Simple Bargaining Agents for Decentralized Market-Based Control

Dave Cliff*, Janet Bruten
Living The Vision
HP Laboratories Bristol
HPL-98-17
February, 1998

E-mail: jlb@hplb.hpl.hp.com
davec@ai.mit.edu

market-based control,
economic agents,
bargaining,
autonomous agent

Market-Based Control (MBC) is a resource allocation and control technique where multi-agent systems are built to resemble free-market economies. The aim is that MBC systems exhibit the same decentralization, robustness, and capacity for self-organization as do real economies. MBC systems are relevant to Artificial Intelligence (AI) and robotics in at least two ways: first, the agents in a MBC system need to be robot-like in their ability to autonomously coordinate perception and action in dynamic and uncertain environments that include other agents; second, MBC systems could be used as the control technologies for robots and other "intelligent" autonomous agents. We critically review a selection of MBC systems. We argue that the MBC systems reviewed here are either implicitly reliant on centralized knowledge, or require human operators and hence are not truly automatic. We identify a major issue in creating truly decentralized and automatic MBC systems: the need for the system's agents to be capable of *bargaining behaviors*. Following this, we briefly summarize our current results and ongoing work in creating multi-agent systems where each autonomous agent has the ability to bargain with other agents. We demonstrate that markets composed of such agents exhibit desirable behaviors, and that such agents could form the basis of *truly* decentralized MBC systems.

*Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, Massachusetts

Simple Bargaining Agents for Decentralized Market-Based Control

Dave Cliff
Artificial Intelligence Laboratory
Massachusetts Institute of Technology
CAMBRIDGE MA 02139, U.S.A.
davec@ai.mit.edu

Janet Bruten
Hewlett-Packard Laboratories Bristol
Filton Road
BRISTOL BS12 6QZ, U.K.
jlb@hplb.hpl.hp.com

ABSTRACT

Market-Based Control (MBC) is a resource allocation and control technique where multi-agent systems are built to resemble free-market economics. The aim is that MBC systems exhibit the same decentralization, robustness, and capacity for self-organization as do real economies. MBC systems are relevant to Artificial Intelligence (AI) and robotics in at least two ways: first, the agents in a MBC system need to be robot-like in their ability to autonomously coordinate perception and action in dynamic and uncertain environments that include other agents; second, MBC systems could be used as the control technologies for robots and other "intelligent" autonomous agents. We critically review a selection of MBC systems. We argue that the MBC systems reviewed here are either implicitly reliant on centralized knowledge, or require human operators and hence are not truly automatic. We identify a major issue in creating truly decentralized and automatic MBC systems: the need for the system's agents to be capable of *bargaining behaviors*. Following this, we briefly summarize our current results and ongoing work in creating multi-agent systems where each autonomous agent has the ability to bargain with other agents. We demonstrate that markets composed of such agents exhibit desirable behaviors, and that such agents could form the basis of *truly* decentralized MBC systems.

1 Introduction

We argue here that much current work in market-based control (MBC) is flawed. MBC is an approach to automatic resource allocation and control that draws inspiration from free-market economics. Real free markets are self-organizing and decentralized, and a primary attraction of MBC is its potential for creating artificial systems that also have these properties. Despite this, we show here that much work in MBC relies on *centralized* processes, often in an implicit manner. One likely reason for this is that centralizing the market price-fixing process avoids the need to create 'trading agents' capable of *bargaining* with each other (i.e., agreeing a price for a transaction). This need to create trading agents capable of bargaining is undeniably challenging, but if the challenge is avoided by centralizing the market then one of the primary advantages of MBC disappears, and so the entire endeavor becomes questionable.

Yet if such trading agents could be developed, truly decentralized MBC systems could be created. Such MBC systems would have strong links with artificial intelligence and robotics in two ways. First, creating the necessary trading agents requires addressing problems that research in AI and mobile robotics has studied for some time. Second, if truly decentralized MBC systems can be created, they could be of significant use in robotics and AI applications. In both cases, developing solutions requires intensive use of simulations of agent-environment interaction, where each trading agent's environment consists primarily of other trading agents.

In addition to presenting our critique of the MBC literature, we demonstrate here that simple trading agents, using machine-learning techniques common in artificial intelligence (AI), can give human-like trading behavior. We show that our agents stabilize at an equilibrium that is predictable from economic theory, and are robust with respect to sudden changes in the market. Thus, our trading agents can be used in decentralized MBC applications.

The rest of this paper is structured as follows. Section 2 discusses the background to this paper; Section 3 presents our critique of current MBC systems; and Section 4 gives a brief overview of our simple trading agents that can bargain in decentralized markets.

2 Background

Recent work in AI and robotics has addressed issues in creating systems that exhibit desirable complex "emergent" behavior. Such systems are typically *distributed* in the sense that they are composed of a number of separate but interacting autonomous subsystems. Crucially, the interactions between the subsystems, and the specifications of the subsystems themselves, are simple in comparison to the behavioral complexity of the overall system. Thus, the complex behavior *emerges* from simple interactions between simple components. Often, inspiration is drawn from biological systems. Two prominent cases of such systems are *artificial neural networks* in AI and *collective behavior* robotics systems.

