DISTRIBUTED ARTIFICIAL INTELLIGENCE

Deshmukh A.
University of Massachusetts-Amherst, USA

Keywords: multi-agent systems, distributed problem-solving, distributed decision-making, cooperation, negotiation, coordination.

Contents

1. Introduction
2. Classification of DAI Issues
3. Multi-Agent Architectures and Environments
4. Distributed Problem Solving
5. Distributed Decision Making
6. Applications
7. Conclusions
Glossary
Bibliography
Biographical Sketch

Summary

Distributed artificial intelligence (DAI) has emerged as a powerful paradigm for representing and solving complex problems. The growth of this field has been spurred by the advances in distributed computing environments and wide spread information connectivity. Although DAI started as a branch of artificial intelligence over twenty-five years ago, it has emerged as an independent research discipline in its own right, representing a confluence of ideas from several disciplines.

This paper presents an overview of the state-of-the-art in DAI and multi-agent systems. Our goal is to introduce the reader to the main concepts and principles related to DAI. Since DAI is an active and rapidly expanding research area, we do not intend to give an exhaustive list of all the research projects. However, this paper should help the reader grasp the fundamental issues and orient them towards relevant research in the area of interest. We discuss key issues related to DAI and classify them into common research areas. We also present an overview of application domains where agents have been used to improve the system performance.

1. Introduction

Distributed artificial intelligence (DAI) is the science of distributing, coordinating and predicting the performance of tasks, goals or decisions in a multiple agent environment. DAI, which started as a branch of artificial intelligence about twenty-five years ago, has grown into an independent research discipline representing the confluence of ideas from artificial intelligence, economics, psychology, sociology, operations research and organizational theory. DAI research has focused on modeling the communication, knowledge and decision making processes that are necessary to sustain societies of

computational agents or amalgamation of people and computers in information economies. This research can be broadly classified into two areas: distributed problem solving and multi-agent systems. Distributed problem solving deals with the issues related to solving a problem by dividing it among a number of cooperative problem solvers which share the computational burden and knowledge of their partial solutions. The interaction between individual nodes or problem solvers is explicitly defined in distributed problem solving environments. Multi-agent systems on the other hand are concerned with the behavior of loosely coupled problem solvers, or agents, that work together to solve a problem beyond their individual capabilities. These independent problem solvers, or computational agents, are autonomous entities that possess the ability to intelligently operate under diverse environmental conditions given their sensory and effectual capabilities. The key concepts of agency are situatedness, autonomy, flexibility and interaction. Since the DAI field is closely associated with the notion of agents, it is often referred to as multi-agent systems research in literature. Both the names will be used in the following discussion interchangeably.

An important characteristic of DAI is that it combines computational resources of a group of agents such that the group intelligence is more than the sum of the individual agents' capabilities. The presence of intelligence in individual agents, however limited, is implicit in this description. The agents use this intelligence, or ability to make informed decisions based on a model of the world and the available information, to coordinate, cooperate, negotiate or compete with other agents. Note that a central controller does not mandate the interactions among agents. The agents may possess cooperative or competitive personalities, based on the application context and the overall goal. For example, modeling resources in an organization with the goal of improving the profitability of the organization may require cooperative agents, whereas, agents representing a market economy, with independent producers and consumers, may be best modeled by competitive agents. Early DAI research had primarily focused on cooperative agents. Systems such as Hearsay-II and ACTOR demonstrated the use of unintelligent agents to achieve group intelligence via concurrency and close coupling in information transfer. Contract Net protocol was the first framework to demonstrate the use of loosely coupled agents to perform tasks, which were too big for any individual agent. This was achieved through the task decomposition process and allocation of sub-tasks to a network of agents. This protocol used contracts between different agents to allocate tasks and coordinate activities. It is important to note that the Contract Net protocol can be applied to both cooperative and competitive agents.

