Economics-Based Approaches in Multi-Agent Systems and Distributed Information Networking*

Jonathan Shapiro

Department of Computer Science
University of Massachusetts at Amherst
jshapiro@cs.umass.edu

Abstract

The field of economics has received increasing attention from researchers in both artificial intelligence and distributed systems. In large-scale open systems, the competing interests of autonomous users must be explicitly considered in system designs to ensure global performance objectives are satisfied. At the same time, reasoning about the tradeoffs involved in the complex resource allocation scenarios typical of such systems is an interesting problem in AI. Economics-based approaches are proving useful to researchers in both these areas and represent an exciting area of interdisciplinary collaboration.

1 Introduction

As networked computer systems have grown in scale and openness, their role as a public infrastructure has prompted growing interest among computer scientists in social aspects of computing. The networked environment is characterized by a high degree of autonomy and decentralized control, and is host to various actors pursuing diverse and possibly conflicting agendas. System designers confronted with the need to balance the individual interests of the user with system-wide goals are finding that this familiar dilemma takes on a social engineering flavor in open systems. It is perhaps not surprising that some researchers have looked to the field of economics for tools to help analyze and provide a foundation for the design of such systems. Interestingly, this attention to economic principles has been growing in both the Artificial Intelligence community and in the Systems community. In Multi-Agent Systems (MAS), there is a growing line of research to explore interactions among agents who are not necessarily cooperative, but instead are endowed with individual interests. In Networking and Distributed Systems, systemic problems such as congestion control in the Internet are strikingly similar to classic problems of resource allocation studied in the social sciences and there is a hope that economics-inspired mechanisms may be applicable in such settings. Most important, however, is the fact that an increasing portion of the real economic life of our society is being conducted in an automated fashion on large open systems, providing a requirements-driven pressure on the design of networks and distributed systems, as well as a challenging application domain for multi-agent systems.

In this research statement, I will argue that the application of economic principles in computer science represents an exciting area of crossover between the AI and systems communities. I will briefly mention some of the interesting applications of economic tools such as the theory of markets and game theory in each sub-discipline and offer an example of a research problem that is interesting for both sub-disciplines.

2 Multi-Agent Systems

The field of multi-agent systems grew out work in distributed problem solving in AI in which agents were assumed to be working cooperatively to solve a problem which none could solve individually[3, 1]. As distributed AI approaches were applied to a greater range of inherently distributed problems, it was often the case that a decision that appeared locally optimal to a single agent did not lead to a globally optimal solution. This tradeoff between local and global optimality has been explored extensively in the cooperative setting[2], leading to the observation that in many environments, agents have real and legitimate differences as to what constitutes optimality. The field has thus broadened to include the study of interactions among self-interested agents, who are not necessarily concerned with global measures of utility, but who nevertheless are instrumental in finding the solution.

Many problems in multi-agent systems can be formulated generally in terms of resource allocation[13]. When agents do not share global objectives, there typically is no allocation that is optimal for all agents. It is interesting to examine these problems in an economic framework in which various well-studied mechanisms offer opportunities to find efficient allocations that preserve globally desirable outcomes. For example, work in automated negotiation[9] directed at enabling small numbers of agents to come to agreements is informed by game theory, auction theory, bargaining theory in economics. As the number of agents scales up, there is significant overhead associated with both the communication required for negotiation and the computation required for a rational agent to behave strategically. For resources that are well modeled as commodities, the theory of markets gives us mechanisms to alleviate both types of complexity, by reducing the amount of information that must be communicated to price data, as well as allowing the individual agents to adopt a "competitive" posture in which the need for complex

---

* This material is based upon work supported by the National Science Foundation under Grant No. IRI-9523419. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

strategic reasoning is reduced. Interesting research questions include how efficient, decentralized, scalable market mechanisms can be implemented in practice to realize these benefits [12]. Game theory is also relevant here for examining how social conventions may emerge among learning agents in the absence of direct communication [10].

