International Journal of Cooperative Information Systems
Vol. 6, No. 1 (1997) 1–2

FORMAL METHODS IN CIS: MULTIAGENT SYSTEMS GUEST EDITORS' INTRODUCTION

MICHAEL N. HUHNS
Department of Electrical and Computer Engineering, University of South Carolina Columbia, SC 29208, USA

and

MUNINDAR P. SINGH
Department of Computer Science, North Carolina State University Raleigh, NC 27695-8206, USA

There has been much progress recently in cooperative information systems: new problems are being attacked, new techniques are being developed. As the field matures and as its applications become deployed in a wider range of domains and locales, there is increasing interest in formal methods to understand, analyze, design, verify, and validate cooperative information systems. Since the science of cooperative information systems (CIS) unabashedly borrows from several traditional disciplines, the techniques and formal methods that it applies are also of a wide variety.

This issue presents some of the best papers in the multiagent systems subfield of CIS. Even within this subset of CIS, it is interesting to see methods whose traditional counterparts are in software engineering, logics of programs, and distributed computing. The papers reflect a wide range of approaches, and are unified only by their relationship to multiagent systems.

Balzer and Tuomela develop a formal theory, with a number of mathematical results, of the joint intentions of groups of intelligent agents. Intentions, along with other cognitive concepts such as knowledge and belief, have been studied as a means of giving high-level specifications of complex systems. Such concepts have drawn much attention in artificial intelligence and distributed computing. Balzer and Tuomela focus on the joint intentions of agents, which are crucial to understand, analyze, and design cooperative behavior. A particular strength of the proposed approach is its reliance on the plans for actions of the agents involved.

Mazure, Saïs, and Grégoire address the critical problem of determining the consistency of knowledge bases when they are combined or interoperated. It is commonplace that individually consistent knowledge bases might be mutually inconsistent. Inconsistency can be quite damaging, especially in classical logic settings, where one might derive any assertion from inconsistent premises. The authors develop an efficient technique that in many practical situations rapidly zeroes in on the conflicting assertions in different knowledge bases.

Fisher and Wooldridge develop a temporal belief logic for reasoning about multiagent systems. They use this logic to reason about systems developed in a programming language for multiagent systems. Fisher and Wooldridge's approach borrows from the rich tradition of logics of programs developed in distributed computing. Their logic is able to express several important properties of multiagent systems, and to formally verify that those properties are realized in the systems at hand.

Brazier, Dunin-Keplicz, Jennings, and Treur present a high-level framework for modeling multiagent systems. Using this framework, a developer can systematically carry out the conceptual design of a multiagent system. The framework provides abstractions to capture both the capabilities of the constituent agents of a system, and the knowledge and coordination required for them to execute the desired tasks. This framework is applied on a working multiagent system for electrical utility control that has been one of the recent industrial successes of CIS technology.

We expect the papers of this issue to have a great impact on the formal methods in CIS. However, we fervently hope that these papers will not be the last word on their subjects, but will provide a solid foundation for the problems that remain unsolved or even unformulated today.