Multiagent Systems: Milestones and New Horizons

Sandip Sen
Department of Mathematical & Computer Sciences
University of Tulsa
600 South College Avenue
Tulsa, OK 74104-3189
sandip@kolkata.mcs.utulsa.edu

Abstract

Research in multiagent systems (MAS), or Distributed AI, dates back to late 70's. Initial work in the area focused on distributed interpretation of sensor data, organizational structuring, and generic negotiation protocols. But several recent developments have helped reshape the focus of the field. Like the rest of AI, the field has matured from being largely exploratory in nature to focusing on formal theories of negotiation, distributed reasoning, multiagent learning, and communication languages. The field is also maturing to the point of developing its first few fielded applications. The recent widespread interest in the internet, the world-wide-web, and intelligent agent applications have further fueled the need for techniques and mechanisms by which agents representing users can effectively interact with other agents in open, dynamic environments. The development of several new international workshops and conferences have helped focus research in the area. The field is poised at a critical juncture with stimulating problems and challenges promising some very exciting developments in the next few years.

Keywords: multiagent systems, Distributed AI, coordination, cooperation, negotiation, intelligent agents, agent architecture
Running title: Multiagent Systems

Research in the area of Distributed Artificial Intelligence (DAI) has focused on developing computational mechanisms by which multiple intelligent and autonomous agents can effectively coordinate. The field is nearing the end of its second decade of existence. As to be expected, this sub-area of AI have matured over time. Recently, new influences and developments in related areas are helping reshape and vitalize the field. To reflect the focus of interest of the active researchers, the field has now adopted the name of Multiagent Systems (in the past this name was used to refer to a sub-area in the field). This development was followed closely by the development of a new international conference in the area [1, 2].

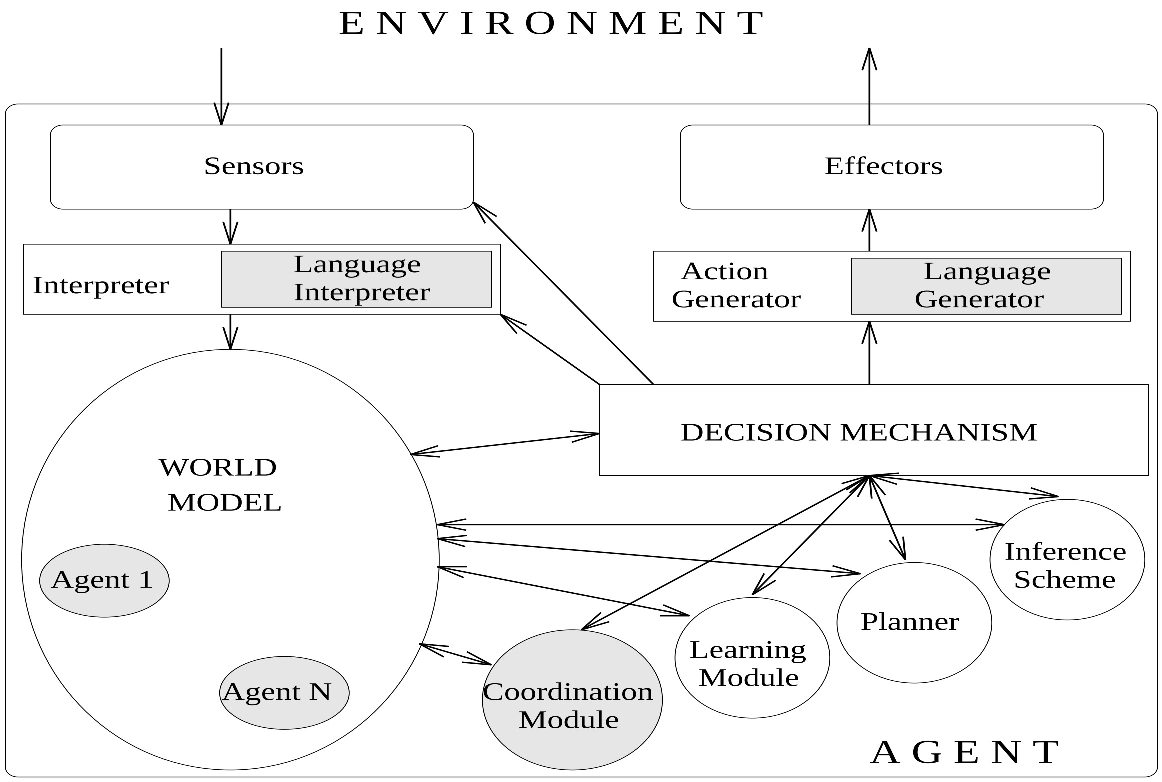

Figure 1: An agent architecture (shaded modules represent components particular to agents in a MAS).

This short overview article is meant to give the reader a feel for core research issues in the field. The interested reader should also look up other more extensive, though somewhat outdated, reviews of the field [3, 4, 5]. Related reviews on intelligent/software agent literature contain more recent references to MAS work [6, 7]

In Figure 1 I present an architecture of an intelligent agent situated in a multiagent environment. I have highlighted parts of the architecture which are unique to agents in a multiagent world. For example, such agents may need to communicate with other agents and hence should be able to generate and understand some communication language. Also, an agent in a multiagent world will benefit from building explicit models of other agents and reasoning about them separately from other aspects of the environment. For example, it is often useful to model the beliefs and desires of other agents to predict their future behavior. The coordination module encapsulates knowledge required to coordinate with other agents. This may include mechanisms for forming shared commitments, protocols for interaction, reasoning with social laws, etc. Other components of this architecture, e.g., the planner and learning module must also be designed to reason explicitly about other agents. These and other considerations make multiagent research issues unique and at times complimentary to single agent (or software agent) research issues [8].