3 Networks and Distributed Systems

Problems of of resource allocation arise in distributed systems at many time scales and levels of abstraction. At the systems level, system nodes vie for fundamental computational resources such as network bandwidth and persistent storage as well as more abstract commodities such as response time. At the level of distributed applications built on top of such systems, these low level resources are bundled together to form more complex resources. Traditional approaches to achieving system-wide performance goals in distributed systems, including most classical distributed algorithms assume that individual nodes are essentially cooperative within the limits of available bandwidth with the result that many distributed solutions developed for closed environments are easily exploited by malicious (or merely self-interested) nodes in an open environment.

The problem is that systems designers have historically considered the architectural implications of autonomy but not the social implications of divergent interests among autonomous domains. In an open system like the Internet, service providers typically would like to optimize measures like utilization, fairness, and stability. But such goals must be traded off against the quality of service demands of the individual consumers. Without design principles that consider the aggregate effect of self-interest, the resulting allocations may be sub-optimal by any measure. Numerous cases of congestion of Internet traffic, for instance, are examples the problem of the "Tragedy of the Commons". Well known in economics, this situation occurs when a public good is overtaxed by a group whose members lack any individual incentive to regulate their usage, despite a clear benefit of conservation to the aggregate. Divergent interests also appear among different service providers, who share an incentive to maintain full connectivity on the one hand, but have a competitive relationship among themselves on the other. It is important that binding agreements between providers reached through automated negotiations be, at the very least, verifiable, if not directly enforcable.

A number of research efforts in the distributed systems and networking communities have attempted to apply the tools economic theory offers designers for building very large scale open systems. The theory of markets can help develop decentralized mechanisms to achieve efficient allocation of resources without requiring consensus among nodes [5, 4, 7, 8]. Game theoretic models of networks are beginning to yield insights[6] about, for example, the effectiveness of congestion control mechanisms in the Internet.

4 Combined Approaches

As an example of how these approaches can be combined, consider a distributed information processing infrastructure for a large-scale electronic publishing enterprise. In this example, publishing requires the use of computational resources in a production environment where data are processed and assembled into information offerings prior to publication, and a publication platform–a cluster of web servers, perhaps–from which information is served to customers. This combined platform is used to produce and publish to a variety of content, perhaps several online information services, for example. Let us assume that new applications (i.e. new collections of processes devoted to producing additional information for publication) may be introduced by users at any time without direct administrative intervention.

The system in this example is more centralized and smaller in scale compared with the national information infrastructure. But even in this limited domain, the divergent interests of system stake holders illustrates the power of an economics-inspired approach. Even when the entire infrastructure is managed by a single entity, the divergent interests of application owners drive the need for efficient resource allocation. Each application generates tasks which consume system resources such as CPU time, disk storage and network bandwidth to execute. Individual application owner are chiefly concerned with the performance of their own applications and are under pressure to deploy additional applications. From the administrative organization's perspective, a central concern is that the system makes certain service guarantees under normal loads and that performance degrades gracefully under heavy loads. The administrators, however, are not in a position to judge what constitutes graceful degradation, since this is largely application dependent. Some application owners may prefer that only their most important applications function at normal performance levels when system load is high, while others may prefer an overall performance degradation across all of their applications.

Of course, without any incentive to conserve resources, most application owners would prefer to have all of their applications functioning at normal levels at all times. One way of providing this incentive [4] is to construct a market for computational resources and allocate a form of currency to each application owner. The total amount of currency in the market is controlled by the administrators to keep the system at a manageable load, but application owners spend their individual allocations as they see fit. Prices of resources reflect demand and a type of decentralized load balancing occurs. Whether such a market can function with manageable overheads on larger scale systems, however, is still an important research question.

If the system is large and complex and the applications running are of sufficient importance then there can be a significant incentive for the individual application owners to reason carefully about the tradeoffs associated with various allocations of their resources. Yet this reasoning must be performed quickly and adapt to a dynamically changing environment. It may only be possible to compute a near-optimal allocation when task deadlines are considered [11]. On the other hand, it may be possible to negotiate with other application owners when time permits to find a joint allocation that is mutually beneficial. Mechanisms that allow agents to hedge against future uncertainty such as leveled commitment protocols [9] and futures markets [14] can expand the space of possible deals to include agreements that improve the payoffs of all parties.