In artificial neural networks, "intelligent" behaviors are exhibited by computational systems composed of many simple "neuron-like" processing units. Each unit typically does nothing more than compute a threshold function of a weighted sum of several scalar numeric input values, producing a scalar output value. Appropriately configured, networks of such units can exhibit surprisingly sophisticated processing behaviors: see e.g. [27, 21].

Collective behaviors are familiar in many animal species: flocking in birds; schooling in fish; swarming in insects; and herding in terrestrial animals are all common examples. In seminal simulation studies, Reynolds [25] demonstrated that such coordinated group behaviors could arise from simple agents moving so as to satisfy a small set of simple constraints. Subsequently, a number of researchers worked on collectives of simple robots coordinating to achieve tasks that might otherwise have been tackled by a single complex robot: see e.g. [3, 20].

Systems can exhibit emergent behavior at multiple levels. For example, a collective of robots could be engaging in activity that is emergent at the group level; yet the controller for each robot could be a neural network, in which case the behavior of an individual robot is itself an emergent consequence of the interactions between the 'neurons' in that robot's controller.

A primary advantage of distributed emergent systems such as these is that they are *decentralized*: there is no central coordinating component or process on which the system relies. Such decentralization offers ***robustness*** in the form of ***graceful degradation***: the absence of centralization can lead to systems that show a progressive loss of performance, rather than sudden catastrophic collapse, when individual subsystems malfunction or fail. Often this robustness is a consequence of the distributed emergent system having some capacity for *self-organization*: there is no need for explicit reprogramming when a subsystem fails; rather, the remaining subsystems alter their activity to accommodate the change in their operating environment. Sometimes this self-organization also extends to the initial programming or calibration of the system.

It is beyond the scope of this paper to give a full review of biologically-inspired emergent computation: see Forrest [12] or Huberman [17] for significant collections. Rather, our primary aim here is to emphasize that there is another type of decentralized system that shows desirable complex emergent behaviors. The distributed emergent system that concerns us here is the *free-market economy*. In the last few years, a small but growing number of researchers have explored the prospect of using natural human economics and market mechanisms as metaphors for constructing computational solutions to difficult problems in resource allocation and control.

The allocation of scarce resources is a topic that has long been studied in economics. In brief, if the quantity of a resource demanded by consumers in a market is greater than the quantity supplied, competition between consumers causes the price of the resource to rise. This can both reduce the quantity demanded (because some consumers can no longer afford it) and also increase the quantity supplied (because some suppliers may be more interested in selling at higher prices). Similarly, when the quantity supplied is greater than the quantity demanded, competition among suppliers can lead to the price falling, which may reduce the quantity supplied and increase the quantity demanded. Thus, according to classical economic theories of markets, the market *equilibrates*: transaction prices approach an ***equilibrium*** value where the quantity demanded matches the quantity supplied.

Advocates of free-market economics claim that the actions of groups of individuals, engaging in simple trading interactions driven by self-interest, can result in good or optimal allocation of resources. Crucially, the market mechanism does this in a decentralized fashion: there is no central control process; rather, the allocation 'emerges' from the interactions of the buyers and sellers.

Thus, economics can act as a valuable source of terminology, inspiration, and metaphors for developing solutions to problems in distributed resource-allocation and control. This theme forms the basis of the growing research field known as ***market-based control*** (MBC).

In brief, the aim of MBC systems is that groups of software agents or 'traders' interact within a market-like framework. In general, the influence from market economics is implemented by dividing the scarce resources into units of 'commodity' and then providing a 'currency' that is exchanged when agents buy or sell the commodity. Some agents act as 'producers' or 'sellers' of the commodity (e.g., an agent may be assigned to a node or link in a telecommunications network, charging for use of that resource), while others act as 'consumers' or 'buyers' (e.g., an agent may be assigned to a data-packet on a network, spending currency in order to route the packet from its source to its destination). In principle, when supply is greater than demand, the price of the commodity will fall; and when demand exceeds supply, the price rises. The aim then is that prices rise and fall, dynamically matching the quantity demanded to the quantity supplied, while either or both of these quantities also vary dynamically.