As the field has matured, some of the old issues have been rephrased and studied in new context, and new issues and problems have cropped up. In particular, the approaches and methodologies being currently used by active researchers have undergone notable changes from the early days of the field. After summarizing the research directions and achievements of the field to date, I will emphasize those research issues and approaches that are likely to be actively investigated by MAS researchers in the next few years. Without a doubt, this focusing of topics will be biased by my own research experience and preferences. I still believe that the reader will be able to glean the significant challenges facing the field at this point in its short history.

MAS research can be broadly classified into the following classes based on the relationships between the agents:

Cooperative agent systems: A group of cooperative agents jointly work on achieving a common goal. The key research issues in these systems are

- how do agents decompose the goal into subgoals that can then be assigned to individual agents based on their capabilities and access to resources,
- how to develop agent organizations (authority relationships) and problem solving protocols (information flow) that enable agents to share results and knowledge in a timely, effective manner
- how do agents maintain coherence and problem solving focus when locally available information can be incorrect, inconsistent, outdated, etc.

Non-cooperative agent systems: A group of self-interested agents interact in a shared environment. Such domains include both systems were agents are adversarial to each other (e.g., bargaining parties, game players) as well as domains where agents are indifferent to each other (e.g., drivers sharing a highway are neither cooperative nor

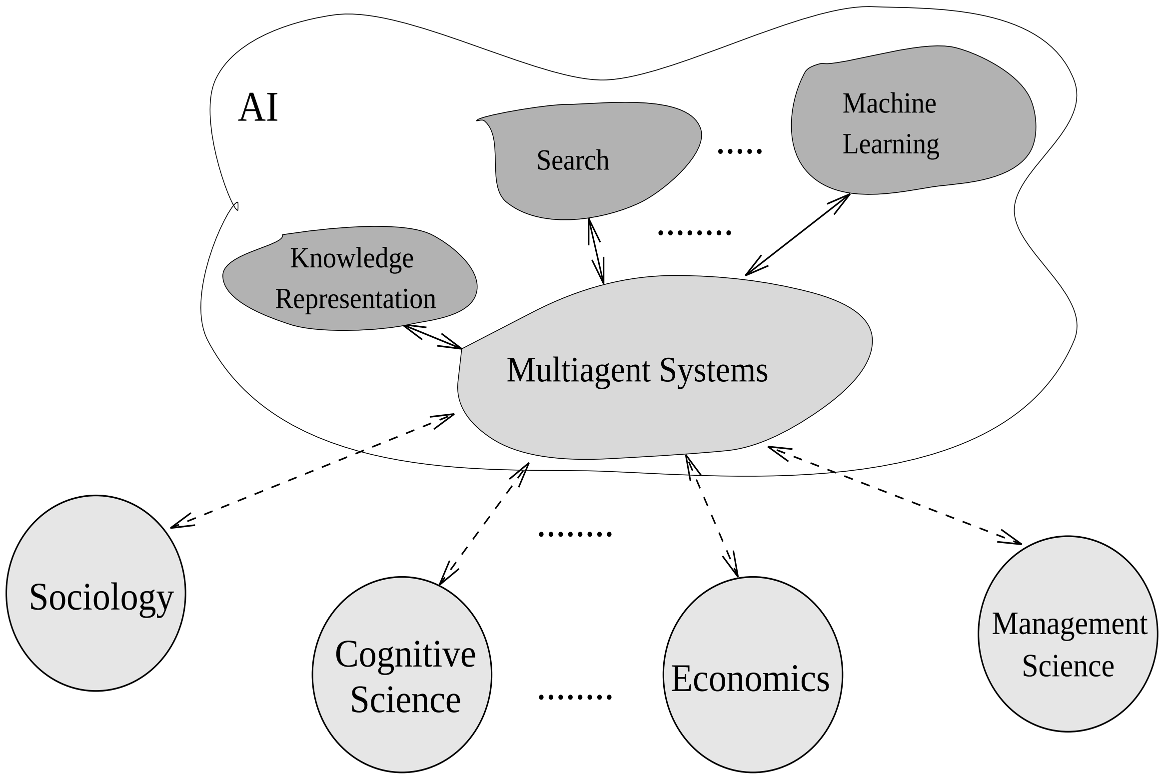

Figure 2: Synergies of multiagent system research with the rest of AI and other related fields.

adversarial: they interact because they share a common resource). In the latter case the key research problems are designing social laws, conventions, and protocols by which each agent can achieve its own goal without significantly affecting the chances of others achieving their goals. Research involving adversarial agents concentrates on issues like modeling the knowledge and behavioral strategies of opponents, learning to exploit opponent weaknesses, and developing interaction rules by which agents can arrive at equilibrium configurations.

