Agents, Self-Interest, and Electronic Markets

Lars Rasmusson and Sverker Janson

Swedish Institute of Computer Science
Box 1263, 164 29 Kista, Sweden

lra@sics.se, sverker@sics.se

Abstract

How free are our software agents to take the best possible care of our interests? How free can we make them? In what sense and to what extent do currently proposed mechanisms and agent behaviors consider self-interest? What current research addresses these issues? What needs to be done?

We address these issues in the context of electronic markets, such as consumer goods markets and (future, more fine-grained) markets for electric power or communication bandwidth.

Electronic markets and self-interested agents

The notion of a *self-interested software agent* is well established in the agent community. Although this notion of agenthood is widely applicable, it is perhaps most obviously useful in the context of electronic markets, where it coincides with the rational, utility maximizing, agent of microeconomic theory (see, e.g., [1]).

Through automated trading by software agents, we expect improvements in the quality of existing markets, such as consumer goods markets, and to reap the benefits of markets as effective instruments of resource allocation also in non-traditional domains, such as fine-grained markets for electric power and communication bandwidth. Agents have the capacity to consider more information, e.g., evaluate thousands of offers for a new car, and may also act in domains where we are disqualified due to speed requirements, e.g., buying the bandwidth we need packet by packet.

Activity in this field has surged the last few years, riding on the Internet bandwagon, but the idea of viewing parts of a distributed computing systems as economic agents has been around at least since the late 70s (see, e.g., [2]). At the time, *computational economies* were mainly regarded as mechanisms to achieve system global goals, i.e., a class of decentralized algorithms for resource allocation built into a system by top-down design.

In systems designed with the market based resource allocation paradigm, resources are allocated based on the current supply and demand for a set of resources. If the consumers have several alternative ways to satisfy their resource demand, the only important factor for them to consider is the price, provided that individual resource differences have been normalized away. It is conceptually easy to understand that if the agents choose to sell/buy resources where the prices are highest/lowest, resources will generally be allocated to where they are expected to be of most use. The paradigm has been proven applicable and efficient for a range of domains.

Work presented at conferences in 1998 includes

- Selling excess CPU time or other resources [3, 4, 5, 6]

- Constraining peak usage of networked resources [7, 8]
- Improving network utilization and economic efficiency [7, 9, 10]
- Differentiated network service levels [11,12]
- Work flow and task allocation [13, 14]
- Data replication and mirroring [15, 16, 17]
- Information searching/selling [18, 19]
- Technical infrastructure for trading agents [20, 21, 22, 23]

The most recent areas where the economic behavior of software agents is being studied are *information economics* and *network economics*. Information economics is concerned with pricing information goods and with the problems that arise when the production marginal cost is zero, i.e. when copying is essentially for free (see, e.g., [18, 19]). Information can be expected to be a commonly traded goods in software agent mediated commerce, since both trading and distribution can be handled by agents.

Network economics deals with pricing the access to a shared resource in which the marginal cost most of the time is zero, but where occasional usage peaks result in congestion and heavy degradation of the service (see, e.g., [24, 7]). As the congestion incurs more cost for the other agents than for the congesting agent, different kinds of economic control is expected to reduce peak usage.

Current research on agents in electronic markets is still to a large extent focused on idealized computational economies, or on the implementation of agents and market infrastructure without an accompanying analysis of their real-life properties. The agents of such systems have no true self-interest, are not competitive, do not speculate, but play the market game using a given strategy, working towards a common goal. We would like to encourage a shift of interest from well-behaved, easy-to-analyze computational economies, to open full-fledged electronic markets in which the participants are free to act in their own best interest.

In the following sections, we will discuss some recent results on the topic of self-interested agents in electronic markets, suggest issues that need further research, and finally offer some conclusions.

Emphasizing self-interest

Agents that are not self-interested are often made to negotiate to achieve a high global (or social) welfare. Such unselfish agents will give up their goals if someone else has a more important goal which is blocked by the unselfish agent. The idea that agents should behave in order to optimize the global welfare of the system has been popular in the distributed AI community, where a distributed system of agents solve a common task for which the system was specifically designed. All agents share the same goal, hence the high global welfare coincides with the agents' own return.

