Making a Case for Multi-Agent Systems

by

Fredrik Ygge and Hans Akkermans

Department of
Computer Science and Business Administration
University of Karlskrona/Ronneby
S-372 25 Ronneby
Sweden

ISSN 1103-1581
ISRN HKR-RES—97/4—SE

Making a Case for Multi-Agent Systems

by Fredrik Ygge and Hans Akkermans

ISSN 1103-1581

ISRN HKR-RES—97/4—SE

Copyright © 1997 by Fredrik Ygge and Hans Akkermans

All rights reserved

Printed by Psilander Grafiska, Karlskrona 1997

Making a Case for Multi-Agent Systems

Fredrik Ygge[‡]

Hans Akkermans[§]

[‡]EnerSearch AB and University of Karlskrona/Ronneby

Department of Computer Science (IDE)

372 25 Ronneby, Sweden

Fredrik.Ygge@ide.hk-r.se, <http://www.rby.hk-r.se/~fredriky>

[§]AKMC and University of Twente

Department of Information Systems (INF/IS)

P.O. Box 217, NL-7500 AE Enschede, The Netherlands

akkermans@ecn.nl, J.M.Akkermans@cs.utwente.nl

Abstract

Multi-Agent Systems (MAS) promise to offer solutions to problems where established, older paradigms fall short. To be able to keep promises, however, in-depth studies of advantages and weaknesses of MAS solutions versus conventional ones in practical applications are needed. In this paper we offer one such study. Climate control in large buildings is one application area where MAS, and market-oriented programming in particular, have been reported to be very successful. We have therefore constructed and implemented a variety of market designs for this problem, as well as different standard control engineering solutions. This paper gives a detailed analysis and comparison, so as to learn about differences between standard versus MAS approaches, and yielding new insights about benefits and limitations of computational markets.

1 Introduction

When new paradigms arise on the scientific horizon, they must prove their value in comparison and competition with existing, more established ones. The multi-agent (MAS) paradigm is no exception. In a recent book on software agents, edited by Jeffrey Bradshaw [Bradshaw, 1997], Norman observes that perhaps “the most relevant predecessors to today’s intelligent agents are servomechanisms and other control devices”. And in-

deed, a number of applications for which MAS have recently claimed success, are close to the realm of what is traditionally called control engineering. One clear example is the climate control of large buildings with many office rooms. Here, Huberman *et al.* [Huberman and Clearwater, 1995] have constructed a working MAS solution based on a market approach, which has been reported to outperform existing conventional control.

A key question is: in what respect and to what extent are multi-agent solutions better than their alternatives? We believe that the above-mentioned application provides a nice opportunity to study this question in more detail. It is practically very relevant, it lends itself to alternative solutions, and it is quite prototypical for a wide range of industrial applications in resource allocation (including the energy management applications we are working on [Ygge and Akkermans, 1996], the file allocation problem of Kurose and Simha [Kurose and Simha, 1989], and the flow problems investigated by Wellman [Wellman, 1993]).

The contribution of this paper is a detailed analysis of a published MAS approach to building climate control and a comparison between this approach and traditional approaches. We also introduce a novel approach to this problem based on a general equilibrium market. From the analysis of these approaches we draw conclusions not only about the suitability of these approaches, but of many other MAS approaches as well.

The structure of the paper is as follows. Section 2 introduces the application: it describes the office environment and gives the physical model for cooling power and the various temperature and outside weather influences. First, we give a standard control engineering solution, based on local and independent integral controllers (Section 3). Next, we consider the market-based approach as put forward by Huberman *et al.* [Huberman and Clearwater, 1995] (Section 4), and we discuss a number of variations on this market design (Section 5), providing a kind of factor analysis for its success. Section 6 then shows the results of an improved control engineering scheme that exploits global data. Finally, we propose a market design of our own, based on local data only, and make a comparison across the different approaches (Section 7). Section 8 puts our results into perspective and summarizes the conclusions.

2 The Office Environment

In this section, we present the mathematical-physical model of the office environment. For those readers who find the mathematics forbidding, we first give a conceptual summary which makes it possible to skip the equations. The offices are attached to a pipe in which the resource (cold air) is transported as in Figure 1. The characteristics of this system are similar to the characteristics of a district heating system, but with offices instead of households. We assume that there are 100 offices in total, and that they are equally distributed towards East, South, West and North.

Figure 1: *Offices and air transportation.*

The thermodynamics of the office environment is actually quite simple. Every office is seen as a storage place for heat, but heat may dissipate to its environment. In the model, the thermodynamic behaviour of an office is equivalent to a simple electrical RC-circuit. Here, voltage is analogous to temperature, and electrical current is analogous to heat flow. C and R then respectively denote heat capacitance and heat losses to the environment.

A good general reference on thermodynamic models as we use here is [Incropera and Witt, 1990]. The heat equations are continuous in time, but are discretized according to standard procedures used in control engineering, described *e.g.*, in [Ogata, 1990]. The AI aspects of thermodynamic model building with component libraries are discussed extensively in [Borst *et al.*, 1997].

2.1 Basic Physical Properties

The resource treated is cooling power. Each office can only make use of a fraction, η , of the available resource at that office, P_{aio} , i.e.

$$P_{cio} \leq \eta \cdot P_{aio}, \quad (1)$$

where P_{cio} is the consumed power. The available resource at one office is equal to the available resource at the previous office minus the consumed resource at the previous office. Throughout the paper we use an η of 0.5. The index i denotes the time interval under observation and the index o denotes the office under observation. This naming convention is used throughout the entire paper. In addition, the index k is also sometimes used for time intervals.

