A Multi-Agent System for Controlling Building Environments

Bernardo Huberman and **Scott H. Clearwater**

Xerox Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, CA 94304, U.S.A.
huberman@parc.xerox.com

Xerox Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, CA 94304, U.S.A.
clearwat@parc.xerox.com

Abstract

A working solution to the problem of thermal resource distribution in a building is demonstrated using a market-based system. In this system computational agents representing individual temperature controllers bid to buy or sell cool or warm air. They do so via a double-blind computerized auction which is moderated by a central computer auctioneer. The auctioneer sees to it that no agent buys resources for more than its bid and no agent sells resources for less than its bid. The market system has been implemented and runs on a regular basis as part of a building energy management system. Results show that the thermal auction leads to an equitable temperature distribution throughout the area under its control without incurring any extra costs such as excessive actuator movement.

Introduction

Multiagent systems are characterized by the interaction of many agents trying to solve problems in cooperative fashion. In these systems, each agent has an individual, local view of the problem while having to successfully interact with other agents in order to find global solutions to the task at hand. One of the difficulties in both designing and understanding multiagent systems comes from the lack of central controls, and the ensuing conflicting, uncertain, and delayed knowledge on the part of the agents. Thus, there is a need for understanding the behavior of large collections of locally-controlled, asynchronous and concurrent processes interacting with an unpredictable environment. This set of issues, which are also found in social and biological communities, lead us to refer to a multiagent systems as ***computational ecosystems*** [1].

The problems faced by multiagent systems are solved on a daily basis in the economic context through the institution of markets. Markets provide examples of multiagent systems that deal successfully with the problem of coordinating asynchronous operations in the face of imperfect knowledge. A computational system set up along market rules can allow the system as a whole to adapt to changes in the environment or disturbances to individual members. This analogy has been exploited in the past for scheduling computational tasks in network environments [2–4]. In effect, a coordinated solution to the allocation problem was obtained using Adam Smith's "invisible hand" [5]. Although unlikely to produce the optimal allocation that would be made by an omniscient controller with unlimited computational capability, price mechanisms perform well compared to other feasible alternatives [6, 7]. As in economics, the use of prices provides a flexible mechanism for allocating resources, with relatively low information requirements [8]: a single price summarizes the current demand for each resource.

In this paper we present a working solution to the problem of thermal resource distribution in a building using a market-based multiagent system. Computational agents representing individual temperature controllers bid to buy or sell cool or warm air. In this application agents interact via a double-blind computerized auction which is moderated by a central computer auctioneer. The auctioneer sees to it that no agent buys resources for more than its bid and no agent sells resources for less than its bid. Results show that the thermal auction leads to an equitable temperature distribution throughout the area under its control without incurring any extra costs such as excessive actuator movement.

Current control schemes fall into one of two classes; independent controllers and centralized controllers. Independent controllers act as agents and utilize tried and true control techniques such as proportional-integral-derivative (PID) control, but totally ignore interactions between agents. This shortcoming reveals itself in the often inadequate quality of control resulting from these types of controllers. The interaction among agents provides additional degrees of freedom with which to improve control.

While in principle an omniscient central controller with access to all the environmental and thermal parameters of a building (i.e. a perfect model) could optimally control it, in practice such knowledge is seldom available to the system [9]. Instead, partial information about local changes in the variables (such as instantaneous office occupancy, external temperature, and computer use) is the only reliable source that can be used for controlling the building. This leads to the notion of a distributed control system, where most of the decisions are made at the local (i.e. office) level, while achieving global desirable behavior.

Using Markets for Distributed Control

Given the similarities between social organizations and computational ecosystems, it is tempting to test the efficacy and desirability of market mechanisms in distributed control systems. Since market devices such as prices and auctions greatly facilitate resource management in human societies, one would expect them to be similarly useful in managing resources for other control applications, e.g. building environmental systems. A price mechanism could allow offices with different thermal needs to have different temperatures.

