Research Statement
Kate Larson

Interests: Artificial Intelligence, Multiagent Systems, Ecommerce, Game Theory, Mechanism Design, Automated Negotiation, Coalition Formation, Models of Bounded Rationality, Resource-Bounded Reasoning

We are entering into an era where distributed open computer systems are being used by self-interested agents – users and software agents that represent them – which have their own private goals and information. A fundamental problem in these diverse multiagent systems is designing protocols, that is *mechanisms*, which produce desirable system-wide outcomes despite the self-interest of the participating agents.

The frameworks of game theory and mechanism design have had significant influence on formal models of multiagent systems by providing tools for designing and analyzing systems in order to guarantee desirable outcomes. However, many game theoretic models ignore computation and deliberation issues. In particular, it is often assumed that agents are fully rational, in that they know how much they value different alternative outcomes. I argue that this knowledge may not be easy to come by. Rather than being fully rational, agents must *deliberate*, that is agents must compute or gather information in order to determine their valuations. These *deliberative agents* have constraints and bounds on their deliberation resources and capabilities which restrict their ability to optimally determine their values for different outcomes and optimally select actions to take. This creates a potentially hazardous gap in the use of classical game theory in computational settings, since deliberative agents may not be motivated to behave in the desired way.

In my dissertation I have focused on understanding and bridging this gap. In particular, I have developed a theory of interaction for self-interested deliberative agents. I have introduced a new, more refined deliberation control method into the resource-bounded reasoning literature. This allows for the explicit inclusion of agents' deliberation actions into a game theoretic model. I proposed a new game theoretic solution concept, the deliberation equilibrium. This solution concept allowed me investigate the impact that constraints on agents' deliberation resources have on agents' strategies. It has also provided a foundation for developing mechanism design tools for settings where there are deliberative agents. In general, my work highlights the deep interaction between game-theoretic incentive issues and computational and deliberation issues. While the tools and techniques that I develop can be used in a wide range of settings, a key goal and application of my research is to design electronic auctions and markets, as well as software agents that act optimally on behalf of real-world parties in these marketplaces.

In the next sections I describe previous and current research in more detail. I discuss the normative model for bounded-rationality that I have developed. Using this model, I have been able to rigorously study and understand the strategic behavior of deliberative agents in classical negotiation mechanisms such as bargaining and auctions. I have studied the possibility of designing new mechanisms (protocols) for deliberative agents which take into account their deliberative capabilities. While not part of my dissertation, I shortly describe work I have done on coalition formation for self-interested agents. I conclude by discussing some research directions I am interested in pursuing.

Normative Models of Bounded Rationality

I have developed a fully normative model for deliberation control, the performance profile tree [3]. It allows agents to fully condition their deliberation on all results and features of execution so far. Having such a fully normative model is of particular importance for designers of mechanisms in multiagent deliberative settings. It gives designers the ability to prove properties about deliberation control strategies (such as the optimality of following a certain deliberation strategy) since the deliberation control method does not rely on heuristics and simplifying assumptions.

Using the performance profile tree as the model of deliberation control, I have explicitly incorporated the deliberation actions of agents into a game-theoretic formalization, and have introduced a game-theoretic solution concept, the deliberation equilibrium. The deliberation equilibrium is a foundation for both analyzing existing mechanisms to see what the strategic behavior of deliberative agents is, and for designing mechanisms for deliberative agents in particular. Additionally, I have also developed tools for measuring the impact that deliberation has on the entire system [6].

In experimental work, I have shown that the performance profile tree is a useful deliberation control method that can be used in practice. It is feasible to build from performance data of algorithms and is feasible to use. It improves the accuracy of deliberation control compared to prior techniques [1].

Deliberative Agents in Classical Negotiation Mechanisms

Two common negotiation settings in multiagent systems, and particularly in markets, are one-to-one negotiation (bargaining) and one-to-many negotiation (auctions). Using the performance profile tree as the model for deliberation control, I have studied the strategic behavior of deliberative agents in different negotiation settings.

In bargaining (one-to-one negotiation) agents must deliberate in order to determine their values for acting independently and coordinating their actions. While they are deliberating, agents must also negotiate as to whether they will coordinate or not. I have analyzed the two most common bargaining mechanisms – single-shot and alternating-offers – and have characterized the deliberation equilibria in each setting [3] [5]. In particular, I developed algorithms for computing all deliberation equilibria in the single-shot bargaining mechanism, and showed that all deliberation equilibria in the alternating-offers bargaining mechanism have the same form as the equilibria in the single-shot mechanism.

Turning attention to auction mechanisms, I have studied the strategic behavior of deliberative agents in all the classical 1-object and multi-object auction protocols. I discovered a new type of strategic behavior – *strategic deliberation* – where agents actively deliberate on each others' deliberation problems. I showed that if agents have a cost associated with deliberation, then strategic deliberating occurred in equilibrium in all the classical auctions [4]. This means that auctions which are "strategy-proof" for rational agents (i.e. robust against strategic behavior since agents have dominant strategies), no longer have this property if agents are deliberative.