However, in open systems it is not possible for the participants to agree on which goal is more important than another as a unique global utility function does not generally exist. One therefore often uses budget constraints rather than the relative importance of goals to settle conflicts. In other words, the highest paying agent gets the goods. This appears to be working remarkably well as the one willing to pay a lot often can find good use for the resource, since paying too much tends to drain the budget and eventually makes the bidder go broke. (This argument assumes that the other bidders are rational and the market is sufficiently large.)

That agents will be self-interested is easy to see. A personal software travel agent should preserve the interests of its client. It should not negotiate to maximize the profit of the airlines, and similarly, a personal bargaining CD-agent should not act to maximize the profit of the record company. To preserve the individual goals in a negotiation, each negotiating party will rather have its own agent looking after its interests. In open infrastructures agents will necessarily be self-interested, since

- If people are free to create their own agents, we expect that they will incorporate self-interested behavior.
- If people must choose between using self-interested agents or agents optimizing the profit for someone else on their expense, we expect they will choose self-interested agents.
- If people only may use agents that are optimizing the outcome of someone else on their expense, we expect agents to be used only for tasks where one is even worse off without an agent.

Cooperative agents give away their resources to other agents when explicitly asked (see, e.g., [25]) and more implicitly in the constraint-based coordination platforms (see, e.g., [26]). Using cooperative agents in a self-interested negotiation context results in absurd behavior of the agents.

Imagine that you have bought a holiday trip, and your hotel is now full so no-one else can go. Another person contacts you and says he *really* wants to have your room. You would probably not just give up the room just because someone you don't know could use your room. However, in cooperative resource allocation, if another agent says it wants the resource *more* than you, the resource is given up. Some consider it "mean" of an agent to refuse to help other agents without compensation, but simple "nice" behavior means that your agent subsidizes the other holiday goers since they do not have to bid competitively. You also subsidize the travel-agency as it can charge more than the end-consumer is willing to pay, which may well decrease the total welfare, if such a thing exists.

If giving up your holiday was to actually increase the global welfare, the seller and other travelers should have been able to compensate you economically for the inconvenience. If they cannot do that, either the global welfare decreases or a global welfare metric does not exist, which means that one outcome is not necessarily better than the other. This line of reasoning has been criticized for not taking into account the agent's happiness or increased self-esteem from helping others [27]. However, it should be obvious that if *happiness* cannot be measured on an absolute scale it is meaningless to say that one outcome is better than another simply because one agent is more happy than another agent is sad.

We should not assume cooperative behavior from the agents, partly because it is irrational, and partly because agents will in general not agree on which is the best social outcome. Unselfish agents will make the system exploitable and vulnerable rather than efficient.

To seal or not to seal bids

A large portion of the work in economic mechanism design consists of game theoretic analyses of interaction or bidding protocols to find out when agents may benefit from speculative behavior. The agent community is eagerly adopting results from this work. *Incentive compatible* protocols promise to make life easy for the agent designer by guaranteeing that the best strategy is not to speculate, but to be honest, thereby freeing them from the burden of adding speculation intelligence, adaptation and recursive world models to the agent.

Occasionally, results are adopted without giving proper attention to the preceding assumptions. This happens occasionally for the *Vickrey auction*, where the winner only pays the price of the second best bid. One supposes that bids are sealed and that an agent has no knowledge of the bid distribution. The analysis shows that bidding too high only risks paying too much without increasing chances of winning at an acceptable price. Bidding too low only decreases chances of winning without lowering the price the agent must pay. Therefore bidding should be incentive compatible and agents should bid honestly.

The Vickrey auction is only guaranteed to work as designed when the bidders have no possibility to communicate on the side and have no information at all about each others' bids. In spite of that, the Vickrey auction is suggested for scenarios where these assumptions clearly can be violated, for instance in iterated auctions or auctions over the Internet [5]. In iterated auctions one can often infer clearing prices from previous prices. One may benefit from bidding too high if that reduces the profit for a competitor. A seller in a Vickrey auction can use the equivalent of shills (fake bidders) to increase the profit. With communication on the side, the buyers can collude against the seller by agreeing on bidding very low.