The similarity of the temperature distribution problem to resource allocation in market economies is reminiscent of the problem of computation distribution and the use of markets for scheduling computational tasks in network environments [2–4]. In effect, a coordinated solution to the allocation problem was obtained using Adam Smith's "invisible hand" [5]. Although unlikely to produce the optimal allocation that would be made by an omniscient controller with unlimited computational capability, price mechanisms perform well compared to other feasible alternatives [6, 7]. As in economics, the use of prices provides a flexible mechanism for allocating resources, with relatively low information requirements [8]: a single price summarizes the current demand for each resource, whether cooling or heating or controlling.

This flexibility is especially desirable when resource preferences and performance measures differ among occupants. For example, just as utilities buy and sell power among themselves, one can envision offices in a building trading for cool or warm air according to their individual needs. A reasonable allocation of thermal resources could thus be brought about by simply having the individual temperature controllers be sensitive to current thermal resource price. Moreover, adjustments can take place continually as occupants change their preferences or leave their offices, and do not require all users to agree on, or even know about, these new uses, thus encouraging an incremental and experimental approach to resource allocation.

While this example motivates the use of market based resource allocation, an actual implementation is required to see how large the system must be for its benefits to appear, and whether any of the differences between simple computer programs running the temperature controllers and human agents doing so pose additional problems. In particular, a successful use of markets requires a number of changes to traditional building control systems. First, the system must provide an easily accessible, reliable market so that buyers and sellers can quickly find each other. Second, temperature controllers must be price sensitive so that they can respond to changes in relative prices among resources. As we show below, these two considerations led to the implementation of a double blind auction system that effectively answers these concerns [10].

System Description

There are two basic components to the energy auction. The first is the collection of agents that represent the players in the energy auction. The second is the auctioneer who oversees the auction. All the agents run within a single program, although this need not be the case. The temperature sensors and air flow actuators are all accessible directly through distributed hardware modules via a network connection.

Data Structures

The agents participating in the thermal market are constructed with the characteristics shown in the table below.

The correlation with other agents is a function of the air duct layout. Usually offices that are nearer can trade air more effectively than those that are farther apart although we have measured a nearly 80% transfer of air from offices that are the farthest apart of all the agents in the auction. To know the correlations between offices is quite complicated and for simplicity we assumed that all the offices could transfer air equally to any other office.

The Auction

The purpose of the auction-based control is to provide the most comfort for a given amount of thermal resource. This contrasts with conventional (PID) controllers that strive to maintain the comfort of a single office. The problem with conventional controllers is that while some offices can be comfortable there are others at the same time that may be much less comfortable. The auction-based controller distributes the thermal resource so that *all* the offices are at the same comfort level. For example, on a very hot day a PID-controlled building will have some offices that are comfortable and some that are too hot, whereas in an auction-controlled building all the offices will be uniformly comfortable.

Auctions are an effective mechanism for redistributing goods based on the ability of participating agents to pay for them. The auction we used in our building

Thermal Market Agent Structure	
Name	name of the office represented by the agent
Setpoint temperature	the desired temperature in the office
Current temperature	the temperature in the office
Money on hand	the amount of money the agent can spend on an auction
Amount of air desired	expressed as a relative amount, could be converted to air flow
Sell price	the price the agent is willing to sell air at; depends on previous auction price as well as setpoint and current temperature
Buy price	the price the agent is willing to buy air at; depends on previous auction price as well as setpoint and current temperature
Correlation with other agents	efficiency with which air flow can be shunted from one office to another