The classical theoretical picture of price-equilibration dictates that the number of trading agents (buyers and sellers) in the market is practically infinite, or very large at least. Yet, in a series of experiments commencing in the late 1950's, Smith (e.g. [28]) demonstrated that markets consisting of surprisingly small numbers of human traders could rapidly converge on the theoretical equilibrium price. Smith demonstrated that stable equilibria could be reached with fewer than twenty traders. In Smith's experiments, human subjects were told to be either buyers or sellers. Each seller was given a number of units of an arbitrary commodity to sell and each buyer was given the right to buy some units, and some currency. For each unit, each trader was given a *limit price* that was private (known only to that trader). A buyer couldn't pay more than her limit price for a unit of the commodity, and a seller couldn't sell a unit for less than her limit price. Typically, different traders had different limit prices: the distribution of limit prices determined the market supply and demand curves. In each experiment, time was divided into discrete periods of 5 to 10 consecutive sessions known as 'days'. At the start of each day, the rights to buy or sell units of commodity were distributed between the subjects. Each day came to a close either when no more traders were willing or able to trade, or when a pre-set time-limit expired (typically a 'day' lasted 5–10 minutes). During each day, the traders operated within a specific market structure: in many of Smith's experiments, the CDA was used, but he also experimented with 'retail' markets (where only sellers quote prices). In the early experiments, the traders communicated orally, but Smith subsequently developed methods where the traders communicated with each other via a network of computer terminals. Smith's work helped establish the field now known as *Experimental Economics.* The intention of the work described here is to develop intelligent trading agents with the bargaining capabilities necessary for small groups of traders to show price-equilibration in market scenarios similar to those Smith studied with human subjects.

The speed or stability of equilibration in a market can be affected by altering its structure, and markets can be structured in many ways. For example, in so-called *English-auction* markets, sellers remain silent while the buyers quote increasing bid-prices: bid-prices ascend until only one buyer remains, at which point a transaction occurs. At the center of this process is an *auctioneer*, without whom the market cannot operate.

Yet real markets operate efficiently without centralized auctioneers. The fastest and most efficient market structure is the *continuous double-auction* (CDA) market. In a CDA market, a seller can quote an *offer-price* at any time, and a buyer can quote a *bid-price* at any time. A transaction occurs whenever one trader accepts an offer or a bid quoted by another trader. At any time, any trader may quote a new price that supersedes that trader's previously-quoted price. The CDA is an attractive market structure: it is fundamentally asynchronous, and examples of distributed CDAs pervade real markets: for example, the CDA is the basis of trading at many international financial derivatives exchanges.

The intention in MBC is that, by building multi-agent systems based on free-market economies, the control systems exhibit the same decentralization, robustness, and capacity for self-organization as real economies.

Superficially, it may appear that little work in MBC is relevant to AI or robotics: examples from the MBC literature discussed later in this paper include allocating bandwidth in telecommunications networks, allocating memory in a computer operating system, controlling an air-conditioning system, distributing pollution controls, and job-shop scheduling. However, this superficial dissimilarity should not be allowed to mask some deep and important issues that are common to MBC systems, AI systems, and robotics. There are two main points to note:

First, the operating environments of trading agents in a decentralized MBC system are similar to those of robot agents in a collective behavior system. In both cases, the agents interact with a complex, dynamically varying environment that includes other agents. For a MBC trading agent, there are issue of co-adaptation, co-operation, and competition; and the value of an action has to be judged in terms of the likely response from the other agents.

Second, when appropriately configured, MBC systems could be used in AI applications including robotics. As more complex robots are developed, issues such as 'attention' (i.e., where next to point directional sensors such as active-vision heads, or which mission goal or subgoal to satisfy next), can be cast in MBC terms, where different processes compete via a market to consume resources that are scarce within the robot. Also, as the sensory bandwidth of robots increases, there is a growing need for specialized operating-systems capable of rapid transport and processing of frame-rate color video data: MBC may offer new solutions to this problem.

As with robotics and AI, the development cycle of MBC systems is often heavily reliant on extensive simulation studies in order to thoroughly debug the system before final implementation in the target environment.

We present here a critical review of a representative selection of MBC systems, concentrating on the degree to which they are truly decentralized. We argue that the MBC systems reviewed here are either implicitly reliant on centralized knowledge, or are not fully automatic insofar as they require human operators (who tend to be expensive and/or error-prone). Thus, we identify a central problem in creating truly decentralized MBC systems: the need for *bargaining behaviors* in the system's agents.

3 Market-Based Control: A Critical Review

To reiterate, two primary reasons for adopting a MBC approach are the prospects of *automation* and *decentralization* of resource-allocation or control processes. For the process to be automatic, it should devolve power from human operators: once the system is operational, human input should be reduced to a minimum. And for the process to be decentralized, there should be minimal reliance on central control mechanisms, processes, or databases (e.g., models of the entire network). Ideally, the system should not rely on the operation of any single critical component or sub-system. In much the same way that the national market for bread does not collapse when one baker goes bankrupt, so the failure of any one trading agent in a MBC system should result in only a minor impairment (if any) to the overall behavior of the system, rather than a total breakdown.

Yet, to the best of our knowledge, no current MBC systems are both automatic and decentralized in the senses just described. In the applications published in the literature, there is a reliance either on centralized 'auctioneer' processes or on human intervention. In the case where a centralized auctioneer process is used, in addition to the brittleness caused by failure of the central process, there are the issues of imposing synchrony on a fundamentally asynchronous system, and the costs of maintaining the central auctioneer's database (i.e., its 'knowledge' of how many traders there are, their interconnections, etc.).