From the description above it should be clear that in order to succeed, any significant multiagent research endeavor has not only got to borrow from the past work in traditional single agent AI, but should utilize relevant techniques and results from more traditional fields like economics, sociology, management sciences, etc. Figure 2 depicts the supportive relationship multiagent system research utilizes from all of these related fields. To illustrate the mutual enrichment of multiagent research and other related disciplines I will use two illustrative synergies:

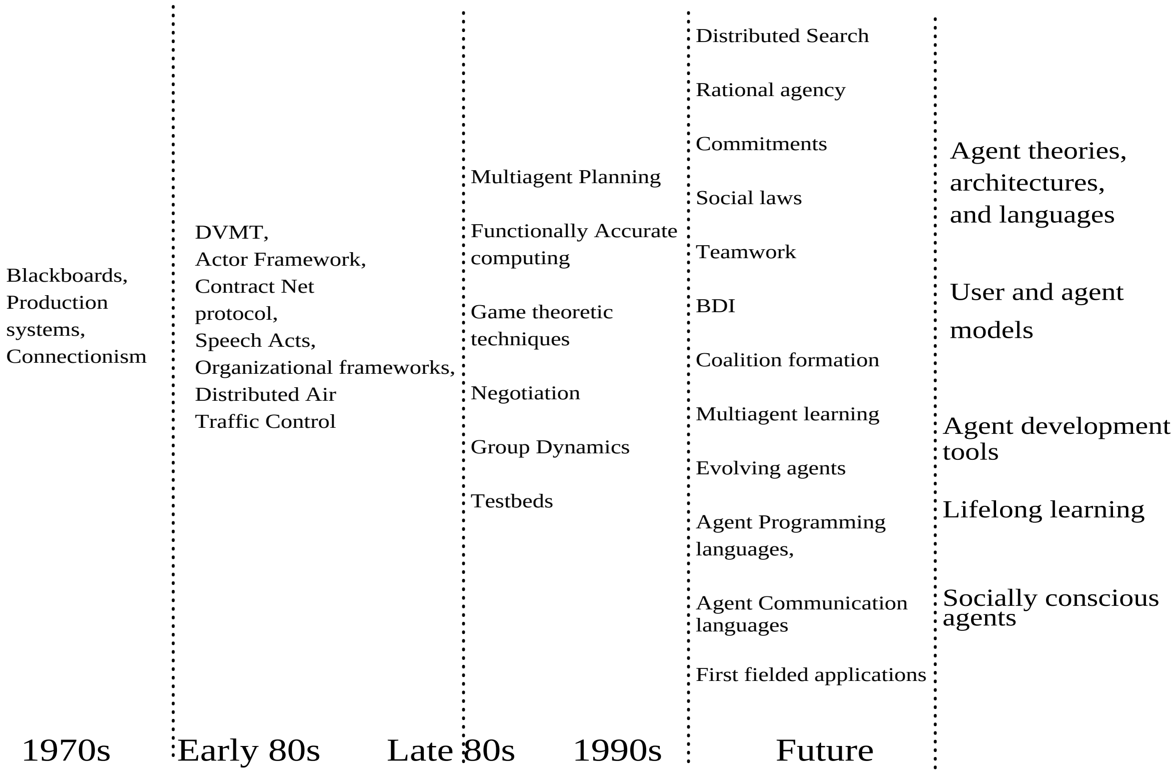

Figure 3: A chronological view of MAS research.

- Well-studied economic principles like market mechanisms, game theoretic concepts of equilibrium, etc. have been used to design agent interaction strategies in a number of domains. Whereas game theory researchers typically assume unbounded computing resources, multiagent systems researchers have been contributing to the game theory literature by developing negotiation schemes where agents have only limited computing power (bounded rationality).

- Within AI, MAS provides a new application domain for machine learning techniques. At the same time particular requirements for MASs violate typical assumptions underlying most machine learning techniques (e.g., environments can no longer be assumed to be stationary when multiple agents learn concurrently) and provides new learning modalities, e.g., use of communication, that can lead to significant contributions to the machine learning literature.

Some of the often cited reasons for building and using a MAS are the following:

1. Some problem domains are inherently distributed in nature, e.g., distributed sensor interpretation, coordination of self-interested agents, etc. Also agents that coordinate with other artificial agents and associated humans open the doors for a whole new genre of computer applications.

2. Agents in a MAS are often designed based on functionality; such systems are modular by design and hence can be easier to develop and maintain.

3. MASs with adaptive capabilities can be robust than centralized systems.

4. MASs provide a very useful framework within which social aspects of intelligent behavior can be modeled, analyzed, and evaluated under a wide variety of domain, behavior, and knowledge assumptions.

In Figure 3 I present an approximate timeline of the major milestones in the field of MAS research. I have grouped the principal research events into one of several groups. The boundaries of the groups are not crisp. But the different time periods represent when different research issues came to the fore in the MAS community. In the following sections, I will discuss these time periods in more detail.

Prehistory

The roots of MAS research can be traced back to early AI systems including blackboard systems [9], connectionism, production systems [10], Minsky's *Society of Mind* concept [11], and related work. Distribution and parallelism of computation are the themes common to these models of problem solving. Though such systems were frequently run on serial machines, they all embodied the potent concept of independent computational entities asynchronously working on shared problems and data. The granularity of the computational entities, however, varied widely from nodes in connectionist networks to rules in production systems to knowledge sources in blackboard systems.

The early years

The first MAS systems can be considered descendants of the above-mentioned programming paradigms. Prominent early work in MAS include the following:

Distributed Vehicle Monitoring Testbed (DVMT): Victor Lesser together with his colleagues and students developed the DVMT framework at the University of Massachusetts, Amherst, and used it for a sustained study of interpretation of distributed sensor data by a group of cooperative, autonomous problem solvers [12, 13, 14]. Each problem solver used a blackboard architecture and had to coordinate with other problem solvers to develop a globally consistent picture of vehicular traffic movement given only local sensor information. The major MAS issues investigated using this testbed were agent organizations, sharing of partial plans, effects of incomplete and inconsistent global knowledge on local problem solving, etc.

