Foreword

Multi-agent systems (MAS) have emerged as a new methodology to address the issues in organizing large-scale software systems. This methodology provides a conceptual model that helps maintaining constraints, a task conventional software engineering is unable to achieve. In recent years, MAS has been used in various areas in computer science and engineering and is becoming a versatile tool that addresses software engineering needs. It also extends the spectrum of computer science research and has drawn more and more attention to a wide range of areas from theoretical studies to practices.

An agent is a software entity that actively seeks ways to complete its tasks. Intelligent agents have the ability to gain knowledge through their problem-solving processes. The study of social behaviors of agents in cognitive science is an important part of the intelligent agent field. Software agents, on the other hand, focus on interaction and collaboration to achieve the goals in a context that changes in a usually unforeseen manner. The necessity of using agents arises from the complexity of large software systems, which bring about design issues that conventional software engineering technology fails to tackle. For instance, mobile agents were proposed to address the needs in the client/server model for the client to be able to migrate to the server side to perform the operation that passive message-passing mechanisms cannot handle efficiently. In a dynamic distributed system, agents with self-adjusting ability can simplify the system architectural design. The design of such a system may be exceedingly complicated in traditional software architecture frameworks or object-oriented modeling.

Agent-oriented modeling yields an unconventional approach to system design, including component definition and system integration. Different applications may impose various requirements on the design and lead to different types of agents. Autonomy is a distinguishing property of an agent. Autonomy entails the agent's capability to survive in a changing environment. An agent has the ability to sense the conditions and make decisions on how to react accordingly. Adaptability requires learning ability necessary for the agent to be able to adjust its decision-making according to past experience. Moreover, an agent-oriented design should address robustness—the system should be reliable when unexpected events occur.

According to their usages, agents can be classified as collaborative agents, interface agents, reactive agents, mobile agents, information agents, and so forth. In the design of large, complex real-life systems, an agent is an abstraction that helps the design of components that tackle different aspects of a problem. Each agent is designed in the most appropriate paradigm for solving its part of a problem. A multi-agent system is used to solve a complex problem that cannot be solved by any single entity of the system. The coordination of independent agents' behaviors is a central part of multi-agent system design. Multi-agent systems are often used to model loosely coupled distributed systems with decentralized control and data allocation. In these systems, communicating agents represent distributed expertise for problem solving. The agents have the ability to process local data efficiently and to communicate with other agents when necessary if the tasks that they are facing are beyond their domain knowledge. Multi-agent systems have been used in a wide spectrum of applications, such as e-commerce, e-learn-

ing, communication, data mining, simulation, robotics, transportation systems, and grid computing. It also initiated theoretical studies in specification and reasoning of MAS systems, languages, modeling methods, knowledge representation and processing, and cognitive sciences.

Multi-agent systems are often classified into two categories based on agents' characteristics: self-interested agents and cooperative agents. Self-interested agents are based on economic concepts in which an agent is assumed as a utility maximiser who always tries to maximize an appropriate utility function. This assumption is widely employed in micro-economics and game theory. Thus, researchers often use economic tools and game theory tools to model agents. Self-interested agents tend to close their private information and fail to react if no benefit is available. Cooperative agents are built so that they engage in behaving cooperatively.

Negotiation is one of the key research topics for multi-agent systems. Due to its nature, agents have to make an agreement and achieve consensus among themselves in order to cooperatively perform shared tasks. Negotiations are often modeled as interaction among self-interested agents. As a traditional result, the contract net protocol has been referred as the basic negotiation mechanism among self-interested negotiation. The contract net protocol is a simple task allocation protocol in which agents announce tasks, make bids, and award the tasks. Nowadays, researchers are trying to build more complex and sophisticated protocols based on auctions. Pareto optimality, incentive compatibility, and revenue optimality are crucial concepts to build such auction mechanisms.

As multi-agent system is a rapidly evolving area, promoting a worldwide exchange among scholars is crucial to helping researchers to rightly position their effort relative to the current trends. This book is one of the early attempts that aim to provide a resource to facilitate research and education endeavors in this area. This book covers a wide spectrum of topics and compiles research results from all around the world. The readers can find new advances in every aspect of current multi-agent systems research, such as mathematical logic, agent-oriented modeling, architectural design, coordination programming, knowledge engineering, machine learning, expert systems, communications, computer networks, parallel processing, grid computing, security, simulation, and robotics. The material in this book covers both theoretical and practical matters of MAS. The publication of this book will doubtlessly benefit the entire community of MAS research. Scientists who work in related areas can use this book as a reference of frontier research themes. University faculty can use this book as supplemental material that reflects trends in advanced research. Managers in industry can use these topics, especially the materials on applications, as a source to find solutions to engineering issues in complex systems.

