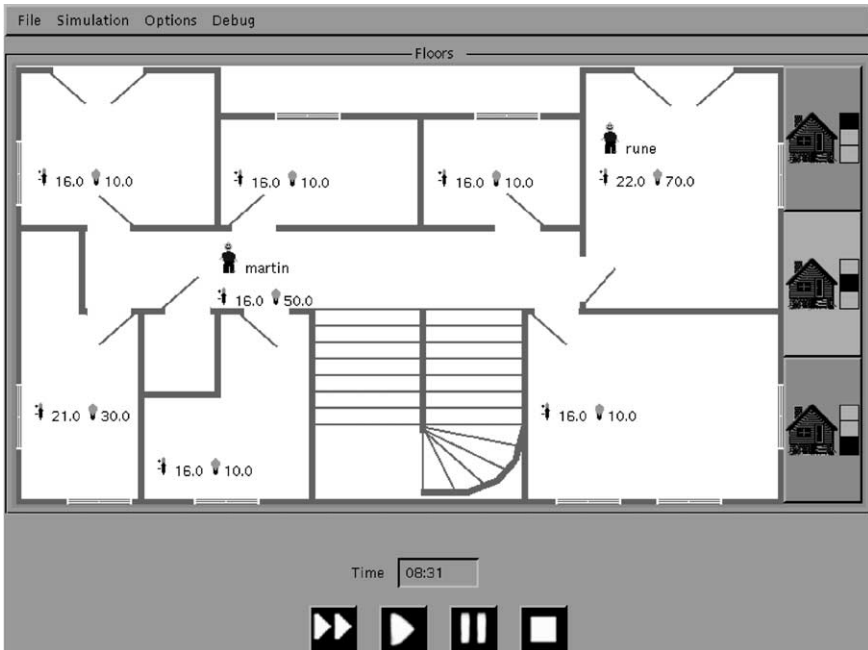


Fig. 3. A snapshot of the environment visualization GUI.

The physical properties of the building were modeled using the thermodynamical models described by Incropera and Witt [7]. These were discretized according to standard procedures (cf. Ogata [11]).

All the thermodynamical characteristics of a room are described by two constants: the thermal resistance, R , which captures the heat losses to the environment, and the thermal capacitance, C , which captures the inertia when heating up/cooling down the entities in the room. (In the quantitative evaluation below we use the sample time 1 min.) The temperature, T_{xi} , in room x at time i is described by

$$T_{xi} = \frac{1}{1 + \frac{1}{R_x C_x}} \left(T_{x(i-1)} + \frac{P_i + T_{out}}{R_x} \right)$$

where P_i is the heating power, T_{out} the outdoor temperature, and $T_{x(i-1)}$ is the temperature in room x 1 min ago.

4.1. The quantitative evaluation

A number of simplifications were made:

- only energy used for heating is taken into account, not for lighting or other purposes,

- constant outdoor temperature is assumed (10 °C),
- radiation from the sun is negligible (i.e., weather cloudy),
- the heat produced by persons in the room is ignored,
- the heat produced by computers, lamps, and fluorescent tubes is ignored.

Note that if we were to take into account any of these aspects, the performance of the MAS approach would probably have been even more favorable compared to the other approaches. For instance, since the MAS approach would take into account and make use of the outdoor sunlight, both energy saving and customer satisfaction would increase if we were to control also the lighting.

The building has five small offices (each used by one person), two large offices (3–5 persons), and one meeting room, and one corridor at each of the three floors. We use $R = 0.1$ and $C = 3000$ for the small offices in the building, $R = 0.05$ and $C = 5000$ for the large offices, and $R = 0.05$ and $C = 3000$ for the meeting room. (Larger rooms have greater losses to the environment than smaller rooms and there are fewer entities to heat up/cool down in the meeting room.) In the small offices there is one 1000 W radiator, whereas in the large offices and the meeting room there are two such radiators.

In the scenario used in the calculations there are 12 persons working in the building who share the following characteristics:

- prefer 22 °C both at their offices and when in the meeting room,
- the working day is normally nine hours with a one hour lunch break, i.e., on average eight hours are spent in the building. However, there is a 20 per cent likelihood that a person does not show up at all during a day (because of meetings in another city, illness, etc.),
- on average there are five meetings in the meeting room each week,
- the length of a meeting is two hours on average.

We assumed that the radiators use a simple (ideal) temperature control algorithm: to raise the temperature, they use the maximal effect (i.e., 1000 W) to heat up the room to the desired temperature. To maintain the desired temperature, they produce just the right amount of heating power. Finally, to lower the temperature, the radiators are turned off.