Moreover, in the few truly peer-to-peer decentralized systems discussed in the literature, ideas from economics are used as a weak metaphor and there is nothing that approximates to a currency or price mechanism. Yet, surely, for the influence from economics to be more than a very weak metaphor, there should be a meaningful price mechanism: a currency should be available for expressing relative utility, indifference, substitution between commodities, and so on. For example, Malone et al [18] report on *Enterprise*, a decentralized system for task scheduling in distributed computing systems, where 'bids' indicate estimated completion times for tasks: this is sufficient for 'transactions' where clients select servers and vice versa, but there is no equivalent of competitive price-bargaining in this system. Also, the recent work of Epstein and Axtell [11] on "artificial societies" has attracted much attention but there is no money or price mechanism in their models [11, p.101], and so their work is of little relevance here. Furthermore, Cliff and Bruten [6, 7] discuss the lack of relevant work in experimental economics, and Cliff [6] notes the lack of related work in "biologically inspired" computing such as in "artificial life" research.

Despite the attractions of MBC, it is important to note that there are a number of theoretical studies that raise difficult issues. In particular, there are indications that the dynamics of some decentralized markets populated by simple traders acting purely to serve their own self-interest may converge to stable but highly sub-optimal equilibria (e.g., [14]), or may exhibit complex chaotic and hyperchaotic dynamics, (e.g., [13, 29]). The extent to which such theoretical models are applicable to real market systems is a matter for current empirical research; and the degree to which these results apply to MBC systems is also not clear, because all MBC systems with which we are familiar exclude one or more fundamental features of real decentralized markets.

Early methodological arguments for the field now known as MBC can be found in a range of disparate publications. It is beyond the scope of this paper to provide a full historical review of the field: in addition to the *Enterprise* project discussed above, which was in turn inspired by earlier ideas in *contract nets* (e.g. [9]), the theme of using ideas from microeconomics in controlling distributed systems was explored in depth by Miller and Drexler in three seminal papers [22, 23, 10].

In the remainder of this section, we briefly discuss five recently-published works in MBC, to demonstrate their centralization or reliance on human operators. All five come from a book edited by Clearwater [4]; which is the first-ever published collection that deals explicitly with MBC.

3.1 Network Bandwidth

Miller et al [24] discuss a system for automated auction of ATM (asynchronous transfer mode) telecommunications network bandwidth. The intention is that bandwidth (a scarce resource) is traded as a commodity by a community of software agents. In times of low network usage, high-bandwidth network connections may have a low 'price': as usage increases, so the demand for bandwidth increases and the price of the resource rises. Once demand has increased, users of the system have to make a simple decision between maintaining the previous level of expenditure (and consequently accepting reduced quality of service) or maintaining the prior quality of service (at a higher price). The system is sophisticated, and its components include mechanisms of banking and currency, bidding agents, auctioneers, delivery agents, and application and user interfaces. However, as is implied by Miller at al [24] and made explicit in a technical report subsequently published by the developers [1, p.21], the system relies on a centralized auctioneer process, known as NetAuctioneer. The developers acknowledge that this is unrealistic because "...it is a single centralized entity which requires a centralized global model of the network, rather than a distributed network of auctioneers each of whom have local knowledge only of parts of the net." [1, p.21].

The need for a global model of the network in NetAuctioneer may incur high maintenance costs in large or complex networks: presumably, any change to the network structure (including sudden failure of nodes or links) has to be reflected in the model for the auctioneer process to operate successfully. Moreover, there is the manifest danger that the entire bandwidth allocation system will collapse if the machine running NetAuctioneer fails.

3.2 Memory Allocation

Harty and Cheriton [16] describe the application of a market approach to memory allocation in a computer operating system. In their system, there is no increase in the price of memory in response to high demand. Rather, they use a tiered pricing system that allows the requestor to indicate the urgency or priority of a request [16, p.152]. When demand exceeds supply, the memory allocation scheme gives initial priority to those applications which have sufficient money to pay for their requested memory: when there is more than one such application, "...the allocation scheme gives priority to those applications that request an amount of space-time that is less than or equal to their 'fair share'." [16, p.153]. Again, this notion of 'fair share' requires a *global* view of the system: the 'fair share' of any one application can only be determined by reference to the needs and expenditure of *all* other applications.

3.3 Air Conditioning

Clearwater et al [5] demonstrate the use of market-based techniques for control of air-conditioning ventilation and temperature in 53 offices within one building. Software agents representing individual temperature controllers bid to buy or sell conditioned air. Unlike conventional building energy management systems, this system can take account of the interactions and connectivity between offices, and results indicate that the market-based system gives better distribution of temperatures and uses fewer resources. Again, a central computerized auctioneer process is employed: "A central auctioneer collects the bids and computes the supply and demand curves and sets the [equilibrium] price for the auction. Agents whose bids were not too high or low have their trade consummated. ... All other traders must wait for another auction to attempt having their bid consummated." [5, pp.256–257].

