AI Magazine Volume 22 Number 2 (2001) (© AAAI)

Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence

A Review

Julie A. Adams

Society has become more dependent on automated intelligent systems; at the same time, these systems have become more and more complicated. Society's expectation regarding the capabilities and intelligence of such systems has also grown. We have become a more complicated society with more complicated problems. As the expectation of intelligent systems rises, we discover many more applications for AI. Additionally, as the difficulty level and computational requirements of such problems rise, there is a need to distribute the problem solving. Although the field of multiagent systems and distributed AI is relatively young, the importance and applicability of this technology for solving today's problems continues to grow.

As the title indicates, *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence* covers the design and development of multiagent and distributed AI systems. The purpose of this book is to provide a comprehensive overview of the field. It is an excellent collection of closely related papers that provides a wonderful introduction to multiagent systems and distributed AI. The book provides not only basic introductory information but also detailed discussions covering the important topics in the field, practical examples and applications, and a section dedicated to the relationship between multiagent systems and various other research areas. This book compiles the important concepts and methodologies required to develop a multiagent system in an understandable, and comprehensive, manner. Not only does the book focus on the known solutions and issues, it also discusses the open questions and dilemmas.

The prologue begins by defining distributed AI as "the study, construction, and application of Multiagent Systems, that is, systems in which several interacting, intelligent agents pursue some set of goals or perform some set of tasks" (p. 1).[1] The collection of papers that compose this book fully supports this definition.

Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence, Michael Wooldridge, The MIT Press, Cambridge, Massachusetts, 1999, 619 pages, ISBN 0-262-73131-2. $60.00.

The design of a system that uses a single agent implies certain requirements, such as the agent's ability to perceive the environment and maintain knowledge about it, reason about said environment, and execute particular actions to solve tasks. The design of a single-agent system, although not simple, is considerably less complicated than that of a multiagent system. Multiagent systems require additional considerations, for example, communication mechanisms, environmental and world knowledge maintenance as well as communication, and societal issues such as which agent is assigned to a particular task. These additional requirements ensure that the design of a multiagent system is much more complex than that of a single-agent system. Many of these considerations are not new. Societal issues and those of distributed communication protocols have existed for years. Each of these topics has been studied in other domains, such as operating systems, networking, and urban development.

Although the multiagent system is significantly more complex, the usefulness of such a system is larger than that for a single agent. In fact, many of the problems that need to be solved in developing multiagent systems must also be solved for domains such as multiple robotic systems, automated distributed transportation systems, aviation situations such as free flight, and even satellite business offices. Multiple robotic systems in some cases can be thought of as a multiagent system. Developments in multiagent systems can be applied to situations in which boat, air, truck, and rail transportation are all utilized to transport goods from one location to another. In the aviation domain, and particularly free flight, similar communication issues exist concerning what information to communicate, when to communicate, and how to resolve conflicts. Finally, a business with a number of satellite offices throughout the country or world faces similar issues in communications, task allocation, and collaboration.

Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence presents the required individual components of the multiagent system and explains the importance of each topic. The structure of the book makes it useful for practitioners from AI and general computer science as well as other areas such as aviation, transportation, and business. The book presents the basics of all the components required to build a multiagent system.

The topics range from the definition of an intelligent agent to distributed problem solving, search algorithms, decision making, learning, and communications. This book is not an individual practitioner's approach to the problem but an introduction to the topic, covering known solutions.

Multiagent Systems begins by providing the necessary background to embark on a more in-depth study of distributed AI. The reader is first provided a well-written introduction regarding the design and implementation of intelligent agents.

An important element of this discussion is Michael Wooldridge's statement that "there is no universally accepted definition of the term agent, and indeed there is a good deal of ongoing debate and controversy on this very subject" (p. 28). There are, in fact, many interpretations of the term *agent*. Maes (1995) defines an autonomous agent to be a computational system that inhabits some complex dynamic environment, is able to sense and act autonomously in the environment, and is able to obtain a series of goals or tasks. Russell and Norvig (1995) define an agent as "anything that can be viewed as perceiving its environment through sensors and acting upon that environment through effectors" (p. 33). Hayes-Roth (1995) defines intelligent agents as entities that continuously perceive a dynamic environment, reason about and interpret their perceptions, solve problems, and determine actions.

Wooldridge defines an agent as "a computer system that is situated in some environment, and that is capable of autonomous action in this environment in order to meet its designed objectives" (p. 29). He further defines an intelligent agent to be "capable of flexible autonomous action in order to meet its design objectives" (p. 32). *Flexible* is interpreted to imply that the agents are reactive to their environment, are able to "exhibit goal-directed behavior," and are able to interact with other agents.

The first chapter ends by providing the reader with various intelligent agent architectures, including reactive, logic based, belief-desire-intention, and layered. Each architecture description includes a well-written definition as well as the pros and cons.