We treat everything in discrete time with a time interval of one minute. For all offices the

temperature, T_{io} , is described by

$$T_{io} = T_{0,o} + \sum_{k=1}^i (P_{hko} - P_{cko})/C_o, \quad (2)$$

where P_{hko} is the heating power and C_o is the thermal capacitance. The heating power is described by

$$P_{hio} = (T_{vio} - T_{io})/R_o, \quad (3)$$

where R_o is the thermal resistance and T_{vio} is a virtual outdoor temperature detailed below.

From Eq. (2) and Eq. (3) we obtain

$$T_{io} = \frac{1}{1 + \frac{1}{R_o C_o}} (T_{i-1,o} + \frac{T_{vio} - P_{cio}}{C_o}), \quad i > 0. \quad (4)$$

2.2 Weather Model

All external weather influences on the office environment are modelled by a virtual temperature, representing outdoor temperature, sun radiation etc. We assume that there is sunshine every day and that the outdoor temperature, T^{outd} , varies from 22 to 35°C according to

$$T_i^{outd} = 22 + 13 \cdot e^{-((i \cdot s - 4) \bmod 24 - 12)^2 / 20}, \quad (5)$$

where s is the length of each time interval expressed in hours, i.e. here $s = 1/60$.

The virtual temperature, T_{vio} , is described by

$$T_{vio} = T_i^{outd} + T_{ro} + T_{dio}, \quad (6)$$

where T_d is a random disturbance, thought to represent small fluctuations caused by e.g. the wind. T_d is Gaussian distributed with an average value of 0 and a standard deviation of 1. T_r is a sun radiation component. For the offices located at the East side T_r is described by

$$T_{ri}^{East} = 8 \cdot e^{-((i \cdot s + 4) \bmod 24 - 12)^2 / 5}. \quad (7)$$

and correspondingly for the South and the West offices

$$T_{ri}^{South} = 15 \cdot e^{-(i \cdot s \bmod 24 - 12)^2 / 5} \quad (8)$$

and

$$T_{ri}^{West} = 8 \cdot e^{-((i \cdot s - 4) \bmod 24 - 12)^2 / 5}. \quad (9)$$

The various temperatures are plotted in Figure 2.

Figure 2: The leftmost plot shows the outdoor temperature, T_i^{outd} . The middle plot shows the sun radiation components, T_r , (with peaks at time 8, 12, and 16 hours for offices located at the East, South, and West side respectively). Finally, the outdoor temperature plus the sun radiation components are plotted to the right.

2.3 Office Temperatures Without Control

In Figure 3, the temperature for a South oriented office is plotted with different thermal resistances, R_o , and thermal capacitances, C_o . For simplicity we assume all R_o to be equal and all C_o to be equal. From this figure we note two things: first, the higher $R_o C_o$ the bigger the lag in temperature changes, and second, the higher $R_o C_o$ the smaller the fluctuations in temperature. For the experiments in this paper we chose $R_o = 10$ and $C_o = 10$. An R_o or a C_o equal to zero implies that $T_{io} = T_{vio}$, as can be seen by letting R_o or C_o approach zero in Eq. (4).

Figure 3: The indoor temperature for an uncontrolled office is plotted for different values of the thermal resistance and the thermal capacitance. Small values of the thermal resistance and capacitance give higher peaks, while higher values give smoother curves.

2.4 Simulations

All solutions discussed in this paper have been implemented in and simulated by the authors using C++ on an IBM compatible PC running Windows95. Furthermore, for guaranteeing that a fair comparison on the same footing was carried out, all simulations have been independently recreated and verified from the paper, as part of a master's thesis project, using Python on an IBM compatible PC running Linux.

3 CONTROL-A: Conventional Independent Controllers

3.1 Integral Control for the Offices

The application of regulating the office temperature has a long tradition in control theory. The most widely known controllers are different variants of the PID controller. The letters PID denote that the control signal is proportional to the error (i.e. the difference between the setpoint and the actual value) — the P; proportional to the integral of the error — the I; and proportional to the derivative of the error — the D. Here, we will use a variant of the integrating controller, i.e., an I-controller, of the form

$$F_{io} = F_{i-1,o} + \beta(T_{io} - T_o^{setp}), \quad (10)$$

where F is the output signal from the controller and β is a gain parameter. For the simulations it is assumed that F_{io} will not be below 0 and that it will not exceed 3. The control signal F_{io} is sent to the actuator and the actual P_{cio} is obtained from

$$P_{cio} = \begin{cases} F_{io}, & F_{io} \leq \eta \cdot P_{aio} \\ \eta \cdot P_{aio}, & F_{io} > \eta \cdot P_{aio} \end{cases} \quad (11)$$

Plots of the office temperatures with different gains are shown in Figure 4. The gain of the controller is not critical for this application. Too high a gain will result in the controller overreacting when the temperature exceeds the setpoint, after which it will be under the setpoint for quite some time. This leads to a larger error than if smaller adjustments are made. Also, the amplitude of the control signal then gets unnecessarily high, but the system does not get dangerously unstable. We note that the maximum deviation here is $\pm 0.06^\circ\text{C}$. Thus, controllers using any of the three gains perform very well. Apart from the plots shown here, also a step response has been analyzed. From a brief analysis we chose a gain of 10.