3.4 Pollution Regulation

Marron and Bartels [19] recount their experiences in developing computer-assisted auctions for the allocation of tradable pollution permits. A pollution permit gives a firm the right to emit some specified amount of pollution over a specific period of time. By giving firms an initial allocation of permits and then allowing them to trade them in an active free market, each firm acts as a participant in a distributed decision-making process. If the cost of reducing pollution by some amount is less than the market-price of a permit to produce that amount of pollution, it is profitable for a firm to reduce its pollution output and sell the corresponding permit. Firms that are unable to reduce pollution can buy extra permits from seller-firms, to avoid punitive financial penalties imposed on firms that produce pollution without a permit. The process is distributed in the sense that there is a reduced emphasis on government regulations that specify detailed standards covering equipment, operations, and so on [19, p.275]. However, once more, a centralized computerized auction mechanism is employed [19, p.283].

3.5 Job-Shop Scheduling

Finally, Baker [2] reports on the development of a fully distributed computer architecture for factory control, based on a market-driven contract net. In this system, a network connects software agents that each control one or more aspects of a manufacturing system, such as particular machines, inventory storage, material-handling, and so on. Also connected to the network are 'sales' agents that allow (human) customers to request a product.

Although there are agents within the contract-net, there is no bargaining, or even a centralized auctioneer: the internal agents calculate 'bids' for various aspects of the manufacturing process involved in satisfying a customer's request, but the decision of *which* bid is acted upon is made by the customer. The sales agent handling a customer's request combines information received from other agents and presents it to the customer. The information given to the customer summarizes many possible ways of satisfying the request: it is a multi-dimensional item of data, which is illustrated [2, p.190] as a 3-d surface relating lot-size (number of units), delivery-time, and cost per unit. It is then the responsibility of the customer to choose a combination of lot-size and delivery-time that gives an acceptable unit cost.

Thus, in Baker's system, the agents contribute to the size-time-cost surface of the final bid in a distributed fashion, but there is no place for an auction mechanism or any bargaining behaviors. The system presents all options to the human user, who is then responsible for deciding which, if any, of the options will be chosen. Thus, any market is *external* to the contract-net in Baker's system.

3.6 Summary

From this brief review, it is clear that although decentralization and self-organization are strong motivating factors for the development of MBC systems, the applications reviewed above are all lacking in some respect: centralized auction processes were used by the first four systems [24, 16, 5, 19]; while Baker's work [2], although inspired by real markets, employs no auction mechanism at all.

Thus, despite the promise of MBC, automatic decentralized systems are yet to be constructed. To do so, there is a need for computational 'bargaining' mechanisms that allow a software agent to decide what price to agree on for a transaction. Clearly, the price quoted by an agent (either a seller or a buyer) is likely to be influenced by the prices quoted by other agents in the market. In the next section, we summarize our ongoing work aimed at developing agents capable of such market-based behavior.

4 ZIP Trading Agents

The emphasis in our work is on creating *simple* autonomous software agents for bargaining in market-based environments. This emphasis on simplicity comes not only from a desire for computational efficiency (important if hundreds or thousands of such agents are active on a network), but also in a speculative attempt at sketching the minimum mechanistic complexity necessary and sufficient for explaining human bargaining behaviors in specific market environments. The potential use of such bargaining agents is not limited to MBC applications. There are at least two other significant application areas in which autonomous software agents with bargaining abilities could be profitably employed: internet commerce, and economic modeling. These issues are discussed further by Cliff [6].

Although it may seem intuitively obvious that some form of 'intelligence' or adaptation is necessary in bargaining agents, Gode and Sunder [15] presented influential results that appear to indicate that *zero-intelligence* agents can exhibit human-like behavior in CDA markets. Gode and Sunder's zero-intelligence trading agents simply generated random prices for bids or offers, subject to the constraint that they could not enter into loss-making deals. However, Cliff and Bruten [6, 8] demonstrated that

Gode and Sunder's result only holds in very specific circumstances and that, in general, some 'intelligence' in the form of adaptivity or sensitivity to previous and current events in the market is necessary. Hence, we give our trading agents adaptive capabilities by employing elementary machine-learning techniques. Because our agents are intended to have minimal intelligence, but not zero intelligence, they are referred to as "ZIP'" traders: ZIP is an acronym for "zero-intelligence-plus".

Space restrictions prevent us from presenting a full discussion of the rationale for the current design of ZIP trader agents, and from presenting full results. The intention here is to briefly summarize key aspects of the design and present some illustrative results. Cliff [6] gives a complete discussion of the design, shows results from many experiments in different styles of market environment, and includes all the C source-code for the system. A recent thesis by van Montfort [30] discusses experiences in using our ZIP traders in spatially distributed markets where there may be potentially hundreds or thousands of traders.