The Actors framework: Carl Hewitt and co-workers developed a model of computation based on a multitude of message passing *actors*, where control is distributed and computation is fine-grained, and inherently parallel in nature [15, 16]. Hewitt proposed the actors framework as a model of problem solving that parallels the interaction of a scientific community of experts [17], and also allows agents to work in *open systems* where the agent group composition and the environment changes over time and agents have to cope with mutually inconsistent knowledge [18, 19].

The Contract Net Protocol: Perhaps no single MAS technique has been as widely used as the Contract-Net problem solving protocol developed by Reid Smith and colleagues [20, 21]. The contract net is a high level protocol designed to solve the *connection problem*. It allows agents with problems to be solved and agents who can solve those problems to find each other and arrive at mutually acceptable contracts. Though the contract net specifies the roles of the communicating parties and the nature, direction, and sequence of information flow between them, it provides almost no constraints on how contracts are developed, when to bid for a contract, how to negotiate, etc. Interestingly, this open-ended nature of the contract net has also been its selling point, allowing it to be flexibly adapted in a number of different application domains [22, 23, 24, 25]. This also permits continued interest in developing auxiliary techniques to better use the contract net [23, 26].

Other early work of significance include work on designing and evaluating agent organizations [27, 28], speech act theory [29], limits of common knowledge in a distributed environment [30], and distributed air traffic control [31].

The last decade

The second generation of influential MAS researchers took center-stage during the first part of the last decade. Some of them were students of the prominent researchers who set the ball rolling in the first place, e.g., Ed Durfee was a student of Victor Lesser. Two prominent researchers of this decade, Les Gasser and Mike Huhns, deserve special credit for their contribution not only as active researchers but also as tireless champions for promoting the cause of the field. Between themselves they edited the three most prominent collections of papers in the field [32, 33, 34]. Huhns has also maintained an e-mail digest for discussing MAS issues (send mail to DAI-List-Request@ece.sc.edu to sign up) which has been an invaluable source of discussion and dissemination of current MAS related information.

The last decade had started with two "hot" areas of investigation within the DAI community. The first of this was multiagent planning. A variety of researchers brought a diverse set of backgrounds, experience, and expertise to attack the problem of agents planning to achieve their goals together with or in the presence of other agents. Some of the notable approaches developed include a centralized planner planning for all agents [35], a single agent resolving conflicts between single-agent plans developed individually [36], and agents sharing partial plans to arrive at globally coherent plans over time [37, 38].

The other area to receive increasing attention was the use of game theoretic techniques to coordinate autonomous agents. Most of this work assumed that each agent is rational in the sense that it chooses actions to maximize its own utility [39, 40] (in contrast to cooperating agents, whose only motivation is to coordinate with other agents in the group to achieve a common goal).

Other work of significance during this period include Katia Sycara's work on using case based reasoning to allow bargaining parties to arrive at a compromise deal [41], and Hubermann and colleagues work on analyzing the dynamics of interactions between a large number of interacting agents following simple behavioral rules [42, 43]. In addition, two DAI testbeds developed during this period gave researchers a common platform to evaluate new coordination schemes and was also used as pedagogical tools in courses on DAI offered around the world [44, 45].

The research issues that came to the forefront over the last five years are more diverse:

Distributed Search: Coordination can be seen as a search for mutually compatible action sequences by agents searching over their local problem solving spaces [46]. Interesting applications of this approach was demonstrated in distributed scheduling [47, 48] and distributed constraint satisfaction [49].

The economics of negotiation: Rosenschein and his students [50, 51, 52, 53] as well as a number of other researchers [54, 55] produced a sustained volume of work on adapting game theory, voting theory, and other mechanisms from economics to address MAS problems. The goal of most of this work is to eliminate strategic thinking or manipulation by individual agents, and developing mechanisms by which even competitive agents can arrive at a compromise solution (equilibrium). The focus on deliberating under bounded computational resources distinguishes this body of work from standard economics literature.

Social reasoning: I am grouping a number of loosely related research ideas under the theme of social reasoning. All of these work deal with improving problem solving in the context of a group of agents. The following list is a sampling of this body of work:

- the design of conventions [56], social laws [57], or domain-independent rules for good teamwork [58],
- identification and resolution of social dilemmas where greedy strategies used by everyone would lead to worse results for all [59],

- mechanisms for arriving at, monitoring, fulfilling, and retracting commitments to oneself or to others [60, 56, 61, 62].
- how agents form shared intentions [63, 64] or collaborative plans [65]; one of the most influential agent architectures to be developed, the Belief-Desire-Intention or BDI architecture [66], has had similar motivations.
- decision procedures by which agents can decide who they should cooperate with [67] or how to form coalitions [54, 68, 69, 70].

Multiagent Learning: Over the last couple of years there has been an increasing realization that coordination strategies designed offline are too rigid and inflexible for dynamic, open environments. As such, there is a need for agents to adapt their behaviors online based on interactions with other agents and feedback from the environment. Learning, adaptation, and evolutionary mechanisms are being used to allow agents to adapt to changing demands of dynamic environments [71, 72, 73, 74, 75, 76].

Fielded Applications: The field also produced the first few fielded applications [77, 78, 79] during this period. Two applications in particular deserve special mention. The first of these is the ARCHON project, which provides an architecture for integrating preexisting "legacy systems", and was used to implement a large scale fielded application for electricity transport management in Europe. The second of this is the Carnot project developed at MCC, which uses distributed, knowledge-based communicating agents that cooperate with each other to allow associated users to efficiently and transparently navigate enterprise-wide heterogeneous information [80]. The prototype system is available for a number of Unix platforms, and is fielded by companies and organizations like Ameritech, Boeing, Department of Defense, Eastman Kodak, etc.