control system belongs to the class known as *double-blind auctions* [11]. In double-blind auctions none of the potential buyers or sellers, i.e. *agents*, knows the value of any bid besides its own. A bid consists of a quantity (in this case thermal units) to be traded at a given price. A *buy bid* means that an agent is willing to buy a certain quantity of thermal units at or below its bid. A *sell bid* means that the agent is willing to sell a certain quantity of thermal units at or above its bid. All agents whose sell bid is at or below the auction price and all agents whose buy bid is at or above the auction price will have their trades consummated and receive the thermal resources they requested. Their accounts are then adjusted by the appropriate amount. The price determined by the auctioneer for this auction is about 9.5, which is the price where the supply first exceeds demand. The auctioneer is a specialized agent that does not represent any office buts acts as the clearinghouse for all energy trades. The auctioneer collects all the bids and determines the *supply* and *demand curves* as a function of bid price. The supply and demand curves, in turn, determine the *volume* that is available for trade at a given price. At the point where the supply first meets the demand the auctioneer sets the *auction price*. Note how the supply rises with price (agents are more willing to sell at a higher price) and the demand falls with price (agents are less willing to buy at a higher price) just as in real auctions.

For example, suppose an agent bids for more cool air because the temperature of its office is above its setpoint. If its bid price is at or above the auction price then the air flow actuator (a butterfly valve called a VAV for Variable Air Volume) for its office will open further to receive more cool air. Correspondingly, an office that is too cool will close its VAV by some amount. Notice that there is no direct agent-to-agent communication in this process.

The decision for an agent to buy or sell is determined from the following formula:

$$t_i = \frac{T_i^{setpt}}{T_i} \cdot \frac{\langle T \rangle}{\langle T^{setpt} \rangle} \quad \begin{cases} > 1, \textit{ seller} \\ < 1, \textit{ buyer} \end{cases} \tag{1}$$

where T_i is the temperature of the ith office, T_i^{setpt} is the temperature set point of the ith office, $\langle T \rangle$ is the average temperature of all the offices, and $\langle T^{setpt} \rangle$ is the average temperature set point of all the offices. The mean temperatures are used to renormalize the temperature to guarantee that there are always some *potential* buyers and sellers.

In our system the auctions are typically run at one minute intervals throughout the day. In other words, 1440 auctions take place everyday.

The actual bid of an agent for a given auction is a non-linear function of the current office temperature, the set point temperature and amount of money the agent has. This function closely follows the work of [12], although our results are not qualitatively dependent on their model. In our system each agent receives an allocation of money to spend for each auction. It should be pointed out that this money is not real, but the agents treat it as such. Whatever money is not spent for a given auction disappears. This "burning" of money after every auction prevents runaway inflation in the price. The amount of money an agent has to bid ranges linearly from a low of 100 when the VAV is fully open to a high of 200 when the VAV is fully closed. The reason for the agent having some money left when its VAV is fully open, even though it cannot open further, is that it can still pay another agent to close its VAV further, so that even a fully open VAV can still get more air. Dribbling money to the agents for each auction prevents them from bidding a day's worth of money on the first auction and having nothing left for the rest of the day.

Unlike human auctions where one gets what one pays for, thermal markets are not as efficient. This is due to the unavoidable inefficiencies in the air distribution system. The air flowing through the duct work is a fluid flowing through pipes or analogous to an electrical current flowing through a resistor network. The duct work has a "resistance" that is a function of its shape that prevents proper distribution of air to the agent that wants it. This means that agents that are topologically in between any pair of buyers or sellers will get more or less air even if they did not explicitly take part in the auction. Thus, a very poorly designed duct system could ruin the performance of the auction. However, this could be overcome by having the agents adapt over time so that they effectively learned the correlations between all the offices. This would reduce the effect of the "tax" on air distribution. The overall effect would shorten the time to thermal equilibrium among the offices.

Our electronic auction also differs from a real auction in that resource is not transferred from one agent to another, rather the resource that *would have* gone to one agent is sent to others by adjusting the VAVs. This is the same thing that happens in a resistor network, there is no direct transfer or current between resistors only a change in current flow based on a change in impedance. However, just as in a real auction, in the unusual case when all the agents have the same temperature there is no trading because all the agents are at their desired setpoints. The thermal market auction requires a non-zero spread in the utility function of the agents for there to be resource trading and therefore control. Of course, all the agents being at their setpoints is usually only a transient situation and once one of the temperatures drifts for whatever reason, there will be some trading.