The Performance Measure We use the same evaluation of the building control as Huberman *et al.* [Huberman and Clearwater, 1995], namely the standard deviation of the de-

Figure 4: The indoor temperature for an office, utilizing a integral controller, is plotted for different controller gains.

viation from the setpoint,

$$StdDev(T - T^{setp}) = \sqrt{\frac{1}{N} \sum_{o=1}^N [(T_{io} - T_o^{setp}) - (\langle T_i \rangle - \langle T^{setp} \rangle)]^2} \quad (12)$$

where T^{setp} is the setpoint temperature, $\langle T \rangle$ is the temperature averaged over all offices, and $\langle T^{setp} \rangle$ is the average setpoint temperature.

Figure 5: The standard deviation, as defined in Eq. (12) for each interval, i , is plotted for different amounts of available resource (130, 140, 150, and 160) with 100 offices. The lower the amount of available resource the higher the standard deviation.

Limitations on Available Resource So far we have assumed that the total amount of available resource is unlimited. Now we assume that there is a maximum value for the power that is inserted into the system. In such a situation offices that are situated close to

the air input will have a sufficient amount of cool air, while those near the end will suffer if totally uncoordinated controllers are used. Thus, the smaller the amount of the total available resource, the larger the standard deviation will be. This is visualized in Figure 5. From the figure we have chosen an upper limit for the total amount of resource of 140.

Discussion As seen from the example, independent integrating controllers perform very well when the resource is unlimited. On the other hand, when there is a shortage of resource the standard deviation increases dramatically.

4 MARKET-A: The Approach by Huberman *et al.*

A MAS solution to the problem of building control has been presented by Huberman *et al.* [Huberman and Clearwater, 1995, Clearwater and Huberman, 1994]. The approach taken is to model the resource allocation problem of the building environment as a computational market where agents buy and sell resource. The non-separability in terms of agents, i.e. the fact that the amount of resource the penultimate agent can use depends on how much resource the last agent receives, is ignored.

The basic idea is that every office is represented by an agent that is responsible for making good use of the resource and to this end trades resources with other agents. The agents create bids which are sent to an auctioneer that calculates a clearing price, and makes sure that no agent has to buy for a price higher than its bid nor has to sell for a price lower than its bid.

Section 4.1 describes the mathematical details of this approach. The reader having no special interest in these details may skip this part and take a look at the results of the corresponding simulations in section 4.2.

4.1 Original Formulation

All formulas in this section were taken from [Huberman and Clearwater, 1995] and [Clearwater and Huberman, 1994].

Trade Volumes First, the decision for an agent to buy or sell is based on

$$t_{io} = \frac{T_o^{setp}}{T_{io}} \cdot \frac{\langle T_i \rangle}{\langle T^{setp} \rangle} \begin{cases} t_{io} > 1, & seller \\ t_{io} < 1, & buyer \end{cases} \quad (13)$$

Then the total trade volume, V , is calculated from

$$V_i = \sum_{o=1}^N |1 - t_{io}|, \quad (14)$$

where N is the number of offices.

Then every agent calculates its request volume, v , according to

$$v_{io} = \alpha \frac{|1 - t_{io}|}{V_i}. \quad (15)$$

When an agent buys or sells its v the actual movement of a valve, called VAV , is computed from

$$\Delta VAV_{io} = f(flow_{io}, v_{io}, VAV_{io}), \quad (16)$$

where f is an empirically determined function which is not explicitly given in the papers [Huberman and Clearwater, 1995, Clearwater and Huberman, 1994].

From this the actual VAV position for each interval is updated according to

$$VAV_{i+1,o} = VAV_{io} + \Delta VAV_{io}. \quad (17)$$

Bids The bids are based on the *marginal utility*¹ of the form described by

$$U(t_{io}/T_o^{setp}, m_{io}) = [U(0, m_{io})]^{(1-t_{io}/T_o^{setp})} = [U(0, m_{io})]^{(1-\frac{\langle T_i \rangle}{T_{io}\langle T^{setp} \rangle})}, \quad (18)$$

with

$$U(0, m_{io}) = u_3 - (u_3 - u_1)e^{-\gamma m_{io}}, \quad (19)$$

and

$$\gamma = \ln \left[\frac{u_3 - u_1}{u_3 - u_2} \right], \quad (20)$$

where $u_1 = 20$, $u_2 = 200$, and $u_3 = 2000$, and m is the amount of money that an agent has, as given by

$$m_{io} = 100(2 - VAV_{io}). \quad (21)$$

By observing Eq. (19) and Eq. (20), we note that the equations can be simplified to

$$U(0, m_{io}) = u_3 - (u_3 - u_1)e^{-\gamma m_{io}} = u_3 - (u_3 - u_1)(e^\gamma)^{-m_{io}} = u_3 - (u_3 - u_1) \left(\frac{u_3 - u_1}{u_3 - u_2} \right)^{-m_{io}}. \quad (22)$$

Since the VAV varies between 0 and 1, the amount of money an agent has varies between 100 and 200, and thus $U(0, m)$ varies between 1999.85632 and 1999.99999. Hence, the

¹What Huberman *et al.* call utility in their work [Huberman and Clearwater, 1995, Clearwater and Huberman, 1994] is generally called marginal utility in micro-economic theory.

notion of gold does not effect the marginal utility in practice, as will be verified by the simulations in Section 5.