Other notable events of recent times are the development of a high-level agent communication language based on speech acts [81] and attempts to develop agent technology standards by the Foundation for Intelligent Physical Agents (FIPA) which has representation from both the academia and industry. The first version of such a specification, FIPA97 ver.1.0, is available on the WWW (`http://drogo.cselt.stet.it/fipa/spec/fipa97.htm`) with specifications for agent management, agent to agent communication, agent/software integration, personal travel assistance, personal assistant, audio-visual entertainment and broadcasting, and network management and provisioning. The maturity of the multiagent systems field is also reflected in the publication of a textbook [82], while another one is currently being compiled by Gerhard Weiß.

A look ahead

Like most of AI, MAS research have shifted focus from building grand unified theories to developing specialized techniques to address requirements for well-defined problem classes. I believe this trend will continue over the next few years primarily due to the need for developing fielded applications. To be successful, MAS technology should complement other well-understood technology developed in computer science. For example, multiagent architectures can be gainfully employed for integrating multiple, stand-alone legacy systems. Another area that is likely to receive increasing attention is the use of multiagent architectures for information retrieval or for developing digitized information repositories. The first genre of these systems are already being tested in university laboratories [83]. Other viable and highly probable application areas for MAS include: electronic commerce over the internet, long distance medical care, electronic help desk platforms, agents for managing, integrating and disseminating information in organization wide information sources.

I believe the following research issues and questions will need to be better addressed in the next few years to meet the challenges facing this field:

Agent architectures: More comprehensive agent architectures need to be developed to integrate pre-compiled knowledge, sensor information, negotiation models, planning capabilities, learning mechanisms, and communication modules. A series of workshops titled Agent Theories, Architectures, and Languages (ATAL) is being used by researchers to focus on this particular problem [84].

User and agent modeling: To be effective as assistants to users, agents need to represent not only hard constraints prescribed by users, but also should be able to represent and reason with soft constraints in the form of preferences and biases of the associated user.

Often these cannot be immediately enumerated and the agent must be able to learn more about the user by building and updating user models based on feedback. Also, agents should be able to update such models as the user changes his or her preferences. Agents in an MAS need to continually interact, communicate, and deliberate with other agents in their environment. As such, agents also need to represent, utilize, and elaborate models of other agents they interact with. The lack of a global context and the sparsity of data makes the latter problem a particularly difficult one.

Security considerations: To be trustworthy, agents must not reveal private information about its user to unauthorized agents. But, its action may inadvertently reveal such information (for example, an information gathering agent may be monitored to find out the interests of its user). Planning actions that will achieve goals but will not expose sensitive information is a critical open problem.

Flexibility and reliability: In a cooperative group, or in a coalition of self-interested agents, an agent will have some commitments to other agents. Fulfilling such commitments makes the agent reliable to the group. But locally detected contingencies or significant opportunities can lead to re-planning which may result in unfulfilled commitments. The need for flexibility can then cause a lack of reliability. There is a need for the development of a formal framework to address this responsiveness versus predictability tradeoff.

Continuous learning: To be effective in an open environment, static behaviors are insufficient. Learning to work with new agents as the group composition changes is a necessity. An intelligent agent must evolve over its lifetime, continuing to learn about its environment and fellow agents from experience. There is a need for exploring new modalities (e.g., using communication) and new approaches (e.g., pro-active versus reactive learning) for multiagent learning.

Conclusion

MAS research is poised at an exciting horizon. Initial uncertainties and hurdles have been overcome and the first fruits of labor has begun to materialize. The vista of the road ahead appears promising but challenging. Opportunities abound given the development of internet related technologies and the ubiquity of personal computers. There is a need for blending of existing technology and application of theoretical frameworks from related fields. Over and above all, there is a need for a concerted effort to tie the research and application issues with a clear identification of immediate and long term impact potentials. Without doubt, these are exciting times, as an active researcher in the field can look ahead with hope of contributing something significant to the information revolution.

Acknowledgments

This work has been supported in part by the National Science Foundation through a Research Initiation Award IRI-9410180 and a CAREER award IRI-9702672.

References

[1] Victor Lesser, editor. *Proceegings of the First International Conference on Multiagent Systems*, Menlo Park, CA, 1995. AAAI Press/MIT Press.

[2] Mario Tokoro, editor. *Proceegings of the Second International Conference on Multiagent Systems*, Menlo Park, CA, 1996. AAAI Press.

[3] B. Chaib-draa, R. Mandiau, and P. Millot. Distributed artificial intelligence: An annotated bibliography. *SIGART Bulletin*, 3(3):20–37, August 1992.

[4] Keith S. Decker. Distributed problem-solving techniques: A survey. *IEEE Transactions on Systems, Man, and Cybernetics*, 17:729–740, 1987.

[5] Edmund H. Durfee, Victor R. Lesser, and Daniel D. Corkill. Trends in cooperative distributed problem solving. *IEEE Transactions on Knowledge and Data Engineering*, 1(1):63–83, March 1989.

[6] Hyacinth S. Nwana. Intelligent agents: Theory and practice. *Knowledge Engineering Review*, 11(3):205–244, September 1996.

[7] Michael J. Woolridge and Nicholas R. Jennings. Intelligent agents: Theory and practice. *Knowledge Engineering Review*, 10(2), June 1995.

[8] Jeffrey M. Bradshaw. *Software Agents*. AAAI Press/The MIT Press, Menlo Park, CA, 1997.

[9] H. Penny Nii. Blackboard systems: The blackboard model of problem solving and the evolution of blackboard architectures. *AI Magazine*, 7(2):38–53, Summer 1986.