Experimental Setting The auction system gets VAV and temperature data from the conventional building control system, performs the auction and sends out the new VAV settings to the control system. The setting is part of the Xerox Palo Alto Research Center and contains 13 offices and 2 open areas constituting approximately 3000 square feet. The experimental set-up is shown in Fig. 1.

The control system used for comparison has been developed and tuned by facility engineers for over two years and has resulted in substantial energy savings and improved comfort. Still, there is room for improvement and this is where the auction-based control comes into play.

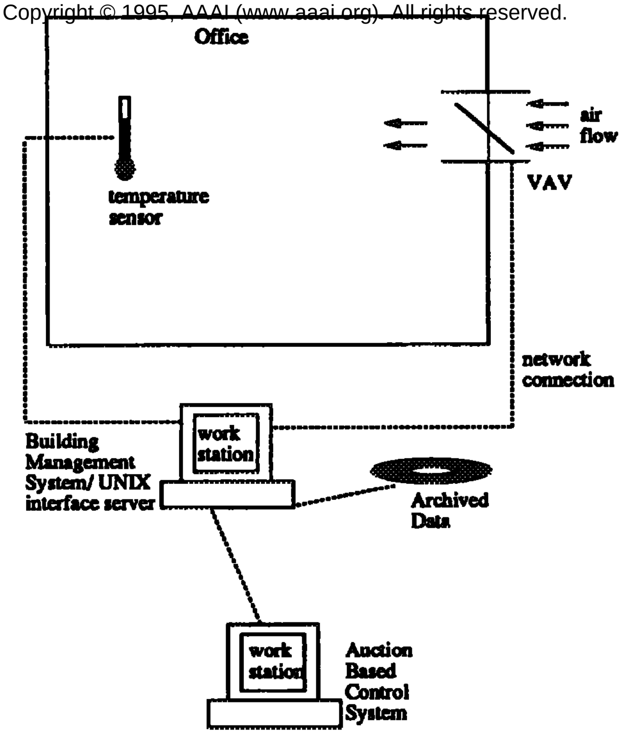

Fig. 1. Schematic diagram of the instrumentation in an office and the connections to the computerized auction.

Results

In this section we describe the results of the auction-based control for controlling office environments and compare the results to those obtained using conventional control techniques.

Since the goal of the auction-based control is to balance resources fairly, we can measure the success of the auction by determining how widely all the office temperatures are distributed from their setpoints at a given time. The standard deviation of a set of numbers, defined as,

$$\sqrt{\frac{1}{N}\sum_{i=1}^{N}(x_i-\bar{x})^2} \quad (2)$$

gives a measure of how widely those numbers are distributed. Thus the effect of the auction will be to narrow the distribution of temperatures from their setpoints among the offices participating in the auction. Note that the auction as currently configured cannot adjust the *average* temperature, $\langle T\rangle$, of all the offices because $\langle T\rangle$ depends on the amount of resource pumped into the area as well as outdoor conditions, which are not a part of the auction.

A striking comparison between conventional and auction-based control is shown in Fig. 2. On the left side of the figure (June 24, 1993), the integral control