The bids are calculated from multiplying the marginal utility with the previous price, $price$, according to

$$B_{i+1,o} = U_{io}(t_{io}/T_o^{setp}, m_{io}) \times price_i. \quad (23)$$

Auction All bids are sent to an auctioneer which calculates a market price where supply meets demand. All agents which requested to buy and which bids are above or equal to the market price are assigned the corresponding amount, and, similarly, all agents which requested to sell and which bids are below or equal to the market price are assigned the corresponding of amount of resource.

Assumptions and Simplifications Since f in Eq. (16) above is not explicitly given and since the relation between the VAV position and P_c is specific to the office environment, we need to make an assumption about how the bid volumes relate to P_c . We take the simplest choice and let Eq. (16) and Eq. (17) be replaced by

$$F_{i+1,o} = F_{io} \pm v_{io}, \quad (24)$$

where plus or minus depends on whether it was an accepted buy or sell bid. P_{cio} is obtained from Eq. (11). This simplification is not crucial to our central line of reasoning.

As pointed out by Clearwater & Huberman [Clearwater and Huberman, 1994], the purpose of the auction is only to reallocate between the agents and not to affect the average temperature. This means that even if the valves are not opened very much, there is plenty of resource in the system, and if the offices are all above the setpoint, no valve can open further without another one closing. In order not to complicate the model further, the simulation is only performed in a time interval where there is a total lack of resource, with the total available resource assigned to the system. Thus, we do not implement a mechanism that inserts and deletes the total resource to the system. We choose to simulate between 3 p.m. and 7 p.m.

We assume the amount of money to be given by

$$m_{io} = 100(2 - \frac{F_o^{max} - F_{io}}{F_o^{max}}). \quad (25)$$

Thus, the amount of money will vary from 100 to 200, as in the original work.

Eq. (23) turns out to produce major problems in simulations. Above we have seen that $U(0, m) \approx 2000$. Equation (18) shows that $U(t_{io}/T_o^{setp}, m_{io})$ is minimized for minimized T_{io} and maximized $\frac{\langle T_i \rangle}{\langle T^{setp} \rangle}$. As we can expect T_{io} to be well above $10^\circ C$ and $\frac{\langle T_i \rangle}{\langle T^{setp} \rangle}$ to be well below 2, we can be sure that $U(t_{io}/T_o^{setp}, m_{io})$ will be well above $2000^{1-\frac{2}{10}} \approx 437$.

Thus, the bidding price will never be below 437. Then, Eq. (23) tells us that the market price will be at least $price_0 \times 437^i$. This leads to numerical overflow after only a few iterations. In practice the market price scales with approximately $price_0 \times 1300^i$ which of course is even worse. We note however that, since all agents multiply their bids with the previous price, this has no effect on the reallocation itself, but only affects the price level. Therefore, we choose to omit multiplying by the previous price in our simulations, noting that this leads to exactly the same allocations but avoids numerical overflow. Hence, the bids are based on Eq. (18) rather than Eq. (23).

The price at each auction is adjusted until supply meets demand. But, since the bids are given using discrete volumes, supply will seldom match demand *exactly* at the clearing price. Normally, there will be a very small excess demand or supply. If there is an excess supply, all buyers that are willing to pay at most the clearing price will buy, but not all sellers that are willing to accept at least the clearing price can sell. In this situation, the sellers are selected randomly from the valid candidates. The same procedure is used when there is a small excess demand.

4.2 Simulations

Figure 6 shows two plots. The upper plot is the standard deviation when independent integrating controllers are used and the lower one shows the agent-based control scheme as defined above. We found that an α , as described by Eq. (15), of 64 led to the smallest overall standard deviation. It is seen that the agent approach offers at least one order of magnitude of improvement.

Figure 6: *Standard deviation for independent controllers (top) and agent-based control (bottom).*

Compared to the independent controllers this is indeed a major advance. It is clear that the market coordination of the controllers leads to a much smaller standard deviation and to higher comfort in the offices. This validates the results of Huberman *et al.* [Huberman and Clearwater, 1995, Clearwater and Huberman, 1994].

5 MARKET-A': A Suite of Variations

In this section we present a suite of variations on the scheme presented in Section 4. The main aim is to see if the scheme can be simplified without loss of performance.

Figure 7: *Standard deviation for the original scheme and a scheme with the gold dependencies removed.*

Deleting the Gold Dependency Plots of the standard deviation with the original scheme as well as with a scheme where all gold dependencies have been removed — by setting $U(0, m)$ in Eq. (19) to a constant value (2000) — are shown in Figure 7. The scheme without gold dependencies performs approximately as good as the original scheme. The reason for this, as pointed out earlier, is that $U(0, m)$'s gold dependency is extremely weak. For the scheme without gold an α of 66 turned out to be optimal.

Figure 8: *Standard deviation for the scheme with the gold dependencies removed, compared to a scheme with both the gold and temperature dependencies removed.*

Deleting the Temperature Dependency The next issue is the marginal utility's dependence on the current temperature, i.e. the dependency of U on T in Eq. (18). In Figure 8 the

standard deviation is plotted for the scheme without gold, mentioned above, and a scheme where the dependency on the temperature has been removed as well. We see that also here the performance is approximately as good as that of the original scheme. In order to remove the temperature dependency from the marginal utility, we simply set all selling prices to 10 and all buying prices to 100. If there is a mismatch between supply and demand, say, supply exceeds demand, the agents that will sell are picked randomly. Here, an α of 65 turned out to be optimal. Note that the temperature is still used to determine both the bidding *volume* and the decision *whether* to buy or sell.