Vickrey auctions vs English auctions

Economic coordination for work-flow uses agents that buy and sell their clients' services (see, e.g., [28]). If the agents represent the individual workers, they are clearly representing self-interested individuals. This domain was explored in a paper about English auctions over multi-dimensional goods [13]. The complex issues of self-interest were avoided by assuming that it is illegal to resell the resources and that agents treat prices as exogenous.

Recent work by Sandholm [29] and by Vulkan and Jennings [13] has made it clear that the Vickrey auction has several drawbacks that makes it unsuitable for open agent systems. The fact that the Vickrey auction relies on secret private value bids makes it susceptible to collusion and unusable for common valued goods. Vulkan and Jennings suggest using the English auction since it uses open bids and proves that for private value goods the auction is incentive safe if prices are completely exogenous. Nothing is said about how the system behaves if agents are allowed to combine and resell goods in which case goods will have common values.

Ungar [30] also suggests the English auction, but for another reason. The Vickrey auction is considered too computationally costly or difficult to use, as it requires you to compute your true reserve price in advance. This is a somewhat controversial claim because since the Vickrey auction is incentive compatible it has been held as strength of the Vickrey auction that no strategic reasoning is required. The reason why computing the reserve price is difficult is not very thoroughly discussed. Rather this is inferred from the fact that most online auctions are using English auctions, and almost none use Vickrey auctions. One reason given is that bidding for several goods in parallel requires nested models of the other agents' preferences and knowledge, as one must predict the allocation in the other auctions to bid correctly in one auction. Iterated auctions like the English auctions, will simplify, although not eliminate, the problem, since in early phases of bidding, rough estimates may be sufficient, and as bidding proceeds, more information will be available.

Price-taking behavior

Network pricing is often analyzed in terms of price-taking consumers, whose demand for each resource will depend only on prices; increased prices reduces traffic [7, 8, 11, 33, 16]. The assumption that consumers are price-taking is a simplification since it does not take into account that competing providers will adjust prices to maintain or reroute traffic. This gives consumers a possibility to speculate and the price-taking assumption will not hold.

In computational economies, sellers are almost always modeled as setting their prices according to a predetermined function, increasing prices with demand, rather than maximizing profit [3, 4, 16, 17]. This results in a nice and smooth dynamic behavior, but the increasing margin-cost pricing strategy is not necessarily a rational choice for the resource providers, since the often exponential increase in marginal costs results in some of the resources not being sold. Lowering prices just enough to empty the stock of resources seems more rational. Kephart et al [18, 19] show that self-interested agents in information economies may generate price wars by undercutting.

The price-taking behavior assumption is also incompatible with the ability of consumers to form coalitions (consumer cooperatives) to cut prices. Collusion in multi-agent systems has been analyzed in terms of game theory (see, e.g., [31, 32]).

General equilibrium markets

In a general equilibrium auction bids are submitted as resource supply and demand functions to an auctioneer, who computes prices by solving the resulting system of equations, such that supply equals demand. If the market only clears once, there is little room for speculative behavior among the agents, as has been demonstrated recently [34]. However, if the market clears several times or if the bidding length is randomized, there is again room for time-based speculation, and it can be beneficent to bid aggressively initially or near market clearing time.

Clients might not want their agents to reveal their private valuations to an external auctioneer as that information can be exploited. In open equilibrium markets, agents inform the auctioneer about their demand and supply only when they receive specific price quotes. This opens up for communication on the side and thereby collusion of the auction. The auctioneer may have to restrict the bidding deliberation time to ensure that the market will clear eventually [13]. This introduces a time-based speculation race even in the price computation of the equilibrium market.

The idea of viewing markets as equilibrium markets is a theoretical simplification. It assumes that agents treat prices as exogenous, i.e., the agents' behavior does not affect the price. This is a very strong assumption, which is not guaranteed to hold when agents can form coalitions outside the market or whenever negotiation is possible.

Issues that require further attention

Speculation is necessary

Our vision of future electronic markets is that they will allow us to create agents that integrate simple resources and services into more complex products. This calls for agents that can act as enterprises, speculatively producing resources that they wish to sell. Such agents will often face the task of combining several resources into one, tailored and packaged for its client. To increase the number of services that are possible to produce, the agents must speculate in order to obtain benefits from mass production (of goods and information), etc.