Each ZIP trader operates by maintaining a *profit margin* that it uses for calculating the price it 'quotes' (offers or bids) in the market: the profit margin determines the difference between the price the agent quotes and the *limit price* for the commodity the agent is trading. For agents designated as sellers (i.e., resource-producers), the limit price is the price below which they may not sell a unit of the commodity. For agents designated as buyers (i.e., resource-consumers), the limit price is the price above which they may not buy a unit of the commodity.

The 'aim' of each ZIP agent is to maximize profit generated by trading in the market. If an agent's profit margin is set too low, it will miss out on potential profit when it makes a transaction with another agent, so all agents are constantly trying to increase their profit margins. But if an agent sets its profit margin too high, it may miss the opportunity to make transactions with other agents, because the price it offers is less attractive than the prices offered by competing agents. Clearly, what it means for the profit margin to be "too high" or "too low" is dependent on the context of the market conditions, and varies dynamically. Thus, the problem of designing a trading agent can be considered as a combination of two issues: the *qualitative* issue of deciding *when* to increase or decrease the profit margin, and the *quantitative* issue of deciding *by how much* the margin should be altered.

For reasons discussed in detail by Cliff [6], each ZIP trader makes the qualitative decision of when to alter its margin on the basis of four factors. The first factor is whether the agent is *active* in the market: agents are active until they have sold or bought their full entitlement of units of the commodity. The remaining three factors concern the last quote by any agent in the market: we refer to this as Q. Each ZIP trader notes whether Q was an offer or a bid, whether Q was accepted (i.e., led to a transaction) or rejected (ignored by the traders in the market), and whether Q's price, $q(t)$, is greater than or less than the price the ZIP trader would currently quote. We refer to the price a ZIP trader i would quote at time t as that trader's *quote-price*, $p_i(t)$, which is calculated from i's limit price $\lambda_{i,j}$ (for i's jth unit of commodity) and i's current profit margin $\mu_i(t)$ using $p_i(t) = \lambda_{i,j}(1 + \mu_i(t))$.

A ZIP seller raises its profit margin whenever Q was accepted and $p_i(t) \leq q(t)$. It lowers its margin only if it is still active and Q was an offer with $p_i(t) \geq q(t)$, or if Q was a bid that was accepted and $p_i(t) \geq q(t)$. Similarly, a ZIP buyer raises its profit margin whenever Q was accepted and $p_i(t) \geq q(t)$, and it lowers its margin when it is active and either Q was a rejected bid with $p_i(t) \leq q(t)$ or Q was an accepted offer with $p_i(t) \leq q(t)$.

The quantitative issue of by how much the profit margin $\mu_i(t)$ should be altered is addressed by using a simple machine-learning algorithm. Specifically, the learning rule we use is *Widrow-Hoff with momentum*, which also underlies back-propagation learning in neural networks [26]. Briefly, this adjusts the actual output of a system toward some *target* output value, at a speed determined by a learning rate β, and with a simple 'memory' or 'momentum' parameter γ. In each ZIP trader the target value $\tau_i(t)$ is given by a stochastic perturbation of $q(t)$, and each trader i uses this in combination with β_i and γ_i to adjust its profit-margin $\mu_i(t)$. The profit-margin update rule is $\mu_i(t+1) = (p_i(t) + \Gamma_i(t))/\lambda_{i,j} - 1$ where $\Gamma_i(t) = \gamma_i\Gamma_i(t-1) + (1-\gamma_i)\beta_i(\tau_i(t) - p_i(t))$, and $\Gamma_i(0) = 0 : \forall i$. For further details of how learning is implemented in ZIP traders, see Cliff and Bruten [6, 7].

To demonstrate the effectiveness of this simple strategy, we present some illustrative results from 22 ZIP traders interacting in a CDA market. In these experiments, we used 11 buyers and 11 sellers, each with the right to engage in one transaction. The limit prices for both the buyers and the sellers ranged from \$0.75 to \$3.25 in steps of \$0.25: the supply and demand curves induced by this distribution of limit prices intersect at an equilibrium price of \$2.00. We start the traders with random initial profit margins, and record the transaction-prices over ten trading sessions or 'days'. Then, at the end of the tenth day, we add a 'shock' increase in demand. We do this using the method developed by Smith [28]: in some of his experiments, at the end of a day a new set of limit prices was distributed to the buyers, sellers, or both. Typically, the human traders would adapt, converging to the new market equilibrium values. This rapid, robust, and decentralized adaptation is one of the attractions of using the CDA as a market organization. Thus, it is important to explore the behavior of ZIP traders when supply or demand alter (either increase or decrease): for ZIP traders to be of genuine use in applications of market-based control or internet-based commerce, they should exhibit smooth and fast convergence to the new equilibrium that results from shifts in supply or demand.