Figure 9: *Standard deviation with the auction scheme removed, compared to the original MARKET-A scheme.*

Deleting the Auction Next, we let the agents assign their bids to themselves without any auction whatsoever. This means that the sum total of the controller outputs, F_{io} , might sometimes exceed the total amount of resource and sometimes be below. The physical model is of course still obeyed so that the total resource actually used, i.e. the sum of the cooling power, P_{cio} , will never exceed the total amount of available resource, as described by Eq. (1). The result of this simulation is shown in Figure 9. Note that this scheme is roughly ten times better than the MARKET-A scheme. Here an α of 17 turned out to be optimal.

Discussion At first glance, it might seem counterintuitive that performance actually *improves* significantly when virtually all the core mechanisms of the market are removed. First we showed that introducing the market improves performance considerably and then we showed that most market mechanisms are superfluous. What, then, is the big difference between the uncoordinated integrating controller, the CONTROL-A scheme, and the scheme without any auction that we ended up with? The simple answer, in our view, is the *access to global data*, in terms of the average temperature and the average setpoint temperature (as seen from Eq. (13) and Eq. (15)), in the MARKET-A and MARKET-A' schemes.

6 CONTROL-B: A Standard Controller Using Global Data

Having concluded that the access to global data is crucial for performance, it is of course of interest to analyze what the performance of an integrating controller, like the one introduced in Eq. (10), *but now incorporating global data*, would be.

We would like the controller to take into account not only its own deviation from its setpoint, but also to consider the deviations of the other offices from their setpoints. Therefore, the controller in Eq. (10) is extended to

$$F_{io} = F_{i-1,o} + \beta[(T_{io} - T_o^{setp}) - (\langle T_i \rangle - \langle T^{setp} \rangle)], \quad (26)$$

where β is set to 10 as previously, and where P_{cio} , as before, is obtained from Eq. (11).

The plot from the simulation with this controller, compared to the MARKET-A simulation, and to the MARKET-A' simulation where the auction was removed, is shown in Figure 10.

Figure 10: *Standard deviation with an integrating controller that utilizes global data compared to the MARKET-A scheme, and the MARKET-A' scheme with the auction removed.*

We see that the standard deviation is approximately the same for the CONTROL-B and the MARKET-A' schemes. Thus, also the CONTROL-B scheme performs roughly ten times better than MARKET-A. An important difference, though, is that CONTROL-B employs the well-known integrating controller for which there are well-understood methods and theories for e.g. stability analysis². In contrast, the MARKET-A scheme is not easily analyzable from a formal theory perspective, since it does not rely on such well established concepts.

²This holds under the assumption that the time scale of the variations of the average temperature is much slower than the fluctuation of the temperature. This is a very reasonable assumption in the present case, however.

7 MARKET-B: A Market with Local Data Only

In the previous sections we have seen that the computational market MARKET-A was outperformed by the global control scheme CONTROL-B. It is an interesting question if a simple and well performing computational market approach can be devised that does *not* depend on having available global information, in contrast to both these schemes. In this section we show that this is indeed the case.

We adopt the concept of resource allocation, as proposed by Wellman [Wellman, 1993], whereby general equilibrium theory is applied to resource allocation. We note that any resource allocation problem with a single resource, like the ones treated in [Ygge and Akkermans, 1996, Kurose and Simha, 1989], can be viewed as a proper market exchange with two commodities: the resource r and money m . We then write the utility function for each agent as $u(r, m) = u(r) + m$. The competitive behaviour for each agent, then, is to trade r for m whenever $\frac{\partial u}{\partial r}$ is below the current market price and vice versa. The *equilibrium price* is a price, p , for r in terms of m , such that supply meets demand. Thus, at market equilibrium all $\frac{\partial u}{\partial r} = p$, except for the agents that are at their boundary values. Agents being at their upper bounds will have a $\frac{\partial u}{\partial r} \geq p$, and agents at their lower bounds will have a $\frac{\partial u}{\partial r} \leq p$. We note that the resulting equilibrium is identical to the Kuhn-Tucker conditions for maximizing the sum of the utilities, when the utilities are concave and independent. Under these conditions, we know that the equilibrium is unique and that it is equal to the solution maximizing the sum of the individual utilities.

7.1 The Relation Between Markets and Independent Controllers Utilizing Global Data

The performance measure for the system is given by Eq. (12). The best system is therefore the one that minimizes this equation. Hence, the most straightforward move one can think of to come to a market model, is to take *this* measure as representing the utility function for the overall system. So, the utility function for the individual agents is ideally related to $[(T_{io} - T_o^{setp}) - (\langle T_i \rangle - \langle T^{setp} \rangle)]^2$.³ However, this is still a formulation containing global information.