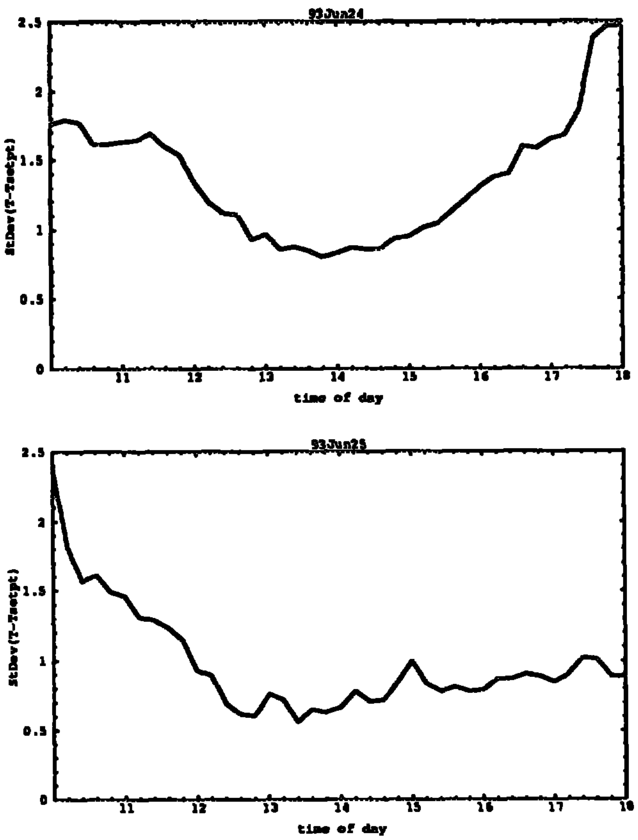

Fig. 2. The standard deviation of the average temperature minus the average setpoint (StDev<T-Tsetpt> on the plots) versus the time of day in hours for two consecutive days. (top) integral control. (bottom) auction-based control. The measurements are averages over a 12 minute time span. Note the narrower spread of deviations from setpoint for the auction-based control. All setpoints were identical on the two days.

was in operation all day and the maximum temperature for that day was $102°F$. On the right side of the figure (June 25, 1993) the auction-based control was in operation starting at 10am and the maximum temperature for that day was $106°F$. Note that it took a couple of hours for the auction to significantly reduce the spread of temperatures. Quantitatively, the average standard deviation from noon to 6pm was $0.74°F \pm 0.11°F$ for the auction and $1.39°F \pm 0.43°F$ for the integral, a significant difference. The "$\pm$" is the standard deviation of the average standard deviation which is a measure of how stable the spread in temperatures was over the course of the experiment.

It is obvious from the data that auction-based control offers significantly better resource allocation even in a highly tuned building management system and does so with greatly reduced oversight.

Games Agent Play While the computational agents participating in the auction have been programmed to "believe" in the value of the money they are bidding, that is not the case for the people who have control of their office's thermostats. This creates an interesting environment where the computational agents are playing as if in a real economy but the people in the offices are playing as if the thermal resources were free. This is quite different from the situation in people's homes where they have to pay for the thermal resource they use. This means that office occupants can "game" the system by adjusting their thermostats to unrealistic values just so that they can be comfortable at the expense of all the other offices in the economy. We have tested this "free-riding" effect by changing various auction parameters such as number of buyers and sellers, trade volume, and price, with different temperature setpoints, one with fixed setpoints and the other with "gamed" setpoints. We noted drastic changes in the auction, for some times there is no trading which means no control. The reason is that the price is driven to very high values by the demands of the unreasonable setpoints. This causes the demand to exceed supply and drives the price higher. Eventually the price reaches a value greater than the amount of money that the agents have. Note however, that the auction did eventually recover and control was regained. The conclusion is that some degree of temperature setpoint can be tolerated but if it becomes too large, other options are needed.

Discussion

In this paper we presented a working solution to the problem of thermal resource distribution in a building using a market-based multiagent system. Computational agents representing individual temperature controllers bid to buy or sell cool or warm air. The market system has been implemented and runs on a regular basis as part of a building energy management system. Results show that the thermal auction leads to an equitable temperature distribution throughout the area under its control without incurring any extra costs such as excessive actuator movement. The relative ease of installation and maintenance of the auction-based control means that it is more likely to be used than systems that are disconnected because the facility engineers won't use them for a number of reasons.

An interesting extension is to turn the "funny" money that the computational agents use into real money. This directly couples energy cost to the people who are using it. This level of submetering is currently not available in most office buildings but may become more prevalent as utility rate structures become more complicated.

