

PREFACE

Why Market-Based Control?

1. The Paradigm

This volume, for the first time, brings together the research from a wide range of fields, all sharing a common market-based conceptual framework. There are no strict definitions as to what constitutes “market-based,” or “auction-like,” or “economically-inspired.” Rather, it is the use of certain features found in a market that we have found useful to describe or motivate our work. These features include decentralization, interacting agents, and some notion of a resource that needs to be allocated.

Market-Based Control is a paradigm for controlling complex systems that would otherwise be very difficult to control, maintain, or expand. The purpose of this volume is to illustrate the utility of market-based control through a series of chapters focusing on different applications. Readers interested in control or management of complex systems, such as those in electrical, control and industrial engineering, operations research, information science, computer science, physics, economics, business administration, and management science, will find topics of interest.

2. The Phenomenon

A very abstract definition of a market is a system with locally interacting components that achieve some overall coherent global behavior. The fascinating aspect of a market is that through the simple interactions of trading, i.e. buying and selling, among individual agents a desirable global effect can be achieved, such as stable prices or fair allocation of resources. Markets don't guarantee an optimal solution but they often achieve satisfactory results. Their ability to facilitate resource allocation with very little information, i.e. price, makes them an attractive solution for many complex problems.

People have used markets for thousands of years to get things done. In human markets, shopping centers get built, new products are designed, all

without a global controller or overseer. The failure of centralized control of economies as demonstrated by the ongoing conversion to market-based economies underlies the difficulty of controlling a complex system through a central means.

All complex systems share common dynamical behaviors. These possible behaviors are: stability, transient instability, persistent oscillation, deterministic chaos, and random motion, which in turn are characterized by parameters of the system. Depending on the values of the parameters the system will be in one of these states. The parameters characterize the communication paths between the entities in the system, the efficiency of the communication, the time delays for communications, the accuracy of the response and the overall response of the entities, and the expectations of the agents. Other characterizations also exist. In a market-based system the “control” emerges from the individual goals of the agents rather than having a central goal imposed from above. Unfortunately, it can be a very difficult problem to know the parameters in the first place, to predict ahead of time what values the parameters should have for the system to be in a particular behavior, or even how to measure the parameters in a real situation.

The predictability issue is the main problem with a central control. While a “good enough” control may be achieved from time to time with heroic efforts, such a system is quite difficult to maintain and expand because the centralized authority has a necessarily rigid and therefore brittle structure. In the real world no omniscient controllers exist.

Contrast the centralized controller with a market-based system where there is no need for any of the agents in the system to know all the parameters of the system in order for the overall system to function smoothly. A market-based system necessarily is easy to expand and maintain since there is no need for a central coordinator.

We usually don’t think of markets as “controlling” anything, but in a sense they do. Markets determine whether a particular product or service is desired and what the level of that desire is. The market decides the value of products and services in a way that is both extremely complicated yet intuitive enough to be used by nearly everyone in everyday life. Markets “control” our economy and therefore what we will pay for things and what we will be paid for our contributions.

Markets are often thought of as being “open” in that there is little explicit control over them. They can be biased in certain directions but no one

has direct control or knows exactly what the outcome of a particular action by any of the components will be.

Other markets, though, are not as open. Consider the market for sulfur dioxide pollution discussed in the last chapter of this volume. This is a market created by the U.S. government that determines the total amount of sulfur dioxide pollution. The goal of the market is to control pollution by allocating the “right to pollute” to companies that are willing to pay for it. Certain companies are granted the right to pollute a given amount but may sell those rights to another company. The pollution market contains the aspects of an open decentralized market (buying and selling pollution rights) and more centralized aspects (government control of total pollution and distribution of initial rights).

Market-based control often contains both centralized and decentralized aspects. Computers give us an easy way to achieve an “economy” that is both under our control and yet able to have the desirable decentralized aspects that we want, such as the ability to scale to much larger sizes with little change in the underlying program.

3. Issues of Acceptance

Market-based control seems best-suited to those applications involving complex resource allocation or load balancing. These include tasks like shop-floor scheduling as discussed in Chapters 8 and 9.

While simulations, case studies, and demonstrations may show the utility of market-based control, there is still the very real problem of getting the paradigm accepted and into general use. Often market-based systems are in competition with monolithic centralized systems. The owners of these systems are frequently reluctant to give up what they perceive is a degree of control even though it may be beneficial to do so in terms of performance. This is somewhat akin to the problem of converting formerly centralized economies to decentralized ones, not an easy task.

3.1. It's Not a “Real” Market

One common criticism about artificial markets is that they are not exactly like financial markets. For example, in the auction-based control system to control conditioned air flow in a building discussed in Chapter 10, “funny money” was

used in the auctions and could not be saved from one auction to the next. In general one may make the “money” used by the agents to be anything one likes. Money can be burned at any time and hoarding may or may not be allowed.

However, these criticisms overlook the fact that these markets *could* be operated akin to financial markets. The reason not to operate artificial markets as realistically as possible is in part due to the complexity, but also because we need not incorporate undesirable market behavior if we “own” the market. We can choose the best features of a market system without some of its undesirable ones, e.g. cheating.