Multi-agent systems are poised to make a significant impact on future computer applications and on how new problem solving or decision architectures are structured. With the advent of distributed computing platforms, with open architectures and heterogeneous operating environments, the concept of agents is clearly gaining importance. Multi-agent systems provide a promise for managing large, geographically distributed networks of information and computational resources efficiently. Another important reason for studying multi-agent architectures is that it parallels the design of human societies. Hence, the study of interaction methods and analysis methods for multi-agent systems can be used to gain a deeper understanding of human societies. DAI offers the promise of higher computational speeds for solving complex problems by using the inherent asynchrony and parallelism in agents. Multi-agent systems are

also more robust to failures since they do not rely on a central controller to coordinate all the tasks. These systems are easily scalable as compared to the monolithic centralized computing environments. Moreover, the components can be reusable, and large multi-agent systems can be designed and tested in parts before deployment.

DAI or multi-agent systems is a rapidly growing research field and this paper attempts to capture the state-of-the-art in the area at this point in time (2002). Section 2 presents a classification of different issues related to multi-agent systems. The following sections discuss DAI issues in detail. Section 3 presents an overview of multi-agent architectures and environments. Section 4 discusses the issues related to distributed problem solving architectures. Section 5 focuses on the area of distributed decision-making and gives an overview of the methods used to aggregate decisions in a multi-agent environment. Section 6 lists the various application areas where DAI has been used effectively. Finally, Section 7 presents a summary of the major topics discussed in the paper and discusses future research directions for the field.

2. Classification of DAI Issues

Design and operation of DAI systems, which perform as expected, present significant challenges. Since the system control is distributed, each agent has limited capabilities and information, data is decentralized, and the computations are asynchronous, the only way to guarantee desired performance of DAI systems is to allow appropriate interactions among agents. Hence, the issues of which agents should interact, how they should interact, what should be their strategies and how to resolve any conflicts that may arise are crucial for DAI systems.

Given the multidisciplinary nature of DAI, there is significant diversity in the issues addressed by the researchers in the field and the tools used. Hence a classification of the issues is useful to the readers who want to get an overview of the state-of-the-art in the field. The issues related to DAI can be studied from different perspectives:

    a) Individual perspective: This perspective focuses on the issues related to the design and understanding of individual agents, decision makers or problem solvers. The characteristics of agents, such as reactive, deliberate, intentional, cooperative, self-interested and social, are studied under these issues. Each agent's internal knowledge structure and knowledge maintenance strategies also fall under this category. Agents also need the ability to reason about their own actions and the needs of other agents. Hence, development of reasoning mechanisms, either logic based or utility based, are important from individual perspective. Finally, the agents need to adapt to changing situations. Agent evolution or adaptation schemes are based on passive or active learning schemes.

    b) System perspective: The system level or group perspective is concerned with issues related to the performance of the entire organization. Multi-agent systems can be approached in two ways: bottom-up methods, where individual agent capabilities are identified that result in appropriate interactions at the group level; and top down methods, where group level norms are developed which constrain the agent level interactions in order to achieve the desired performance. Organizational structure plays an important role in system level

behavior. The organizational structure is not fixed in most DAI systems, hence the mechanisms for updating the roles and relationships between agents are important areas of research. Agents use coordination methods, negotiations, distributed planning schemes and cooperation strategies in order to achieve group level coherence. The issues related to information exchange are intimately tied with these strategies.

c) Implementation perspective: The implementation perspective deals with issues related to building multi-agent systems, test-beds, design tools and applications. DAI system designers have to ensure interoperability, ability to deal with spatial distribution, varying degrees of agent autonomy, ability to function as an open system and domain independence when developing industrial applications.

From the above discussion it is evident that the following problems are associated with the design of any DAI system, and hence are areas of active research in the multi-agent community:

a) How to formulate multi-agent architectures? The research in this area focuses on dividing larger tasks into subtasks that can be efficiently processed by multiple decision-makers or problem solvers. The crucial issues in creating DAI systems are:

- Task decomposition and allocation of subtasks to individual agents. The results of each sub-task also need to be aggregated in order to obtain the desired final goal.
- Reasoning by agents about their own tasks and actions, representation of their intent to other agents, and interaction between different agents.
- Recognition of conflicts or inconsistent goals of different agents and appropriate methods for their resolution.
- Coordination of interaction so that the group works as a coherent team to achieve a larger goal.