4.2. The results—energy saving

Four different approaches were compared:

1. *The thermostat approach*: This is the current method of controlling the environmental parameters of the Villa Wega building (and most other buildings in the industrialized world). The people working in the building set the

desired temperature manually. However, since most people do not lower the temperature in their offices when they go home, we assume that the temperature is always set to 22 °C both in the offices and in the meeting room.

2. *The timer-based approach*: This is a bit more sophisticated (in fact, it may well be the smartest approach in current use). A timer starts raising the temperature at 7 a.m. to 22 °C in all rooms, and at 7 p.m. it starts to lower the temperature to 16 °C, i.e., the thermostat is set to 22 °C and 16 °C respectively.
3. *The reactive MAS approach*: When a person is in the building, the temperature of her office is set to 22 °C, and when she is not, the temperature is set to 16 °C. Similarly, when the meeting room is empty the temperature is set to 16 °C, and otherwise to 22 °C.
4. *The pro-active MAS approach*: makes use of the electronic diaries of the persons working in the building in order to heat up the rooms to the preferred temperature in advance. (Thus, it requires that the individuals keep their electronic diaries on their PDAs updated.)

The results are described in Table 1.

Thus, compared to first approach, we save almost 40% energy by using the MAS approach and almost 12% compared to the timer-based approach. Note also that the pro-active approach is only slightly more energy consuming than the reactive, but will increase customer temperature satisfaction (see Section 4.3).

4.3. The results—customer satisfaction

The saving of energy was only one goal of our system. Now we turn to the evaluation of how the MAS meets the second goal of increased customer satisfaction. We will here concentrate on the persons who work in the small offices. We use a simple linear model of the degree of satisfaction with respect to temperature where 16 °C corresponds to 0% satisfaction and 22 °C corresponds to 100% satisfaction.

Table 1
The average weekly energy consumption of the four control approaches

Control approach	Average weekly energy consumption
1. Thermostat	221.8 kW h
2. Timer-based	154.3 kW h
3. Reactive MAS	136.2 kW h
4. Pro-active MAS	137.0 kW h




Fig. 4. The distribution of working time of the persons involved (the height of a bar corresponds to the probability that the person is working in the building during that hour).

In order to make an appropriate comparison we have to specify the distribution of working time of the persons involved. We have assumed the distribution illustrated in Fig. 4 on weekdays, i.e., no work at all during weekends.

The results are described in Table 2.

Thus, we see that the thermostat approach yields the maximal degree of customer satisfaction since it keeps the desired temperature at all times. However, as we have seen, the price for this is a very high energy consumption. The current method to lower the energy consumption, i.e., using a timer-based approach, on the other hand, has a significantly lower degree of customer satisfaction than the MAS-based approaches. In addition, the MAS-based approaches enable us to control the trade-off between energy saving and customer satisfaction in a much more sophisticated manner than, e.g., the timer-based approach. Notice also that the assumed distribution of working time is quite favorable to the timer-based approach. For instance, if we were to include over-time work during weekends the results would be much worse while the performance of the MAS-based approaches would be the same as before.

Admittedly, this evaluation of customer satisfaction is very coarse, ignoring many aspects that would influence the degree of satisfaction experienced by the people working in a building actually equipped with such a system, e.g., personal integrity issues, and the pressure of keeping the diary updated. Privacy and integrity are very interesting aspects, and believe that empirical studies based on field-tests are necessary to assess these aspects properly. That said,

Table 2

The average degree of temperature satisfaction of the four control approaches

Control approach	Average degree of temperature satisfaction (%)
1. Thermostat	100.0
2. Timer-based	91.8
3. Reactive MAS	97.7
4. Pro-active MAS	100.0

the very identification of a MAS being able to act as an adjusting system between conflicting goals is a result in its own right.

Also, developing more complex “contracts” with the customers would be a possibility for the providers. For instance, the contract could be stated in the following way: the average “satisfiability factor” (a qualitative measure of subjective utility) should be 0.95 and/or it should not drop below 0.5 for more than 10 min. Such contracts would open up further possibilities to save energy, or performing intelligent load balancing by agent-controlled Demand Side Management (DSM).

5. Conclusions

We have given a high-level description of a project aimed at investigating the usefulness of the multi-agent systems approach for the design of control systems for intelligent buildings. The use of the agent approach was initially motivated by the close mapping that it offered between the entities of the application domain and the entities of the software. The concurrent non-deterministic nature of the activities inside the building was another factor that led to the development of concurrent autonomous entities. In fact, the intelligent office building domain matches all the characteristics of the domains for which agent-based systems has been found useful. For instance, Parunak [15] argues that an agent approach is appropriate for applications that are modular, decentralized, changeable, ill-structured, and complex.