Figure 1A shows a transaction-price time-series from one experiment. From the random initial profit-margins, the ZIP traders rapidly self-organize to give transaction prices that approach and stabilize at the predicted equilibrium of $2.00. At the end of Day 10, an increase in demand is imposed: the demand curve is shifted upwards by adding $0.50 to each buyer's limit price (the equilibrium price increases to $2.25), and the experiment continues for another five days. After this "shock change" in the market, the traders rapidly adapt to stable trading at the new equilibrium price. Figure 1B shows the average results from fifty such experiments. Similar results for increases in supply are shown by Cliff [6]. These figures clearly demonstrate that ZIP-trader CDA markets can self-organize, are capable of rapidly adapting to new equilibrium values resulting from changes in supply or demand, and are well-damped in the sense that a shock-change in supply or demand does not induce severe transients before the system settles. Thus, in experimental markets such as these, the results from ZIP traders are very similar to those from Smith's [28] human subjects: a point explored in detail by Cliff [6].

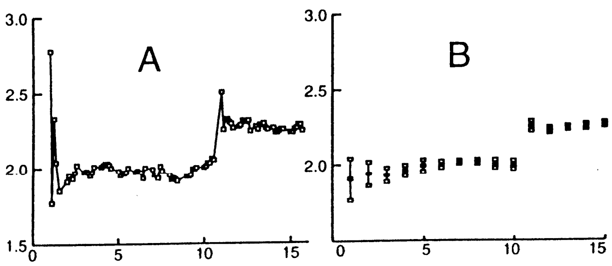

Figure 1: **A: Transaction-price time series for one experiment with a sudden increase in demand. The initial market has an equilibrium price of $2.00. After 10 trading sessions or 'days', demand is increased (raising the equilibrium price to $2.25) and the experiment continues for another 5 days. In each day, 5 or 6 transactions occur. B: Mean ZIP transaction prices per day,, averaged over 50 such increased-demand experiments. Solid points indicate the average mean transaction price for the day, with hollow points above and below indicating plus and minus one standard deviation.**

Cliff and Bruten [6, 7] show similar results illustrating ZIP traders operating successfully in CDA markets where Gode and Sunder's [15] zero-intelligence traders fail. These examples include markets where there are asymmetric supply and demand functions and imbalances between the number of buyers and sellers. The recent work of van Montfort [30] shows our ZIP traders operating successfully in more sophisticated markets: where there may be spatial structure or segmentation in the market (e.g., the traders are distributed over some space, and each trader can only transact with other traders in its local spatial neighborhood); and where agents are permitted to engage in arbitrage (i.e., selling units when the price is high, with the intention of buying them back when the price falls; or similarly buying units of a commodity for subsequent re-sale).

Thus, current indications are that our ZIP traders form a firm foundation for further work in decentralized MBC. In the current ZIP specification, the profit-margin update rule is discrete-time: our current research is directed at extending our approach to continuous-time domains, allowing for fully asynchronous operation. Evaluation and testing of all future developments requires extensive simulation experiments in a wide variety of market structures. The importance of this is highlighted by our critique [8] of Gode and Sunder's [15] zero-intelligence traders: we demonstrated problems with zero-intelligence traders by testing them in a wider variety of market simulations than had originally been used by Gode and Sunder. Thus, in MBC research such as this, it is worth investing effort in refining the simulation methods and techniques employed, to maximize computational efficiency, because many simulation experiments will be required.

5 Conclusion

The primary aims of this paper have been: (1) to emphasize that, for MBC systems to be truly distributed and self-organizing, it is necessary to develop computationally efficient mechanisms that endow autonomous 'trader' software agents with the ability to interact through bargaining behaviors; (2) to argue that current work in MBC often evades or avoids this issue; and (3) to demonstrate that the bargaining abilities of ZIP traders give desirable collective behavior in CDA markets.

Because CDA trading agents need to interact with a dynamic environment that includes other agents, where information may be noisy, uncertain, or delayed, and which is unforgiving of mistakes, there are similarities between the requirements of a trading agent and those of a mobile robot. Adaptation or learning will be necessary, possibly on multiple time-scales, to allow the agents to adjust to changes in market conditions: methods developed in artificial intelligence may offer a solution.

We presented a critique of recent MBC applications, demonstrating that the need for bargaining behaviors was avoided either by introducing a centralized auctioneer, by relying on a human operator, or by having no price mechanism. We also noted that, although Gode & Sunder [15] argue to the contrary, CDA trading agents *do* need to have some 'intelligence' or adaptive capability [8].

We described simple bargaining agents that adapt their profit margins to the prevailing market conditions, via elementary machine learning techniques. We demonstrated that groups of these "ZIP" software agents can, in a CDA market, exhibit collective behaviors comparable to those of human traders: transaction-prices rapidly approach the equilibrium price, are stable about the equilibrium, and respond well to shock changes. Simple bargaining mechanisms such as these could realistically be used in future MBC applications, allowing for truly distributed systems with the same decentralization, robustness, and self-organizing properties as real free-market economies.

Thus, the primary contributions of this paper are our critique of current MBC systems and our demonstration that the simple ZIP traders have the capabilities necessary to form a foundation on which future, truly decentralized, MBC systems can be built.