In general equilibrium markets, the agents must express their bids as absolute utility functions, assuming gross substitutability [35]. This assumption is too strong for most real world applications, and incompatible with our vision, as it assumes that the utility of one resource does not depend on whether one holds another resource or not. If we allow reselling, we also allow for time-based speculation, and for instance many of the assumptions needed for Walrasian general equilibrium auctions fail.

If demand depends on the availability of other resources, if resources can be resold, or if equivalent resources can be obtained at a later time, the prevalent assumption that simple private value bids can be used does not hold. Most of the trading performed with agents will be performed at market price with both time based ("Should I buy now or later?") or value based ("Should I increase my bid or will the seller decrease his offer?").

Trading concurrently in several markets

One can not assume that all resources are traded at the same market as this would be an enormous bottle neck (and give too much power to the owner of the market). This complicates the process of combining resources. Agents may end up with too much resources (as in [3]). Unless these resources can be resold, they are lost. In a dynamic system, this will cause low resource utilization. Therefore, in large systems reselling may be necessary even though it causes speculation opportunities.

In multimarket scenarios, we can expect to find agents that act as speculators with the intention to exploit price inefficiencies, just as in the "real" world. This is not a bad thing as it adds liquidity to the market. (Increased liquidity generally simplifies the reasoning for the agents as it stabilizes prices.) However, frictionless electronic trading also poses a risk for speculation races.

Both anonymous and identity based trading

Anonymous trading in stock markets is considered to increase the liquidity, especially since large traders can participate incognito. Large traders, if identified, risk altering the prices in a for them negative way by trading.

Privacy, or integrity, is another reason for allowing anonymous trading. Today, sellers log information about their customers to present more pertinent and personalized information and offers with the intention of increasing and improve sales. Unfortunately, this information can also be misused (reduce the outcome of negotiations), and the client has little or no control over what information is being logged.

To get the benefits of directed and personalized deals, personal assistants can help the client to be anonymous in negotiations and in the search for information. Instead of having the seller register personal information about a user, private information is maintained by the personal assistant who asks the seller for information and filters it according to the clients wishes.

Identity based trading is necessary for trades where the participants risk that the other party defaults, or "defects" from the deal. In such situations, the identity is necessary to establish trust [36]. The agents should handle the establishment of trust via third parties, recommendations, etc. They also need social protocols for revoking trust if an agent betrays its trust.

As trust is necessary for some kinds of trades, trusted identities will become a commodity. This is particularly easy for software agents as an identity most probably will be proved by digital signatures or something that can be copied.

Conclusions

Although the appearence of electronic markets has created a need for agents that act in their clients' best interests, still most agent economies are computational economies rather than economies of self-interested agents.

Many agent systems are based on distributed programming protocols that rely on cooperative behavior. Their performance and results can not be expected to hold in a real market as self-interested agents have no a priori reason to give up a good deal without compensation. If we are ever to be able to delegate tedious work to software agents in open systems, the focus of future work must be on understanding how selfish agents should allocate resources, using speculation as needed, and how they should avoid potential problems arising from speculation.

Some requirements on open markets with truly self-interested agents

- Agents should not be constrained by predefined bidding strategies
- Agents should not be assumed to bid their reserve price
- Agents should be able to resell resources

Issues that need further investigation

- Speculation and reducing risk from speculation
- Strategies for bidding on combinations of goods
- Decentralization of markets

- Trust as a commodity

We expect that research in electronic markets in the coming years will increasingly emphasize truly self-interested agents.