We have presented a general multi-agent system architecture, which we argue can be easily adapted to almost any building. Moreover, the agent system was designed to allow for dynamic re-configuration of the agents, without any disruptions of the operation of the system. This is a useful feature when changes in the building infrastructure or of the persons in the building occur. Finally, we evaluated the approach by means of qualitative computer simulations and quantitative analyses based on thermodynamical models. Our results indicate that the approach is viable and that considerable energy savings are possible while at the same time providing added value for the customer. In addition, the approach enables a much more fine-grained control of the trade-off between energy saving and customer satisfaction than is possible with current approaches.

It is also worth mentioning that an agent-based approach permits even more advanced control mechanisms than previously mentioned in this paper. For instance, it is possible to let the agents take into account that the price of energy is not constant. As an example, we are considering to integrate the system with another agent-based system, used to control a district heating system using demand side management [17]. We have also been experimenting with more complex functionality, e.g., when a person enters the building in the morning,

her monitor is switched on and the coffee machine starts making coffee. While the study of such functionalities is beyond the scope of the work presented here, we believe that they will be very important when developing future intelligent buildings.

Acknowledgments

We thank Marko Krejic for helping us with hardware-related issues, Dr. Nikolaos Skarmas and Professor Keith L. Clark who helped us with the implementation of the MAS, and Dr. Fredrik Ygge for help with the thermodynamical modeling.

References

- [1] M. Boman, P. Davidsson, H.L. Younes, Artificial decision making under uncertainty in intelligent buildings, in: 15th Conference on Uncertainty in Artificial Intelligence, Morgan Kaufmann, 1999, pp. 65–70.
- [2] M. Coen, Design principles for intelligent environments, in: 15th National Conference on Artificial Intelligence, 1998, pp. 547–554.
- [3] M. Colley, G. Clarke, H. Hagrass, V. Callaghan, A. Pounds-Cornish, Intelligent inhabited environments: cooperative robotics and buildings, in: 32nd International Symposium on Robotics, Seoul, Korea, 2001.
- [4] P. Davidsson, M. Boman, Saving energy and providing value added services in intelligent buildings: a MAS approach, in: Agent Systems, Mobile Agents, and Applications, LNCS, vol. 1882, Springer-Verlag, 2000.
- [5] S. Feldmann, K. Kyamakya, A. Zapater, Z. Lue, An indoor Bluetooth-based positioning system: concept, implementation and experimental evaluation, in: International Conference on Wireless Networks, 2003.
- [6] R. Hasha, Needed: a common distributed-object platform, IEEE Intelligent Systems (March–April) (1999) 14–16.
- [7] F.P. Incropera, D.P. Witt, Fundamentals of Heat and Mass Transfer, third ed., Wiley and Sons, 1990.
- [8] F.G. McCabe, K.L. Clark, April: Agent process interaction language, in: M.J. Wooldridge, N.R. Jennings (Eds.), Intelligent Agents, LNCS, vol. 890, Springer-Verlag, 1995, pp. 324–340.
- [9] M. Mozer, The neural network house: an environment that adapts to its inhabitants, in: AAAI Spring Symposium on Intelligent Environments, 1998, pp. 110–114.
- [10] D. Musliner, B. Pell (Eds.), Agents with adjustable autonomy, AAAI Spring Symposium, Technical Report SS-99-06, AAAI, 1999. ISBN 1-57735-102-9.
- [11] K. Ogata, Modern Control Engineering, second ed., Prentice-Hall, 1990.
- [12] H. Ottosson, H. Akkermans, F. Ygge (Eds.), The ISES Project, EnerSearch, 1998, 91-9753567-0-0.
- [13] N. Skarmas, Agents as objects with knowledge base state, Ph.D. Thesis, Imperial College, Department of Computing, January 1997.
- [14] K. Thapa, S. Case, An indoor positioning service for Bluetooth Ad Hoc networks, in: MICS 2003, Duluth, MN, USA.
- [15] H. Van Dyke Parunak, Industrial and practical applications of DAI, in: G. Weiss (Ed.), Multiagent Systems, MIT Press, 1999.

- [16] H. Verhagen, M. Boman, Norms can replace plans, in: IJCAI'99 Workshop on Adjustable, Autonomous Systems, 1999.
- [17] F. Wernstedt, P. Davidsson, A multi-agent system approach to monitoring and control of district heating systems, in: Developments in Applied Artificial Intelligence, LNCS, vol. 2358, Springer-Verlag, 2002.