Any energy savings that may accrue will depend on the particular control strategy. If the overall goal of the strategy is to maintain an average temperature then there is likely to be little savings. If the strategy is to pump as much thermal resource into least comfortable space so that it becomes comfortable, then auction-based controls will lead to energy savings because of its superior ability to balance resource [13].

Perhaps the most interesting extension is to make the thermal market a multi-commodity market. During the summer and winter the building is run in a purely cooling or purely heating configuration. However, in fall and spring there are days when both heating and cooling occur simultaneously. This is done with dual duct (one cool and one hot set of duct work) systems or reheat systems. In situations like these there are two resources to be traded. Our preliminary results with this set-up show that once again auction-based control provides better control than the conventional system [14]. Finally we should mention that the use of our market based controls can be easily carried over to other situations where energy balance in diverse settings is needed.

Acknowledgments

The authors wish to thank Rick Costanza, Jim DesRivieres, Mike Dixon, Scott Elrod, and Gene Hall for their efforts in making the hardware and software of the experimental setup a reality.

References

[1] Bernardo A. Huberman and Tad Hogg. The behavior of computational ecologies. In B. A. Huberman, editor, *The Ecology of Computation*, pages 77–115. North-Holland, Amsterdam, 1988.

[2] T.W. Malone, R. E. Fikes, K. R. Grant, and M. T. Howard. Enterprise: A market-like task scheduler for distributed computing environments. In B. A. Huberman, editor, *The Ecology of Computation*, pages 177–205. North-Holland, Amsterdam, 1988.

[3] Carl A. Waldspurger, Tad Hogg, Bernardo A. Huberman, Jeffery O. Kephart, and W. Scott Stornetta. Spawn: A distributed computational economy. *IEEE Trans. on Software Engineering*, 18(2):103–117, February 1992.

[4] I. E. Sutherland. A futures market in computer time. *Communications of the ACM*, 11(6):449–451, June 1968.

[5] Adam Smith. *An Inquiry into the Nature and Causes of the Wealth of Nations*. Univ. of Chicago Press, Chicago, 1976. Reprint of the 1776 edition.

[6] Donald Ferguson, Yechiam Yemini, and Christos Nikolaou. Microeconomic algorithms for load balancing in distributed computer systems. In *International Conference on Distributed Computer Systems*, pages 491–499. IEEE, 1988.

[7] James F. Kurose and Rahul Simha. A microeconomic approach to optimal resource allocation in distributed computer systems. *IEEE Transactions on Computers*, 38(5):705–717, 1989.

[8] Friedrich A. Hayek. Competition as a discovery procedure. In *New Studies in Philosophy, Politics, Economics and the History of Ideas*, pages 179–190. University of Chicago Press, Chicago, 1978.

[9] Z. Cumali. Global optimization of hvac system operations in real time. *ASHRAE Transactions*, 94:1729–1744, 1988.

[10] Scott H. Clearwater and Bernardo Huberman. Thermal markets for controlling building environments. *Energy Engineering Journal*, 91:26–55, 1994.

[11] D. Friedman and J. Rust. *The Double Auction Market: Institutions, Theories and Evidence*. Addison Wesley, Reading MA, 1992.

[12] Kenneth Steiglitz, Michael L. Honig, and Leonard M. Cohen. A computational market model based on individual action. In Scott H. Clearwater, editor, *Market-Based Control: A Paradigm for Distributed Resource Allocation*, chapter 1, page (in press). World Scientific, Singapore, 1995.

[13] Scott H. Clearwater, Rick Costanza, Mike Dixon, and Brian Schroeder. Saving energy using market-based control. In Scott H. Clearwater, editor, *Market-Based Control: A Paradigm for Distributed Resource Allocation*, chapter 10, page (in press). World Scientific, Singapore, 1995.

[14] S. Clearwater, R. Costanza, and B. Schroeder. Two commodity thermal markets in building energy management. *Energy Engineering Journal*, page (in press), 1995.