Our current research is directed at testing the capabilities of ZIP traders in more realistic and challenging environments, and extending them to the point where they can be used in resource allocation and control applications where time is continuous, trading is asynchronous, and information propagates with uncertainty and delays. At that point, the implementation of truly decentralized MBC systems in industrial-scale applications becomes a realistic prospect. It is likely that such systems will both draw from and find use in AI and robotics. Clearly, testing and evaluating such MBC systems will require intensive simulation experiments.

REFERENCES

[1] Agorics, Inc. Real-time video delivery with market-based resource allocation. Technical Report ADd004P, Agorics, Inc., Los Altos, CA, 1996.

[2] A. Baker. Metaphor or reality: A case study where agents bid with actual costs to schedule a factory. In [4] pp.184–223, 1996.

[3] R. Brooks & A. Flynn. Fast, cheap and out of control: A robot invasion of the solar system. *Jou. British Interplanetary Soc.*, 42(10):478–485, Oct. 1989.

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

[5] S. Clearwater, R. Costanza, M. Dixon, & B. Schroeder. Saving energy using market-based control. In [4], pp.253–273, 1996.

[6] D. Cliff. Minimal-intelligence agents for bargaining behaviors in market-based environments. Technical Report HP–97–91, Hewlett Packard Laboratories, Bristol, 1997.

[7] D. Cliff & J. Bruten. Less than human: Simple adaptive trading agents for CDA markets. Technical Report HP–97–155, Hewlett Packard Laboratories, Bristol, 1997.

[8] D. Cliff & J. Bruten. More than zero intelligence needed for continuous double-auction markets. Technical Report HP–97–157, Hewlett Packard Laboratories, Bristol, 1997.

[9] R. Davis & R. Smith. Negotiation as a metaphor for distributed problem solving. *Artificial Intelligence*, 20(1):63–109, 1983.

[10] K. Drexler & M. Miller. Incentive engineering for computational resource management. In [17], pp.231–266, 1988.

[11] J. Epstein & R. Axtell. *Growing Artificial Societies: Social Science from the Bottom Up.* MIT Press, 1996.

[12] S. Forrest. *Emergent Computation: Self-Organizing, Collective, and Cooperative Phenomena in Natural and Artificial Computing Networks.* MIT Press, 1991.

[13] N. Glance & B. Huberman. Dynamics with expectations. *Physics Letters A*, 165:432–440, 1992.

[14] N. Glance & B. Huberman. The dynamics of social dilemmas. *Sci. American*, 270(3):76–81, Mar. 1994.

[15] D. Gode & S. Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Jou. Political Economy*, 101(1):119–137, 1993.

[16] K. Harty & D. Cheriton. A market approach to operating system memory allocation. In [4], pp.126–155, 1996.

[17] B. Huberman, ed. *The Ecology of Computation.* Elsevier, 1988.

[18] T. Malone, R. Fikes, K. Grant, & M. Howard. Enterprise: A market-like task scheduler for distributed computing environments. In [17], pp.177–205, 1988.

[19] D. Marron & C. Bartels. The use of computer-assisted auctions for allocating tradeable pollution permits. In [4], pp.274–299, 1996.

[20] M. Matarić. Designing emergent behaviors: From local interactions to collective intelligence. In J.-A. Meyer, H. Roitblat, & S. Wilson, eds, *From Animals to Animats 2: Proc. Second Int. Conf. Simulation of Adaptive Behavior (SAB92)*, pp.432–441, MIT Press, 1992.

[21] J. McClelland & D. Rumelhart, eds. *Parallel Distributed Processing. Volume 2.* MIT Press, 1986.

[22] M. Miller & K. Drexler. Comparative ecology: A computational perspective. In [17], pp.51–76, 1988.

[23] M. Miller and K. Drexler. Markets and computation: Agoric open systems. In [17], pp.133–176, 1988.

[24] M. Miller, D. Krieger, N. Hardy, C. Hibbert, & E. Tribble. An automated auction in ATM network bandwidth. In [4], pp.96–125, 1996.

[25] C. Reynolds. Flocks, herds, and schools: A distributed behavioral model. *Computer Graphics*, 21(4):25–34, 1987.

[26] D. Rumelhart, G. Hinton, & R. Williams. Learning internal representations by error propagation. In [27], pp.318–362, 1986.

[27] D. Rumelhart & J. McClelland, editors. *Parallel Distributed Processing, Volume 1.* MIT Press, 1986.

[28] V. Smith. An experimental study of competitive market behavior. *Jou. Political Economy*, 70:111–137, 1962.

[29] J. Thomsen, E. Mosekilde, & J. Sterman. Hyperchaotic phenomena in dynamic decision making. *Jou. Sys. Analysis & Modeling Simulation (SAMS)*, 9:137–156, 1992.

[30] G. van Montfort. Economic agents for controlling complex systems. Master's thesis, Delft University of Technology Faculty of Technical Mathematics and Informatics, September 1997. Unpublished.