References

1. Tuomas Sandholm, Negotiation among Self-Interested Computationally Limited Agents. Ph.D. Dissertation. University of Massachusetts at Amherst, Department of Computer Science, 1996.
2. James F. Kurose, Mischa Schwartz and Yechiam Yemini, A Microeconomic Approach to Decentralized Optimization of Channel Access Policies in Multiaccess Networks, Proceedings of the 5th International Conference on Distributed Computing Systems, Denver, Colorado, May 13-17, 1985.
3. Spyros Lalis, Christos Nikolaou, Dimitris Papadakis and Manolis Marazakis, Market-driven Service Allocation in a QoS-capable Environment, Proceedings of the First Int. Conf. on Information and Computation Economics ICE'98, Charleston, SC, USA, October, 1998.
4. Yair Amir, Baruch Awerbuch, R. Sean Borgstrom, A Cost-Benefit Framework for Online Management of a Metacomputing System, Proceedings of the First Int. Conf. on Information and Computation Economics ICE'98, Charleston, SC, USA, October, 1998.
5. Ori Regev and Noam Nisa, The POPCORN Market - an Online Market for Computational Resources, Proceedings of the First Int. Conf. on Information and Computation Economics ICE'98, Charleston, SC, USA, October, 1998.
6. Jonathan Bredin, David Kotz and Daniela Rus, Market-based Resource Control for Mobile Agents, Proc. of the Second Int. Conf. on Autonomous Agents, AA'98, Minneapolis/St. Paul, May, 1998
7. John M. Schinizlein, How Can Routers Help Internet Economics?, Proceedings of the First Int. Conf. on Information and Computation Economics ICE'98, Charleston, SC, USA, October, 1998.
8. C. Courcoubetis, V.A. Siris and G.D. Stamoulis, Network control and usage-based charging: Is charging for volume adequate?, Proceedings of the First Int. Conf. on Information and Computation Economics ICE'98, Charleston, SC, USA, October, 1998.
9. D. Reininger, D. Raychaudhuri and M. Ott, Market Based Bandwidth Allocation Policies for QoS Control in Broadband Networks, Proceedings of the First Int. Conf. on Information and Computation Economics ICE'98, Charleston, SC, USA, October, 1998.
10. Prasad Chalasani, Jha Somesh, Shehory Onn and Katia Sycara, Query Restart Strategies for Web Agents, Proc. of the Second Int. Conf. on Autonomous Agents, AA'98, Minneapolis/St. Paul, May, 1998
11. Kihong Park, Meera Sitharam and Shaogang Chen, Quality of Service Provision in Noncooperative Networks: Heterogenous Preferences, Multi-Dimensional QoS Vectors, and Burstiness, Proceedings of the First Int. Conf. on Information and Computation Economics ICE'98, Charleston, SC, USA, October, 1998.
12. Peter C. Fishburn and Andrew M. Odlyzko, Dynamic Behavior of Differential Pricing and Quality of Service Options for the Internet, Proceedings of the First Int. Conf. on Information and Computation Economics ICE'98, Charleston, SC, USA, October, 1998.
13. Nir Vulkan and Nicholas R. Jennings, Efficient Mechanisms for the Supply of Services in Multi-Agent Environments, Proceedings of the First Int. Conf. on Information and Computation Economics ICE'98, Charleston, SC, USA, October, 1998.

14. Nagao Ogino, Connection Establishment Protocol Based on Mutual Selection by Users and Network Providers, Proceedings of the First Int. Conf. on Information and Computation Economics ICE'98, Charleston, SC, USA, October, 1998.

15. Tracy Mullen and Jack Breese, Experiments in Designing Computational Economies for Mobile Users, Proceedings of the First Int. Conf. on Information and Computation Economics ICE'98, Charleston, SC, USA, October, 1998.

16. Mehmet Karaul and Yannis A. Korilis, Ariel Orda, A Market-Based Architecture for Management of Geographically Dispersed, Replicated Web Servers Proceedings of the First Int. Conf. on Information and Computation Economics ICE'98, Charleston, SC, USA, October, 1998.

17. A. Anastasiadi and S. Kapidakis, C. Nikolaou, J. Sairamesh, A Computational Economy for Dynamic Load Balancing and Data Replication, Proceedings of the First Int. Conf. on Information and Computation Economics ICE'98, Charleston, SC, USA, October, 1998.

18. Jaaka Sairamesh and Jeffrey O. Kephart, Price Dynamics of Vertically Differentiated Information Markets, Proceedings of the First Int. Conf. on Information and Computation Economics ICE'98, Charleston, SC, USA, October, 1998.