[10] A. Newell and H. A. Simon. *Human Problem Solving*. Prentice Hall, 1972.

[11] Marvin Minsky. *The Society of Mind*. Simon and Schuster, 1985.

[12] Edmund H. Durfee, Victor R. Lesser, and Daniel D. Corkill. Coherent cooperation among communicating problem solvers. *IEEE Transactions on Computers*, C-36(11):1275–1291, November 1987. (Also published in *Readings in Distributed Artificial Intelligence*, Alan H. Bond and Les Gasser, editors, pages 268–284, Morgan Kaufmann, 1988.).

[13] Victor R. Lesser and Daniel D. Corkill. Functionally accurate, cooperative distributed systems. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-11(1):81–96, January 1981.

[14] Victor R. Lesser and Daniel D. Corkill. The Distributed Vehicle Monitoring Testbed: A tool for investigating distributed problem solving networks. *AI Magazine*, 4(3):15–33, Fall 1983. (Also published in *Blackboard Systems*, Robert S. Engelmore and Anthony Morgan, editors, pages 353–386, Addison-Wesley, 1988 and in *Readings from AI Magazine: Volumes 1–5*, Robert Engelmore, editor, pages 69–85, AAAI, Menlo Park, California, 1988).

[15] Gul Agha. *Actors: A Model of Concurrent Computation in Distributed Systems*. MIT Press, Cambridge, MA, 1986.

[16] Carl Hewitt. Viewing control structures as patterns of passing messages. *Artificial Intelligence*, 8(3):323–364, Fall 1977.

[17] William A. Kornfeld and Carl E. Hewitt. The scientific community metaphor. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-11(1):24–33, January 1981. (Also published in *Readings in Distributed Artificial Intelligence*, Alan H. Bond and Les Gasser, editors, pages 311–320, Morgan Kaufmann, 1988.).

[18] Carl Hewitt. Open information systems semantics for distributed artificial intelligence. *Artificial Intelligence*, 47(1-3):79–106, 1991.

[19] Carl Hewitt and Jeff Inman. DAI betwixt and between: From "intelligent agents" to open systems science. *IEEE Transactions on Systems, Man, and Cybernetics*, 21(6):1409–1419, December 1991. (Special Issue on Distributed AI).

[20] Randall Davis and Reid G. Smith. Negotiation as a metaphor for distributed problem solving. *Artificial Intelligence*, pages 63–109, 1983.

[21] Reid G. Smith. The contract net protocol: High-level communication and control in a distributed problem solver. *IEEE Transactions on Computers*, C-29(12):1104–1113, December 1980.

[22] S. E. Conry, K. Kuwabara, V. R. Lesser, and R. A. Meyer. Multistage negotiation for distributed constraint satisfaction. *IEEE Transactions on Systems, Man, and Cybernetics*, 21(6):1462–1477, December 1991. (Special Issue on Distributed AI).

[23] Toumas Sandholm. An implementation of the contract net protocol based on marginal cost calculation. In *Proceedings of the Eleventh National Conference on Artificial Intelligence*, pages 256–262, Menlo Park, CA, 1993. AAAI Press/MIT Press.

[24] Sandip Sen and Edmund H. Durfee. A contracting model for flexible distributed scheduling. *Annals of Operations Research*, 65:195–222, 1996.

[25] H. Van Dyke Parunak. Manufacturing experience with the contract net. In Michael N. Huhns, editor, *Distributed Artificial Intelligence*, Research Notes in Artificial Intelligence, chapter 10, pages 285–310. Pitman, 1987.

[26] Tuomas W. Sandholm and Victor R. Lesser. Issues in automated negotiation and electronic commerce: Extending the contract net framework. In *First International Conference on Multiagent Systems*, pages 328–335, Menlo Park, CA, 1995. AAAI Press/MIT Press.

[27] Mark S. Fox. An organizational view of distributed systems. *IEEE Transactions on Systems, Man, and Cybernetics*, 11(1):70–80, January 1981. (Also published in *Readings in Distributed Artificial Intelligence*, Alan H. Bond and Les Gasser, editors, pages 140–150, Morgan Kaufmann, 1988.).

[28] Thomas W. Malone. Modeling coordination in organizations and markets. *Management Science*, 33(10):1317–1332, 1987. (Also published in *Readings in Distributed Artificial Intelligence*, Alan H. Bond and Les Gasser, editors, pages 151–158, Morgan Kaufmann, 1988.).

[29] Philip R. Cohen and C. Raymond Perrault. Elements of a plan-based theory of speech acts. *Cognitive Science*, 3(3):177–212, 1979.

[30] Joseph Y. Halpern and Yoram Moses. Knowledge and common knowledge in a distributed environment. In *Third ACM Conference on Principles of Distributed Computing*, pages 50–61, 1984.

[31] Stephanie Cammarata, David McArthur, and Randall Steeb. Strategies of cooperation in distributed problem solving. In *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, pages 767–770, Karlsruhe, Federal Republic of Germany, August 1983. (Also published in *Readings in Distributed Artificial Intelligence*, Alan H. Bond and Les Gasser, editors, pages 102–105, Morgan Kaufmann, 1988.).

[32] Alan H. Bond and Les Gasser. *Readings in Distributed Artificial Intelligence.* Morgan Kaufmann Publishers, San Mateo, CA, 1988.

[33] Michael Huhns, editor. *Distributed Artificial Intelligence.* Morgan Kaufmann, 1987.

[34] Les Gasser and Michael N. Huhns, editors. *Distributed Artificial Intelligence*, volume 2 of *Research Notes in Artificial Intelligence.* Pitman, 1989.