Thus, we want to obtain a *purely local* reformulation, by getting rid of the terms with $\langle T \rangle$ containing the global information. We might replace them, though, by terms relating to the changes in the local resource. In doing so, we take inspiration from the standard controller equations Eq. (10) and Eq. (26), indicating that we get good results (for unconstrained resources) with the update equation $F_{io} = F_{i-1,o} + \phi_{io}$, when ϕ_{io} has the form $\phi_{io} = \beta(T_{io} - T_o^{setp})$. The intended interpretation of ϕ_{io} is to represent the output that the local controller would have delivered if acting independently with unconstrained resource.

³Since utilities are expressions of preference orderings, they are invariant under monotonic transformations.

In the market setting each agent updates its control signal F_{io} through $F_{io} = F_{i-1,o} + \Delta F_{io}$, where ΔF_{io} is determined by the outcome of the market. Since the resource is only redistributed among the agents, we have that $\sum_{o=1}^N \Delta F_{io} = 0$. Accordingly, as a step in the design of a MARKET-B scheme, based on local data only, we employ the following definition.

Definition 1. Assume as the utility function for the individual office agents:

$$u(\Delta F_{io}) = -\alpha_o^2 (\Delta F_{io} - \phi_{io})^2, \quad (27)$$

where α_o is a strength parameter for each office representing its physical properties such as R_o and C_o . The proper choice of α_o is discussed in Proposition 2 later on.

Proposition 1. A general equilibrium market in which the agents hold the utility function of Eq. (27) is equivalent to an integrating controller, the resource update equation of which exploits global information and is described by:

$$\Delta F_{io} = \phi_{io} - \frac{1}{\alpha_o^2 \cdot \langle 1/\alpha^2 \rangle} \langle \phi_i \rangle, \quad (28)$$

given that no agent is at its boundary.

Proof. At equilibrium, all $u'(\Delta F_{io})$ will be equal, and thus it will hold for every office that $\Delta F_{io} - \phi_{io} = \frac{\alpha_N^2}{\alpha_o^2} (\Delta F_{iN} - \phi_{iN})$. Summing the equations yields $\sum_{o=1}^{N-1} \Delta F_{io} - \sum_{o=1}^{N-1} \phi_{io} = \alpha_N^2 (\Delta F_{iN} - \phi_{iN}) \sum_{o=1}^{N-1} \frac{1}{\alpha_o^2}$. Together with the resource constraint, $\sum_{o=1}^N \Delta F_{io} = 0$, this gives $\Delta F_{iN} = \phi_{iN} - \frac{1}{\alpha_N^2 \cdot \langle 1/\alpha^2 \rangle} \langle \phi_i \rangle$. For reasons of symmetry, this equation holds for all offices.

Corollary. For the special case where all α_o are equal, and $\phi_{io} = \beta(T_{io} - T_o^{setp})$, Equation (28) becomes exactly the CONTROL-B scheme as captured by Eq. (26). Consequently, a computational market based on the utility function $u(\Delta F_{io}) = -\alpha_o^2 [\Delta F_{io} - \beta(T_{io} - T_o^{setp})]^2$, will always lead to a resource allocation identical to that of the global CONTROL-B scheme.

In sum, we see that *local data plus market communication is equivalent to independent control utilizing global data*. Even though our proof was based on the assumption that the agents are never at their boundary values, it will be a very close approximation in most practical applications. It should also be mentioned that managing the boundaries is not required for a successful implementation. As seen above, omitting the management of boundaries in the current application leads to CONTROL-B, which was shown to have a very high performance.

7.2 Finding an Optimal Utility Function

In this section we show how an optimal utility function is constructed in the *constrained* case, from an optimal controller for the *unconstrained* case.

Proposition 2. If T_{io} is a linear function of ΔF_{io} ,⁴ and if ϕ_{io} minimizes Eq. (12) in the unconstrained case, then in a market where the utility functions are described by Eq. (27), with $\alpha_o = \frac{\partial T_{io}}{\partial \Delta F_{io}}$, the associated market equilibrium minimizes Eq. (12) in the constrained case, if the resource can be independently allocated among the agents.

Proof. Minimizing Eq. (12) boils down to minimizing $f(T_{i1}, T_{i2}, \dots, T_{in}) = \sum_{o=1}^N [(T_{io} - T_o^{setp}) - (\langle T_i \rangle - \langle T^{setp} \rangle)]^2$, since this is just a monotonic transformation. Due to the fact that ϕ_{io} minimizes Eq. (12), it holds that $\frac{\partial f}{\partial \Delta F_{io}} = 0$ for $\Delta F_{io} = \phi_{io}$. This gives that $T_{io}(\phi_{io}) = T_o^{setp} - (\langle T_i \rangle - \langle T^{setp} \rangle)$. Thus, f can be rewritten as $\sum_{o=1}^N [T_{io}(\Delta F_{io}) - T_{io}(\phi_{io})]^2$. Since T_{io} is a linear function of ΔF_{io} , we have that $T_{io}(\Delta F_{io}) = T_{io}(\phi_{io}) + \frac{\partial T_{io}}{\partial \Delta F_{io}} \cdot (\Delta F_{io} - \phi_{io})$, and hence f becomes $\sum_{o=1}^N \left(\frac{\partial T_{io}}{\partial \Delta F_{io}} \right)^2 (\Delta F_{io} - \phi_{io})^2$. Thus $f = -\sum_{o=1}^N u_o$. Since we know (as pointed out in the beginning of Section 7), that the equilibrium maximizes the sum of the utilities we see that it minimizes f .