The need to automate the reasoning and negotiation over these tradeoffs is what makes this problem interesting from the multi-agent systems perspective. Developing market and negotiation mechanisms that can operate efficiently on the scale of the global telecommunications infrastructure is an important challenge for distributed systems research. While the complimentary nature of these trends in research is easy to identify, there is still more to be done in the way of integrating work in these two sub-disciplines. Some reasoning and negotiation techniques in multi-agent systems have high overheads, which, although possibly acceptable at the application level in distributed systems, may not be directly applicable as mechanisms for system-level resource allocation. A systems perspective is needed here to clearly identify the computational tradeoffs associated with these new economics-based approaches so that system designers can choose among them in a principled fashion. As these techniques become more widely implemented in distributed systems, there will be a growing need for individual computational agents representing the interests of individuals and organizations that can reason efficiently about the complex tradeoffs inherent in resource allocation for real problems ranging from bandwidth reservation to negotiation in distributed electronic commerce applications. Automating decision making in these areas offers an opportunity for intelligent systems to make a significant contribution in an exciting application domain.

References

[1] K. S. Decker, E. H. Durfee, and V. R. Lesser. Evaluating research in cooperative distributed problem solving. In L. Gasser and M. N. Huhns, editors, *Distributed Artificial Intelligence (Vol. II)*, pages 485–519. Kaufmann, San Mateo, CA, 1989.

[2] E. H. Durfee and V. R. Lesser. Using partial global plans to coordinate distributed problem solvers. In *Proc. of the 10th IJCAI*, pages 875–883, Milan, Italy, 1987.

[3] E. H. Durfee, V. R. Lesser, and D. D. Corkill. Cooperative distributed problem solving. In A. Barr, P. R. Cohen, and E. A. Feigenbaum, editors, *The Handbook of Artificial Intelligence (Volume IV)*, pages 83–148. Addison-Wesley, Reading, MA, 1989.

[4] D. F. Ferguson, C Nikolaou, J. Sairamesh, and Y. Yemini. Economic models for allocating resources in computer systems. In S.Clearwater, editor, *Market-Based Control: A Paradigm for Distributed Resource Allocation*. World Scientific Publishing Co., 1996.

[5] D.F. Ferguson, C. Nikolaou, and Y. Yemini. An economy for managing replicated data in autonomous decentralized systems. *Proc. Int. Symp. on Autonomous Decentralized Sys*, 1993.

[6] E. Friedman and S. Shenker. Learning and implementation on the internet. http://econ.rutgers.edu/home/friedman/research/decent.ps.

[7] R. A. Gagliano and P. A. Mitchem. Valuation of network computing resources. In S.Clearwater, editor, *Market-Based Control: A Paradigm for Distributed Resource Allocation*. World Scientific Publishing Co., 1996.

[8] J. F. Kurose and R. Simha. A microeconomic approach to optimal resource allocation in distributed computer systems. *IEEE Trans. on Computers*, 38, May 1989.

[9] T. Sandholm and V. R. Lesser:. Advantages of a leveled commitment contracting protocol. *AAAI/IAAI*, pages 38–45, 1996.

[10] Y. Shoham and M. Tennenholtz. On the emergence of social conventions: Modeling, analysis, and simulations. *Artificial Intelligence*, 94:139–166, 1997.

[11] T. Wagner, A. Garvey, and V. Lesser. Design-to-criteria scheduling: Managing complexity through goal-directed satisficing. *Proceedings of the AAAI-97 Workshop on Building Resource-Bounded Reasoning Systems*, 1997.

[12] M. P. Wellman. Market-oriented programming: Some early lessons. In S.Clearwater, editor, *Market-Based Control: A Paradigm for Distributed Resource Allocation*. World Scientific Publishing Co., 1996.

[13] M. P. Wellman and P. R. Wurman. Market-aware agents for a multiagent world. *Robotics and Autonomous Systems*, 24:115–125, 1998.

[14] H. Yamaki, M. P. Wellman, and T. Ishida. A market-based approach to allocating qos for multimedia applications. In *Proceedings: International Conference on Multi Agent Systms*, pages 385–392. 1996.