[35] Jeffery S. Rosenschein. Synchronization of multi-agent plans. In *Proceedings of the National Conference on Artificial Intelligence*, pages 115–119, Pittsburgh, Pennsylvania, August 1982. (Also published in *Readings in Distributed Artificial Intelligence*, Alan H. Bond and Les Gasser, editors, pages 187–191, Morgan Kaufmann, 1988.).

[36] Michael Georgeff. Communication and interaction in multi-agent planning. In *Proceedings of the National Conference on Artificial Intelligence*, pages 125–129, Washington, D.C., August 1983. (Also published in *Readings in Distributed Artificial Intelligence*, Alan H. Bond and Les Gasser, editors, pages 200–204, Morgan Kaufmann, 1988.).

[37] Edmund H. Durfee and Victor R. Lesser. Using partial global plans to coordinate distributed problem solvers. In *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*, pages 875–883, Milan, Italy, August 1987. (Also published in *Readings in Distributed Artificial Intelligence*, Alan H. Bond and Les Gasser, editors, pages 285–293, Morgan Kaufmann, 1988.).

[38] Edmund H. Durfee and Victor R. Lesser. Partial global planning: A coordination framework for distributed hypothesis formation. *IEEE Transactions on Systems, Man, and Cybernetics*, 21(5), September 1991. (Special Issue on Distributed Sensor Networks).

[39] M.R. Genesereth, M.L. Ginsberg, and J.S. Rosenschein. Cooperation without communications. In *Proceedings of the National Conference on Artificial Intelligence*, pages 51–57, Philadelphia, Pennsylvania, 1986.

[40] Jeffrey S. Rosenschein and Michael R. Genesereth. Deals among rational agents. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, pages 91–99, Los Angeles, California, August 1985. (Also published in *Readings in Distributed Artificial Intelligence*, Alan H. Bond and Les Gasser, editors, pages 227–234, Morgan Kaufmann, 1988.).

[41] Katia Sycara-Cyranski. Arguments of persuasion in labour mediation. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, pages 294–296, Los Angeles, California, August 1985.

[42] Tad Hogg and Bernardo A. Huberman. Controlling chaos in distributed systems. *IEEE Transactions on Systems, Man, and Cybernetics*, 21(6), December 1991. (Special Issue on Distributed AI).

[43] J. O. Kephart, T. Hogg, and B. A. Huberman. Dynamics of computational ecosystems: Implications for DAI. In Michael N. Huhns and Les Gasser, editors, *Distributed Artificial Intelligence*, volume 2 of *Research Notes in Artificial Intelligence*. Pitman, 1989.

[44] Les Gasser, Carl Braganza, and Nava Herman. MACE: A flexible testbed for distributed AI research. In Michael N. Huhns, editor, *Distributed Artificial Intelligence*, Research Notes in Artificial Intelligence, chapter 5, pages 119–152. Pitman, 1987.

[45] Thomas A. Montgomery and Edmund H. Durfee. Using MICE to study intelligent dynamic coordination. In *Proceedings of the Second Computer Society International Conference on Tools for Artificial Intelligence*, pages 438–444, November 1990.

[46] Edmund H. Durfee and Thomas A. Montgomery. Coordination as distributed search in a hierarchical behavior space. *IEEE Transactions on Systems, Man, and Cybernetics*, 21(6):1363–1378, November/December 1991.

[47] Sandip Sen and Edmund Durfee. Unsupervisesd surrogate agents and search bias change in flexible distributed scheduling. In *First International Conference on Multiagent Systems*, pages 336–343, Menlo Park, CA, 1995. AAAI Press/MIT Press.

[48] K. Sycara, S. Roth, N. Sadeh, and M. Fox. Distributed constrained heuristic search. *IEEE Transactions on Systems, Man, and Cybernetics*, 21(6):1446–1461, December 1991. (Special Issue on Distributed AI).

[49] M. Yokoo, E. Durfee, T. Ishida, and K. Kuwabara. Distributed constraint satisfaction for formalizing distributed problem solving. In *Proceedings of the Twelfth International Conference on Distributed Computing Systems*, pages 614–621, 1992.

[50] Eithan Ephrati and Jeffrey S. Rosenschein. The Clarke Tax as a consensus mechanism amoung automated agents. In *Proceedings of the Ninth National Conference on Artificial Intelligence*, pages 173–178, Anaheim, CA, July 1991.

[51] Sarit Kraus, Eithan Ephrati, and Daniel Lehmann. Negotiation in a non-cooperative environment. *Journal of Experimental and Theoretical Artificial Intelligence*, 3(4):255–282, 1991.

[52] Jeffrey S. Rosenschein and Gilad Zlotkin. *Rules of Encounter*. MIT Press, Cambridge, MA, 1994.

[53] J. S. Rosenschein and G. Zlotkin. Designing conventions for automated negotition. *AI Magazine*, pages 29–46, Fall 1994.

[54] Tuomas W. Sandholm and Victor R. Lesser. Coalition formation among bounded rational agents. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 662–669, August 1995.

[55] Michael P. Wellman. A market-oriented programming environment and its application to distributed multicommodity flow problems. *Journal of Artificial Intelligence Research*, 1(1):1–22, 1993.

[56] Nick R. Jennings. Commitments and conventions: The foundation of coordination in multi-agent systems. *The Knowledge Engineering Review*, 8(3):223–250, 1993.

[57] Yoav Shoham and Moshe Tennenholtz. On the synthesis of useful social laws for artificial agent societies (preliminary report). In *Proceedings of the National Conference on Artificial Intelligence*, pages 276–281, San Jose, California, July 1992.