Discussion. Previously, we saw that the independent controller CONTROL-B that incorporates global data, *viz.* the average temperatures, performs very well. In this section we positively answered the question if one can construct a market, MARKET-B, that is based on local data only and that performs as good or even better.

In this section a market approach based on general equilibrium theory was used. This is of course not the only available mechanism for resource allocation in MAS. It seems interesting to try out other mechanisms, like the contract net protocol [Davis and Smith, 1983], and see if they perform better. However, proposition 2 tells us that, if we treat this problem of building control as separable in terms of agents, like done by Huberman *et al.*, there is no better scheme.⁵ For example, if we assigned all the resource to an auctioneer, that on its turn would iteratively assign the resource in small portions to bidders bidding with their true marginal utility, we would end up with something close to the competitive equilibrium, but we can not do better than MARKET-B. Furthermore, this would be an extremely inefficient way to arrive at equilibrium compared to other available methods [Ygge and Akkermans, 1996, Andersson and Ygge, 1997]. That is, we can use different

⁴From the thermal model discussed in Section 2 (especially Eqs. (4) and (11)) we note that this is indeed the case for a reasonably wide range. Another, and general, reason that this assumption is adequate, resides in the fact that linearization around a working point (here, the minimum of f) is a commonly used and successful procedure in control engineering. Accordingly, the proposition can be read as giving the exact result, when we cut off the Taylor expansion after the linear term. We mention in passing that the error made in f due to this cut-off is only of *third* order with respect to the resource.

⁵We note that, as mentioned earlier, the problem is actually not really separable in terms of agents and therefore better solutions may exist if this is taken into account. However, we believe that for managing non-separability, centralized algorithms are likely to be even more competitive compared to distributed ones. Another observation is that in this paper we have only investigated the case of using the current available resource as the interesting commodity, as done by Huberman *et al.*, and found an optimal mechanism for that. At the same time, extending the negotiations to future resources as well, could potentially increase performance, but this is a different problem setting with different demands on available local and/or global information items, such as predictions.

mechanisms for achieving the competitive equilibrium, but we can never hope to find a mechanism that would do better than the MARKET-B scheme.

8 Discussion and Conclusion

We believe that both the approach and the results, as presented in this paper, pose a general challenge to the MAS community. Multi-agent systems offer a very new way of looking at information technology, with potentially a big future impact. They may lead to what Kuhn calls a ‘paradigm shift’ [Kuhn, 1970]. However, when new paradigms arise on the scientific horizon, they must prove their value in comparison and competition with existing, more established ones. The MAS paradigm is no exception.

We have therefore deliberately played the role of the devil’s advocate in this paper. Any new paradigm in science sparks off enthusiasm. But this tends to lead to a parochial view, on its turn bringing along exaggerated claims and promises, and thus to wrong expectations from society. Pride and prejudice are dangers lurking around the corner, even in the field of MAS. These are not imaginary dangers, because the history of Artificial Intelligence and computer science is full of examples. Take, as an example, the field of knowledge-based systems. Knowledge systems actually are a success story. They are now in everyday use around the world. Yet, we are very far from the ‘copying and automating human experts’ image that was pictured to the public in the early eighties. According to some, such ill-conceived expectations contributed to the so-called ‘AI Winter’ later in the decade. Still, knowledge systems constitute a positive example, as they represent the single most important commercial and industrial offspring from AI. But their role has become much more modest, as intelligent support tools or assistants, usually functioning as knowledge-intensive components embedded within much larger conventional information systems.

So, the key question to be answered by the MAS community is: in what respect and to what extent are multi-agent solutions better than their more conventional alternatives? And the key message of this paper is that arguing in favour of the multi-agent systems paradigm does require much more careful analysis. As yet, there are hardly any studies in the MAS literature of the kind carried out in the present paper. But as we have shown in technical detail, old paradigms such as conventional control cannot be *that* easily dismissed. Paradigm shifts and scientific revolutions are brought about, not only through noisy crowds (although these can indeed be very helpful), but by lots of hard everyday technical work.

Abstract considerations alone, concerning what agents are, or on the general nature of autonomy, rationality, and cooperation, are not sufficient to achieve this. Therefore, we have taken a different approach, one that is bottom-up and application-oriented. In this way, we aim at obtaining experimental data points on the basis of which convincing MAS claims can be established. The data point considered in this paper is climate control within

a large building consisting of many offices. This is a quite prototypical application relating to the general problem of optimal resource allocation in a highly distributed environment. This class of problems has already received much attention in the MAS area. Reportedly, this type of application is very suitable for market-oriented programming [Huberman and Clearwater, 1995]. On the other hand, we have devised some more conventional control engineering solutions, as well as alternative market designs. Important conclusions of our investigation are:

- The market approach by [Huberman and Clearwater, 1995] indeed outperforms a standard control solution based on local, independent controllers. So, the MAS approach indeed yields a working solution to this type of problem.
- However, if conventional control schemes are allowed to exploit global information, as does the market design of Huberman *et al.*, they perform even better.
- We have proposed an alternative market design that uses local data only, and still performs as good as a control engineering scheme having access to global information.
- So, our general conclusion is that “*local data + market communication = global information*”. The important difference is that in computational markets this global information is an *emergent property* rather than a presupposed concept, as it is in standard control.