Once one understands intelligent agents, it is natural to consider many agents working together. Clearly, a single agent works within an environment, but the question becomes, How does this environment and the tasks change when there are multiple agents? If there are a large number of agents in the system, does this collection of agents form a society? Such questions are important in distributed AI and should be addressed early when discussing this topic. The discussion of multiagent systems and societies of agents addresses these questions by familiarizing the reader with a comprehensive compilation of multiagent environment characteristics. The discussion then shifts to agent communications, covering the topics of coordination, message meaning, message types, level of communication, and communication formats. All these topics are important introductory concepts. Huhns and Stephens also provide an overview of agent interaction protocols that includes coordination protocols, cooperation protocols, contract net, blackboard systems, negotiation, and belief maintenance. Their chapter provides the motivation for the remaining chapters with a discussion of the need for multiagent systems.

In general, this introductory material provides the necessary background, and it prepares readers for the remainder of the book and their continued study of the topic. The presentation of this information is such that even a student new to the subject should be able to understand the basic concepts. The book continues with a concentration on the various problem domains within distributed multiagent systems.

The first of the more specific topics is distributed problem solving and planning. Distributed problem solving is a fundamental topic in a discussion of multiagent systems. The point of the system is to solve a problem using multiple agents that are distributed throughout the environment. Planning is also a fundamental topic: Important not only in a single-agent system, it becomes even more so in a multiagent system. The agents' capabilities, current availability, and intentions must be considered when multiple agents are meant to solve a problem. If agent 1 is currently busy completing a task, then this same agent cannot necessarily be available to assist agent 2 with a new task. Durfee begins his coverage of this topic with a discussion that attempts to solve a problem in a distributed system. The discussion entails task and result sharing between agents. One of the primary concerns is the relationship between communications and load sharing to accomplish the task. The discussion naturally flows into one about distributed planning using centralized as well as decentralized schemes.

The next major topic is agent search algorithms. Yokoo and Ishida make the point that "search is an umbrella term for various problem solving techniques in AI" (p. 165), thus setting the tone for the discussion. In principle, the search task is similar in concept between single-agent tasks and multiple-agent tasks. The difference lies in the complexity of the task. As more agents join the system, the search task becomes more complicated, potentially slowing the search capability. If the multiple agents are expected to solve problems and complete tasks in real time, then the appropriate search algorithms that provide real-time functions for single agents must be extrapolated and modified for multiple- agent systems. The authors discuss constraint-satisfaction problems, path-finding problems, and two-player games. Each type of problem, as well as the potential algorithms, is discussed in detail.

Decision making is a difficult enough problem when only one entity is expected to come to a decision. When there are multiple entities participating in the decision process, the problem becomes more complicated. The individual entities can have their own individual needs and goals. In a multiple-agent system, individual goals and needs can undermine the goals and needs of the overall system. The system must combine such goals and needs into a rational solution that is appropriate for the multiagent system but might not meet all the needs and goals of the individual agents. The

task of reaching a decision when multiple entities are involved is significantly more complicated than a single entity reaching a decision about its current task. Sandholm authored the presentation on distributed rational decision making. The primary discussion pertains to the necessary elements of agent negotiation, beginning with a description of evaluation criteria for negotiation protocols. The ensuing discussion concerns negotiation by voting, auctions, bargaining, marketing mechanisms, contract nets, and coalition formation. The discussion reviews both the positive and the negative aspects of each protocol.

The ability of a system to learn about the environment in which it is working and the task it is assigned is important. The system should able to build on itself and its capability if it can successfully solve new problems. In a single-agent system, questions regarding how to acquire knowledge, where to store knowledge, how to store knowledge, when to access knowledge, and so on, are fairly straightforward. This single agent is only as intelligent as its capabilities to perceive, interpret, and learn about the environment and task. When the problem is expanded to a multiple-agent system, then these questions become significantly more difficult to answer and implement. If the individual agents maintain their own knowledge, then the other agents do not necessarily have access to this information and, therefore, must learn it themselves. If the agents are permitted to communicate to each other the results of their learning, then there is the possibility that the information provided is outdated or incorrect because of changes in the dynamic environment. An alternative is to provide a central location for all knowledge, but there are associated problems with this method as well. One of the largest issues is the communication bottleneck between the central knowledge store and the individual agents. Sen and Weiss begin their discussion of multiagent system learning by generally defining the important differences that occur between learning mechanisms. They then describe the algorithms available for agents to learn how to coordinate their activities and the agent coordination required by these algorithms. They continue with learning about and from the other agents to improve individual performance. The all-important topic of coordinating learning with communication, including reducing communication by learning and improving learning through communication, is also covered.