[58] Milind Tambe. Agent architectures for flexible, practical teamwork. In *Fourteenth National Conference on Artificial Intelligence*, pages 22–28, Menlo Park, CA, 1997. AAAI Press/MIT Press.

[59] Tad Hogg. Social dilemmas in computational ecosystems. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 711–716, August 1995.

[60] Cristiano Castelfranchi. Commitments: from individual intentions to groups and organizations. In Victor Lesser, editor, *Proceedings of the First International Conference on Multi–Agent Systems*, pages 41–48, San Francisco, CA, 1995. MIT Press.

[61] Tuomas W. Sandholm and Victor R. Lesser. Advantages of a leveled commitment contracting protocol. In *Proceedings of the Thirteent National Conference on Artificial Intelligence*, pages 126–133, August 1996.

[62] Munindar P. Singh. A conceptual analysis of commitments in multiagent systems. Technincal Report TR-96-09, Department of Computer Science, North Carolina State University, May 1996.

[63] Philip R. Cohen and Hector J. Levesque. Teamwork. *Nous*, 25(4):487–512, 1991.

[64] Munindar Singh. *MULTIAGENT SYSTEMS: A Theoretical Framework for Intentions, Know-How, and Communications*. Springer Verlag, New York, NY, 1994.

[65] Barbara Grosz and Sarit Kraus. Collaborative plans for complex group action. *Artificial Intelligence*, 86(2):269–357, 1996.

[66] Anand S. Rao and Michael P. Georgeff. Bdi agents: From theory to practice. In *First International Conference on Multiagent Systems*, pages 312–319, Menlo Park, CA, 1995. AAAI Press/MIT Press.

[67] Sandip Sen. Reciprocity: a foundational principle for promoting cooperative behavior among self-interested agents. In *Proceedings of the Second International Conference on Multiagent Systems*, pages 315–321, Menlo Park, CA, 1996. AAAI Press.

[68] Onn Shehory and Sarit Kraus. Task allocation via coalition formation among autonomous agents. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 655–661, August 1995.

[69] Onn Shehory and Sarit Kraus. A kernel-oriented model for coalition-formation in general environments: Implementation and results. In *Proceedings of the Thirteent National Conference on Artificial Intelligence*, pages 133–140, August 1996.

[70] Gilad Zlotkin and Jeffrey S. Rosenschein. Coalition, cryptography, and stability: Mechanisms for coalition formation in task oriented domains. In *Twelfth National Conference on Artificial Intelligence*, pages 432–437. MIT Press/AAAI Press, 1994.

[71] Thomas Haynes and Sandip Sen. Evolving behavioral strategies in predators and prey. In Gerhard Weiß and Sandip Sen, editors, *Adaptation and Learning in Multi–Agent Systems*, Lecture Notes in Artificial Intelligence, pages 113–126. Springer Verlag, Berlin, 1996.

[72] Sandip Sen, Mahendra Sekaran, and John Hale. Learning to coordinate without sharing information. In *National Conference on Artificial Intelligence*, pages 426–431, 1994.

[73] Sandip Sen (editor). Adaptation, coevolution and learning in multiagent systems. Technical Report SS-96-01, AAAI, Menlo Park, CA, 1996. Papers from the 1996 AAAI Spring Symposium.

[74] Sandip Sen (editor). Multiagent learning. Technical Report WS-97-03, AAAI, Menlo Park, CA, 1997. Collected Papers from the AAAI-97 workshop.

[75] Gerhard Weiß. Learning to coordinate actions in multi-agent systems. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 311–316, August 1993.

[76] Gerhard Weiß and Sandip Sen, editors. *Adaption and Learning in Multi–Agent Systems*. Lecture Notes in Artificial Intelligence. Springer Verlag, Berlin, 1996.

[77] B. Chaib-draa. Industrial applications of distributed ai. *Communications of the ACM*, 38(11):49–53, November 1995.

[78] N.R. Jennings. *Cooperation in Industrial Multi-Agent Systems*. World Scientific Series in Computer Science, Vol. 43, 1994.

[79] H. van Dyke Parunak. Applications of distributed artificial intelligence in industry. In O'Hare and N. Jennings, editors, *Foundations of Distributed Artificial Intelligence*. Wiley Inter-Science, 1994.

[80] Munindar P. Singh, Philip E. Cannata, Michael N. Huhns, Nigel Jacobs, Tomasz Ksiezyk, Kayliang Ong, Amit P. Sheth, Christine Tomlinson, and Darrell Woelk. The carnot heterogeneous database project: Implemented applications. *Distributed and Parallel Databases Journal*, 5(2):207–225, April 1997.

[81] T. Finin, R. Fritzson, D. McKay, and R. McEntire. KQML as an agent communication language. In *Proceedings of the Third International Conference on Information and Knowledge Management*, New York, NY, 1994. ACM Press.

[82] G.M.P. O'Hare and N.R. Jennings, editors. *Foundation of Distributed Artificial Intelligence*. Wiley Inter-Science, New York, NY, 1994.

[83] Edmund H. Durfee, Daniel L. Kiskis, and William P. Birmingham. The agent architecture of the university of michigan digital library. *IEE/British Computer Society Proceedings on Software Engineering*, 144(1):61–71, February 1997.

[84] Michael J. Woolridge and Nicholas R. Jennings, editors. *Intelligent Agents: ECAI-94 Workshop on Agent Theories, Architectures, and Languages*. Springer-Verlag, New York, NY, 1995.