19. Gerald J. Tesauro and Jeffrey O. Kephart, Foresight-based pricing algorithms in an economy of software agents, Proceedings of the First Int. Conf. on Information and Computation Economics ICE'98, Charleston, SC, USA, October, 1998.

20. Juan A. Rodriguez-Aguilar, Francisco J. Martin, Pablo Noriega, Pere Garcia and Carles Sierra, Competitive Scenarios for Heterogeneous Trading Agents, Proc. of the Second Int. Conf. on Autonomous Agents, AA'98, Minneapolis/St. Paul, May, 1998.

21. Peter R. Wurman, Michael P. Wellman and William E. Walsh, The Michigan Internet AuctionBot: A Configurable Auction Server for Human and Software Agents, Proc. of the Second Int. Conf. on Autonomous Agents, AA'98, Minneapolis/St. Paul, May, 1998.

22. John E. Collins, Ben Youngdahl, Scott Jamison, Bamshad Mobasher and Maria Gini, A Market Architecture for Multi-Agent Contracting, Proc. of the Second Int. Conf. on Autonomous Agents, AA'98, Minneapolis/St. Paul, May, 1998.

23. J. Eriksson, N. Finne and S. Janson, SICS Marketspace - An Agent-Based Market Infrastructure, Proc. of the Workshop on Agent Mediated Electronic Trading, AMET-98, Minneapolis/St. Paul, May, 1998.

24. Jeffrey K. MacKie-Mason and Hal R. Varian, Pricing the Internet, in Public Access to the Internet, Brian Kahin and James Keller, eds. Cambridge, MA, MIT Press, 1995.

25. Keith S. Decker and Victor R. Lesser, Designing a Family of Coordination Algorithms, Technical Report 94-14, Univ. of Massachusetts, MA, USA, 1995.

26. Jean-Marc Andreoli, Uwe M. Borghoff and Remo Pareschi, Constraint-Based Knowledge Brokers, Proc. of Parallell Symbolic Computation, PASCO’94, Linz, Austria,1994.

27. Cristiano Castelfranchi and Rosaria Conte, Limits of Strategic Rationality for Agents and M-A Systems, ModelAge Workshop, Certosa di Pontignano, Italy, January 1997.

28. N. R. Jennings, T. J. Norman and P. Faratin, ADEPT: An Agent-based Approach to Business Process Management, ACM SIGMOD Record 27 (4), 1998.

29. Tuomas Sandholm, Limitations of the Vickrey Auction in Computational Multiagent Systems, Second International Conference on Multiagent Systems, ICMAS-96, Kyoto, Japan, December, 1996.

30. L.H. Ungar, D.C. Parkes and D.P. Foster, Cost and Trust Issues in On-Line Auctions, Proc. of the Workshop on Agent Mediated Electronic Trading, AMET-98, Minneapolis/St. Paul, May, 1998.

31. Jeffrey S. Rosenschein and Gilad Zlotkin, Rules of Encounter: Desiging Conventions for Automated Negotiation among Computers, MIT Press, Cambridge, Mass., USA, 1994.

32. E. Ephrati, Gliad Zlotkin and J.S. Rosenschein. Meet Your Destiny: A Non-manipulable Meeting Scheduler, Conference on Computer Supported Cooperative Work, North Carolina, October, 1994.

33. Dale O. Stahl, Andrew B. Whinston and Keqiang Zhan, A Simulation Study of Competitive Internet Pricing: AOL Flat Rates versus GSW Usage Prices, Proceedings of the First Int. Conf. on Information and Computation Economics ICE'98, Charleston, SC, USA, October, 1998.

34. Tuomas Sandholm and Fredrik Ygge, On the Gains and Losses of Speculation in Equilibrium Markets, Proc. of the Fifteenth International Conference on Artificial Intelligence, IJCAI '97, Morgan Kaufmann, 1997.

35. Michael P. Wellman, A Market-Oriented Programming Environment and its Applications to Distributed Multicommodity Flow Problems, Journal of Artificial Intelligence Research 1(1993):1-23.

36. Lars Rasmusson, Socially Controlled Global Agent Systems, Master's Thesis, 96-x-20, DSV, Royal Inst. Of Technology, Stockholm, October, 1996.