Each of these discussions provides the proper background and understanding for the reader to continue with the material on computational organization theory. Carley and Gasser begin by defining organization in the context of multiagent systems and proceed to define computational organization theory. They discuss how agents, organization design, the tasks, and technology can be applied to model multiagent systems. Organizational dynamics are discussed as well as methodological issues. The authors also provide an overview of the available tools for developing such models. An interesting aspect of this discussion is the numerous parallels in the application of computational organization theory to all types of organization, both artificial and real.

Formal methods based on logic-based representations have been an important element in ensuring correctness when creating systems. The purpose of such representations is to provide a high-level understanding of the systems to be created. The representations can be used to debug and validate systems. Because multiagent systems is a relatively new field, the development of such representations is necessary to ensure that existing and future systems are correct. Such representations can assist future researchers when establishing starting points for their own work. The final discussion of the more detailed topics focuses on logic-based representation and reasoning for distributed AI. Singh, Rao, and Georgeff do an excellent job of giving the reader the necessary background material to understand this important topic. The presentation provides a detailed introduction to the cognitive aspects involved with multiagent systems, specifically, belief-desire-intention, agent coordination, agent communications, and social primitives. This section also presents the various tools and systems for distributed AI, although this discussion seems a bit disjoint from the previous presentation.

A very practical and necessary chapter presenting industrial and practical applications of distributed AI then follows. Parunak has done an excellent job of providing emphases to applications beyond the academic research labs. He discusses why and where distributed AI should be used in industry as well as the constraints of agent-based system development in the industrial setting, constraints that are not present in most research development environments. Finally, he discusses commercially available development tools and their applicability to certain types of agent application.

The final section of the book presents related topics and indicates their relationship to multiagent systems and distributed AI. The topics are groupware and computer-supported cooperative work, distributed decision-support models, concurrent programming, and distributed control algorithms.

The editor, Gerhard Weiss, and the individual authors have done an excellent job of integrating each individual contribution. They have ensured a high level of term and topic coordination across each topic. Such integration and coordination are uncommon in compilations. This cooperation has led to the comprehensive coverage of the topic in a manner that should appeal to many readers.

In general, almost any reader is likely to find something interesting in this book. Individuals working in the field of multiagent systems and distributed AI will find this book contains a complete, comprehensive coverage of the field. Individuals who have an interest in learning more about multiagent systems will find this book appealing because of the introductory nature of the first few chapters. Even those interested in the general field of AI but not necessarily multiagent systems should find this book of interest. The book demonstrates many links between agent systems and the more general fields of AI such as planning, learning, and perception.

This book is also an excellent choice

as a course textbook. The gentle introduction to multiagent systems provides an excellent basis for more senior-level undergraduate students and any graduate student. The book then moves into more detailed issues and concerns that take the student beyond the basic capabilities.

Each chapter contains a series of exercises that allow the reader to fully test his/her understanding of the topic. These exercises have various levels of difficulty, from those intended to determine comprehension of the particular topic through programming assignments to open research questions. The inclusion of such a broad spectrum of exercises supports a goal of the editor to make the book a possible teaching text. However, although the book does contain excellent exercises, it does not provide solutions for the exercises nor is there an associated web page (that I could find) for solutions. This omission is an important one. The book also contains a comprehensive glossary that was developed through the joint effort of all the authors. Because the authors consider this book a possible teaching text, the glossary is a useful feature.

Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence is not simply a collection of papers on a similar topic. This book provides a clear, well-integrated introduction to multiagent systems and distributed AI. A student with even a brief introduction to AI can read this book and understand the contents. The real application examples help the reader understand how the concepts can be applied and provide the necessary groundwork for the provided exercises. I recommend this book to anyone with an interest in multiagent systems and distributed AI.

Notes

1. The authors use the terms *multiagent system* and *distributed AI* interchangeably.

References

Hayes-Roth, B. 1995. An Architecture for Adaptive Intelligent Systems. *Artificial Intelligence* (Special Issue on Agents and Interactivity) 72(1–2): 329–365.

Maes, P. 1995. Artificial Life Meets Entertainment: Life-Like Autonomous Agents. *Communications of the ACM* 37(7): 108–114.

Russell, J., and Norvig, P. 1995. *Artificial Intelligence: A Modern Approach*. Englewood Cliffs, N.J.: Prentice Hall.

Julie A. Adams received a B.S. and a B.B.A. from Siena College in 1989 and 1990, respectively, as well as an M.S.E. and a Ph.D. from the University of Pennsylvania in 1993 and 1995, respectively. She joined the faculty of the Department of Computer Science at Rochester Institute of Technology in September 2000. Previous industrial experience includes Eastman Kodak Company (1996–2000) and the Honeywell Technology Center (1995–1996). Her research interests include complex man-machine systems, multiagent systems, robotics, and AI. Her e-mail address is jaa@cs.rit.edu.