These results provide the opportunity to more clearly state what the added value of market-oriented programming is in this type of applications. There is a scientific role for debunking here. As we have seen, one *can* model the situation in the more traditional terms of control engineering. A similar argument regarding distributed resource allocation could be construed, by the way, vis-à-vis mathematical optimization (*cf.* our discussion in [Ygge and Akkermans, 1996]).

In our analysis we have focused on the market approach. It is tempting to ask whether things are different when a non-market MAS approach is followed, say, using the contract net [Davis and Smith, 1983]. As we argued, the answer in our opinion is a straightforward *no*. Arguing for non-market approaches is not a way out, but a dead-end street. The goal in the considered class of problems is to find the optimal distributed solution. Alternative MAS approaches, market as well as non-market ones, only change the multi-agent dynamics on the way to this goal. This might be done in a better or poorer way, but it is not possible to change this goal itself. The goal state in any MAS approach is, however formulated, equivalent to market equilibrium, the yardstick for having achieved it is given by some quantitative performance measure as we discussed, and both are stable across different MAS approaches.

The agent and market approach does give a highly intuitive and natural picture of problems in distributed environments — even when at a strictly algorithmic level it effectively

leads to the same end result as alternative approaches. We do believe that this is a value in itself: for conceptual modelling, understanding, explanation, knowledge transfer. Moreover, models and pictures that have conceptual simplicity are more easily generalized to more complicated situations, e.g., to allocation of multiple resources in multicommodity markets [Wellman, 1993]. Because of its focus on local information, communication and action, the agent paradigm is more flexible than centralized approaches. This does not show up very clearly in the case analyzed here, because the underlying physics and mathematics has been strongly simplified (for example, all offices and thermodynamic processes are assumed to be equal). This will generally not be true in reality, and that will necessitate a big modelling effort in centralized approaches as standard control engineering. For the present case, we have shown that the MAS approach offers working solutions of equal quality (see the simulation experiments and, for a formal and general result, our Proposition 1). As a point in favour of a MAS approach, we note that it can treat resource constraints relatively easily, see especially our Proposition 2.

So, when we have strong heterogeneity, large scale, and relative independence, we believe that the agent approach starts to pay off. This is visible for example in the application of power load management we are working on ourselves [Ygge and Akkermans, 1996]. It does have similarities with the application discussed in this paper, but for the reasons mentioned the current industrial state of the art based on central control approaches is quite limited. Notwithstanding this, it does not diminish the need for thorough analysis — of (market) failures *and* successes.

Acknowledgments

We thank Rune Gustavsson at the University of Karlskrona/Ronneby and Hans Ottosson at EnerSearch AB for all their support. A special acknowledgment goes to Olle Lindeberg whose very detailed comments on the draft papers led to significant improvements in the simulations and whose ideas also helped us in the design of the market in Section 7. We also thank Bengt Johannesson, who, as a part of his master's thesis, went through our paper in detail and recreated all the simulations. We thank Arne Andersson, Eric Astor and the SoC team at the University of Karlskrona/Ronneby for useful comments and discussions of draft material.

References

[Andersson and Ygge, 1997] A. Andersson and F. Ygge. Managing large scale computational markets. 1997. In preparation.

- [Borst *et al.*, 1997] W.N. Borst, J.M. Akkermans, and J.L. Top. Engineering ontologies. *International Journal of Human-Computer Studies*, 46:365–406, 1997. ISSN 1071-5819.
- [Bradshaw, 1997] J.M. Bradshaw. *Software Agents*. AAAI Press/The MIT Press, Menlo Park, CA, 1997.
- [Clearwater and Huberman, 1994] S. Clearwater and B. Huberman. Thermal markets for controlling building environments. *Energy Engineering*, 91(3):25–56, 1994.
- [Davis and Smith, 1983] R. Davis and R. G. Smith. Negotiation as a metaphor for distributed problem solving. *Artificial Intelligence*, (20):63–109, 1983.
- [Huberman and Clearwater, 1995] B. A. Huberman and S. Clearwater. A multi-agent system for controlling building environments. In *Proceedings of the First International Conference on Multi-Agent Systems ICMAS'95*, pages 171–176, Menlo Park, CA, June 12–14 1995. AAAI Press / The MIT Press.
- [Incropera and Witt, 1990] F.P. Incropera and D.P. De Witt. *Fundamentals of Heat and Mass Transfer, third edition*. Wiley and Sons, New York, 1990. ISBN 0-471-51729-1.
- [Kuhn, 1970] T.S. Kuhn. *The Structure of Scientific Revolutions*. The University of Chicago Press, Chicago, IL, 1970. Second Edition.
- [Kurose and Simha, 1989] J. F. Kurose and R. Simha. A microeconomic approach to optimal resource allocation in distributed computer systems. *IEEE Transactions on Computers*, 38(5):705–717, 1989.
- [Ogata, 1990] K. Ogata. *Modern Control Engineering, Secod Edition*. Prentice-Hall, Englewood Cliffs, NJ, 1990. ISBN 0-13-589128-0.
- [Wellman, 1993] M. Wellman. A market-oriented programming environment and its application to distributed multicommodity flow problems. *Journal of Artificial Intelligence Research*, (1):1–23, 1993.
- [Ygge and Akkermans, 1996] F. Ygge and J.M. Akkermans. Power load management as a computational market. In M. Tokoro, editor, *Proceedings of the Second International Conference on Multi-Agent Systems ICMAS'96*, pages 393–400, Menlo Park, CA, December 9–14 1996. AAAI Press.