Takayuki Ito & Chunsheng Yang
February 2007

Takayuki Ito *received a BE, ME, and doctor of engineering from the Nagoya Institute of Technology (1995, 1997, and 2000, respectively). From 1999-2001, he was a research fellow of the Japan Society for the Promotion of Science (JSPS). From 2000-2001, he was a visiting researcher at the University of Southern California Information Sciences Institute. From April 2001 to March 2003, he was an associate professor of Japan Advanced Institute of Science and Technology (JAIST). He joined Nagoya Institute of Technology as an associate professor of Graduate School of Engineering in April 2003. He is a founder of Wisdom Web Co., Ltd., on July 13, 2004. From 2005 to 2006, he is a visiting researcher at the Faculty of Arts and Sciences, division of engineering and applied science, Harvard University, and a visiting researcher at Sloan School of Management, Massachusetts Institute of Technology. He received the Best Paper Award of AAMAS 2006, the 2005 Best Paper Award from Japan Society for* Software Science and Technology, *the Best Paper Award in the 66th Annual Conference of Information Processing Society of Japan, and the Super Creator Award of 2004 IPA Exploratory Software Creation Projects. His main research interests include multi-agent systems, computational mechanism design, group decision support systems, and agent-mediated electronic commerce.*

Chunsheng Yang*'s interests include data mining, case-based reasoning, network privacy and security, intelligent agent system, and multi-agent systems. After receiving his doctorate in September 1995 from National Hiroshima University in Japan, he worked at Fujitsu, Inc., in Japan, as a senior engineer. Since January 1998, he has been with the National Research Council Canada as a research*

officer. Dr. Yang has authored more than 30 papers and book chapters published in the referred journals and conference proceedings. He was the program co-chair for the 17th International Conference on Industry and Engineering Applications of Artificial Intelligence and Expert Systems.

* * *

A person living at the beginning of the 21st century is surrounded by countless hardware and software artifacts. The mainframe computers of the classical era were designed to be shared by multiple users. The personal computers and the corresponding interactive programs assumed a ratio of approximately one user per computer. This was reflected in the interactive applications, which required the constant dialog between the user and the program. Since the turn of the millennium, the ratio had clearly turned to the favor of computers or computer-like devices (such as PDAs, programmable cell phones, or programmable embedded devices).

We cannot afford any more to pay individual attention to every hardware or software component. The only way out from this dilemma is to increase the autonomy of the individual software and hardware components, that is, to use autonomous agent technologies.

While originally proposed as a more flexible answer to the large, monolithic expert systems of the 1980's, autonomous agents have emerged as a discipline of their own, spawning a bewildering array of applications. From the agent-controlled spaceships to the agents participating in the stock market or in online auctions, agents are a regular part of our life. The implementation techniques vary from the highly formal, such as modal logic or partially observable Markov decision processes, to reactive, behavioral models, such as Brooks' subsumption architecture. Biologically inspired models, such as neural networks, swarm architectures, or models of flocking were applied with success in various agent applications.

In short, autonomous agents is a field of high vitality, a melting pot of the most advanced ideas from artificial intelligence, economics, game theory, robotics, simulation, linguistics, biology, and many other fields.

The current book provides the reader with a representative snapshot of the various currents of the field. It is a welcome and timely addition to the library of any developer or scientist who wants to get an understanding of the current frontiers of agent research. It can also serve as the basis for an advanced topics class in autonomous agents. While other books have concentrated on providing an undergraduate level introduction, this book maps the current frontiers of the agent research.

The chapters, written by experts from all over the world, show the diversity of applications and challenges faced by agent researchers. We cannot be but impressed by the range of applications discussed, from mobile robots to grid computing, from management of health care records to supply chain management. The major challenges of the agent field, such as the problems of collaboration, coalition formation, and security, are addressed in various chapters of the book. Finally, the reader will be able to grasp the underlying unity in diversity of the field, how the various techniques are contributing to hardware and software, which can pursue their goals more autonomously, thus become more performant and ultimately more useful to their human users.

Ladislau Bölöni
February 2007

Ladislau Bölöni *is an assistant professor at the School of Electrical Engineering and Computer Science of University of Central Florida. He received a PhD from the Department of Computer Sciences, Purdue University (May 2000). He received a Master of Science from the Department of Computer Sciences, Purdue University (1999) and a diploma engineer degree in computer engineering (Hon.) from the Technical University of Cluj-Napoca, Romania (1993). He received a fellowship from the Computer and Automation Research Institute of the Hungarian Academy of Sciences for the 1994-95 academic year. He is a senior member of IEEE, member of the Association for Computing Machinery (ACM), American Association for Artificial Intelligence (AAAI), and the Upsilon Pi Epsilon honorary society. His research interests include autonomous agents, Grid computing, and wireless networking.*