b) How to enable agents to communicate with each other? Communications can occur at various levels, from passing data between multiple agents to sharing intent and negotiating over complex organizational issues. The main issues related to communication in multi-agent systems are:

- Lower level information and data transfer protocols.
- Languages which specify context and intent of the messages
- Deciding when to communicate and how much information to communicate.
- Selection of different forms of information transfer, such as blackboards, focussed addressing, multi-casting or global broadcasting.

c) How to ensure coherence in local decisions made by agents or problem solvers in order to guarantee desirable global outcomes? The primary areas of research related to this issue are:

- Using organizational structures to guide the local decision-makers.
- Using negotiations and contracting mechanisms to coordinate activities in order to improve system level objectives.
- Developing coordination mechanisms, which consider task representation schemes of other agents in order to schedule local tasks.
- Formulating appropriate reward or utility functions for self-interested agents in order to facilitate coherence in their actions.

- Studying effects of multi-agent learning on the overall functioning of the system.

d) How do we measure performance of DAI systems? Understanding the behavior of distributed systems presents significant challenges, especially since the global behavior is an emergent property of the agent interactions. The key research topics in this area are:
   - Understanding the dynamics of multi-agent interactions and developing methods to mitigate unsteady or chaotic behavior in the system.
   - Developing formal methods to describe multi-agent systems. These methods also facilitate diagnosis of interaction protocols and system level performance estimation.
   - Incorporating uncertainty in agent decision processes in order to improve the performance of multi-agent systems in an open environment.

e) Which application domains are best suited for DAI? How to design, implement and operate practical DAI systems? The issues related to application of DAI systems are:
   - Design of computational and agent development platforms for rapid deployment of DAI systems.
   - Documenting savings due to multi-agent approaches in real world applications.
   - Distributing decision-making capabilities or problem solving requirements among agents in an application domain.

The following sections discuss some of the fundamental DAI issues in detail. Our intent is not to provide an exhaustive list of all the systems and tools developed. This discussion should help the reader grasp the fundamental aspects of DAI research and orient them towards the current research in their areas of interest.

-
-
-

TO ACCESS ALL THE **15 PAGES** OF THIS CHAPTER,
Visit: http://www.eolss.net/Eolss-sampleAllChapter.aspx

Bibliography

Bond A. H. and Gasser L., editors (1988) *Readings in Distributed Artificial Intelligence*, Morgan Kaufmann, San Mateo, CA. [Collection of early research papers on different aspects of DAI].

Bradshaw J. M. editor (1997) *Software Agents*, The MIT Press, MA. [A recent compilation of state-of-the-art research in multi-agent systems].

Chaib-draa B., Moulin B., Mandiau R. and Millot P. (1992) "Trends in distributed artificial intelligence", *Artificial Intelligence Review*, vol. 6, no. 1, pp. 35-66. [Presents the prevalent directions in DAI research community].

Durfee E. H. (1988) *Coordination of Distributed Problem Solvers*, Kluwer Academic Publishers, Boston, MA. [Presents a distributed problem solving environment and several applications of DAI in that area].

Jennings N. R., Sycara K. and Wooldridge M. (1998) "A Roadmap of Agent Research and Development", *Autonomous Agents and Multi-Agent Systems*, vol. 1, pp. 7-38. [Current state of the research and future directions for agent research].

O'Hare G. M. P. and Jennings N. R., editors (1996) *Foundations of Distributed Artificial Intelligence*, John Wiley & Sons, NY. [First comprehensive overview of DAI research issues].

Weiss G. editor (1999) *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*, MIT Press, Cambridge, MA. [Most up-to-date review of multi-agent systems research].

Biographical Sketch

Abhijit V. Deshmukh is on the faculty of the Department of Mechanical and Industrial Engineering at the University of Massachusetts, Amherst. He has received a Ph.D. in Industrial Engineering from Purdue University. His research interests are in design and analysis of distributed systems. He is a member of IIE, INFORMS and ASME.