Other market-based systems, such as those discussed in Chapter 5, *do* allow for more complex behaviors because the agents are not owned by the same entity and do have selfish interests that cannot simply be programmed away.

3.2. Security

One issue in the computer market paradigm is that of security. In other words, how do we prevent unscrupulous agents from taking advantage of the system? In human markets the answer is to have some kind of regulatory agents, e.g. police. However, in computer markets this need not be the case. Consider a computer market within a single organization that schedules shop floor activity, as discussed in Chapters 8 and 9. Since the supervisor owns the computer agents there is no reason to include unscrupulous behaviors in the agents’ characteristics. The difference between human markets and computer markets is that in the former we have no control over what the human agents might do, whereas in the latter we can program whatever behavior we like into the computer agents and avoid the complications that would occur with greedy or selfish agents. In other words, the computer market will behave much more like the idealized scenarios of economists. In a market-based system where we don’t own all the agents, such as described in Chapters 5 and 11, there is the issue of security that must be addressed.

3.3. Explanation

Another issue facing the widespread acceptance of market-based control is, why does it work? This is the same perplexing question that every economist

must face when asked why a market-based system works. The answer must necessarily be an incomplete one.

One often hears financial market pundits uttering “explanations” of why the markets moved in a certain way on a certain day, “the dollar fell,” “there was a freeze in Florida,” “the president is optimistic about the future,” and so on. But, do these explanations truly tell what drove individual trades among agents? Indeed Adam Smith spoke of an “invisible hand” that guides markets — hardly a rigorous explanation, but as fitting today as it was two hundred years ago.

Artificial intelligence researchers have also probed into the problem of providing explanations, but their answers are only partly satisfactory. For example, expert systems encode the knowledge of an expert at some task, often in the form of IF-THEN type rules. If the user has a question about why the expert system came up with a particular answer, the response is a traceback of all the rules that fired. This is the “explanation.” For very complicated expert systems the interaction among rules can be very complicated and the resulting traceback quite confusing.

Other explanatory knowledge may be in the form of a program where the explanation is the algorithm itself. On a more detailed level one can look at the particular parameter values that caused the program to give its answer. Similarly, in a financial market we can ask, who traded what and when? But these explanations are so detailed that the forest is lost for the trees. So, while many systems attempt to provide explanations, their answers are often times of questionable utility. With a market-based system we gain performance at the expense of not knowing *exactly* what will happen in a given situation.

4. Outline of the Book

The chapters in this volume are intended to present some examples of the many possible applications of market-based control. By illustrating these applications we can begin to answer the question “Why Market-Based Control?” I have deliberately avoided including papers relying on formal proof so that the possibilities of market-based control are more clearly accessible to a wider audience. Most of the chapters involve simulations of market-based control — which means there are still plenty of opportunities for real implementations! A few of the chapters show market-based control in real applications, with notable improvements over conventional controls. Whenever

possible there are direct comparisons of the market-based control with a competing paradigm. Strengths and weaknesses of both paradigms are discussed. There are several topics in the chapters which makes for a natural organization from the more abstract to the more concrete.

The first chapter, by Steiglitz *et al.*, shows results from simulations of a market-based system and the various dynamical situations that can exist. He finds, for example, that adding speculators can help to stabilize a system resulting in more effective resource allocation.

The second chapter, by Gagliano and Mitchem, discusses simulations in a computer network.

The third chapter, by Kuwabara *et al.*, shows simulations of market-based control applied to a communication network. Mechanisms to reduce oscillatory behavior due to time delays and abrupt changes are discussed.

Chapter 4, by Wellman, discusses some the lessons and pitfalls when trying to program a market-based system. He includes some lessons on scaling behaviors.

Chapter 5, by Miller *et al.*, discusses an automated auction for allocation of network bandwidth. In particular, they discuss automating a ATM network bandwidth allocation.

Chapter 6, by Harty and Cheriton, gives simulation results for a market approach to operating system memory allocation. Their approach utilizes memory money to lease a certain amount of memory for a specified amount of time.

Chapter 7, by Ferguson *et al.*, shows market-based control applied to computer systems. They show how competitive economic models provide algorithms and tools for allocating resources in distributed computer systems.

Chapter 8, by Baker, shows a case study for factory scheduling. The simulation study is based on an actual factory and prices using a Market-Driven Contract Net.

Chapter 9, by Tilley, is another factory simulation, this time to machining task allocation. He presents a model for the evaluation of an auction-based scheduling system that functions by passing messages over a network.

Chapter 10, by Clearwater *et al.*, shows a working implementation of an auction-based system for distributing air more equitably in a building. The "thermal market" does a better job at providing comfort while at the same time saving energy.

The book concludes with a chapter by Marron and Bartels of a “smart” auction to aid in the auctioning of air pollution rights. They find that the computer auction provides better results than from “eye-balling.”

Casting such a wide net over such a seemingly diverse set of fields will hopefully yield an eclectic collection of work whose applicability and popularity will grow as the world becomes more complicated and interconnected. I hope that readers will find these chapters of use in their own applications.

Scott H. Clearwater
Palo Alto, California
September, 1995