Mechanism Design for Deliberative Agents

Since classic market mechanisms are not robust against strategizing by deliberative agents, I studied the possibility of designing new market mechanisms (for example, auctions) for such agents. I proposed a set of intuitive properties which are desirable in mechanisms for deliberative agents [7]. First, mechanisms should be *non-deliberative*: the mechanism should not solve the deliberation problems for the agents. Secondly, mechanisms should be *deliberation-proof*: agents should not be motivated to deliberate on others' valuations. Third, the mechanism should be *non-deceiving*: agents should be motivated to truthfully reveal their values obtained through deliberating. Finally, the mechanism should be *sensitive*: agents' actions should affect the outcome.

I derived a fundamental impossibility result [7]: no direct-revelation mechanism satisfies the four properties, where by direct-revelation it is meant that in a single step agents directly report any information to the mechanism, which then determines an outcome. Moving beyond direct-revelation mechanisms, I showed that no *value-based* mechanism (that is, mechanisms which are not necessarily direct-revelation but where the agents are only asked to report information on valuations - either partially or fully determined ones) satisfies these four properties.

Coalition Formation for Self-Interested Agents

In work on coalition structure generation for self-interested agents, I developed an anytime algorithm that searches through the space of possible coalition structures and establishes a tight bound from optimal within the minimal amount of search, and showed that any other algorithm has to search strictly more [8]. If additional time remains, the algorithm searches further, establishing a progressively tighter bound. My experiments showed that, in the average case, most anytime algorithms that search the space of coalition structures do much better than the theoretical worst case, however they tend to be very sensitive to the underlying instance distribution [2].

Future Research Directions

In the short term there are several concrete problems that have arisen from my dissertation which I intend to work on. In particular, I plan to investigate approaches for overcoming the impossibility result discussed earlier. It may be possible to relax some of the desired properties, while still being able to achieve desirable outcomes from a designers perspective. For example, by allowing the mechanism do some minimal amount of deliberation control for agents, it may be possible to avoid strategic deliberation and deception by the agents. In another direction, it may be possible to design markets where agents that are “good” at solving problems, do the deliberating for other agents. Determining how such a market should be structured so that the “right” problems are solved is an interesting avenue I am considering exploring. Preference elicitation is another research area in which I think that the approaches I have taken may prove fruitful.

In the long run I am certain that there will be no shortage of inspiring problems. My work lies in the intersection of artificial intelligence, economic theory, and algorithms. This is an exciting and growing research area with numerous applications including the design of electronic marketplaces, distributed optimization, and networks. In particular, with the advent of electronic commerce there is a need to understand systems in which incentives and computational issues interact in non-trivial ways, which can then lead to the building of robust applications.

I am excited about building a research group to study problems in these areas, and am looking forward to developing research collaborations with both computer scientists and game theorist/economists in an open-minded academic environment. My long term research goal is to achieve a thorough understanding of the effects computation and bounded rationality have on incentives.

References

[1] Kate Larson and Tuomas Sandholm. Performance profile trees. to be submitted early 2004.

[2] Kate Larson and Tuomas Sandholm. Anytime coalition structure generation: An average case study. *Journal of Experimental and Theoretical AI*, 11:1–20, 2000. Early version appeared in the International Conference on Autonomous Agents (AGENTS), pp. 40–47, 1999.

[3] Kate Larson and Tuomas Sandholm. Bargaining with limited computation: Deliberation equilibrium. *Artificial Intelligence*, 132(2):183–217, 2001. Short early version appeared in the Proceedings of the National Conference on Artificial Intelligence (AAAI), pp. 48–55, Austin, TX, 2000.

[4] Kate Larson and Tuomas Sandholm. Costly valuation computation in auctions. In *Theoretical Aspects of Rationality and Knowledge (TARK VIII)*, pages 169–182, Siena, Italy, July 2001.

[5] Kate Larson and Tuomas Sandholm. An alternating offers bargaining model for computationally limited agents. In *First International Conference on Autonomous Agents and Multiagent Systems*, Bologna, Italy, July 2002.

[6] Kate Larson and Tuomas Sandholm. Miscomputing ratio: The social cost of selfish computing. In *The Second International Conference on Autonomous Agents and Multiagent Systems*, Melbourne, Australia,

2003. Early version appeared in the AAAI-02 workshop on Game-Theoretic and Decision-Theoretic Agents, Edmonton, Canada.

[7] Kate Larson and Tuomas Sandholm. Strategic deliberation and truthful revelation: An impossibility result. submitted, under review, 2003.

[8] Tuomas Sandholm, Kate Larson, Martin Andersson, Onn Shehory, and Fernando Tohmé. Coalition structure generation with worst case guarantees. *Artificial Intelligence*, 111(1–2):209–238, 1999. Early version appeared at the National Conference on Artificial Intelligence (AAAI), pages 46–53, 1998.