PANEL SESSION ON DISTRIBUTED ARTIFICIAL INTELLIGENCE

Patrick O. Bobbie

Division of Computer Science
University of West Florida
Pensacola, FL 32514
904-474-2550
(email: pbobbie@uwf.bitnet)

Abstract The next generation of Distributed Computing Systems (DCS) will involve networks of large numbers of *intelligent* agents. Problems would assume different forms of representation and implementation, and would be assigned to different agents in the network. The agents will be expected to act autonomously, competitively, and collaboratively via controlled communication and usage of resources required to complete specific problems.

INTRODUCTION

There is a growing belief that the next generation of *intelligent* distributed computing and distributed AI systems will require a synergistic and complementary effort from DCS and DAI areas. In light of this introspection, the DAI panel discussion focuses on two broad, complementary areas. Thus, the DAI discussion strives for two major goals, viz. to provide a forum for information sharing and, to a large extent, bring out (1) the problems, methodologies, and solutions in these two important areas and (2) DCS and DAI research communities to discuss and understand the intricacies of intelligent cooperative problem-solving.

FOCI OF PANELISTS

Michael Huhns of MCC, Austin, Texas, elaborates on DAI as a generalization of distributed computing in which task allocation is based on the knowledge, intelligence, and reasoning capability of nodes, rather than just on the availability of computational resources. The panelist views DAI as a more general discipline in that it can be heuristic, can involve separate planning and execution phases, and sometimes can require negotiation among the nodes. Furthermore, the nodes can be cooperative or competitive. In spite of these distinctions, DAI requires the results and techniques of distributed computing in order to be efficiently and easily implemented. In sum, the viewpoint will be focused on the use of knowledge, intelligence, and reasoning capabilities to support the planning, execution, and negotiations of distributed computing nodes (in solving generalized computing problems).

Behrooz Shirazi of UT Arlington, Texas, focuses on heterogeneous learning systems involving the combination of neural and symbolic processing. Specifically, the panelist elaborates on (1) symbolic systems (e.g. knowledge-bases plus rule-based inference and learning) which are quite effective when all the features of a physical system they model are exactly and precisely known and (2) the connectionist (neural net) model which has a self-learning capability and can process incomplete or noisy data. The panelist proposes an integration of these two models through an evolutionary learning technique.

Edmund Durfee of Univ. of Michigan, AI Lab., Ann Arbor, focuses on mechanisms for synchronizing and controlling cooperative / competitive agents in a distributed system. In the purview of the panelist, DCS generally assume that distributed tasks are independent or are related in simple ways such as through precedence orderings, while research in DAI has had to examine much more diverse forms of relationships between the activities at different sites. The panelist highlights some techniques developed in DAI that promote this type of dynamic coordination, with a view both to how these techniques can become part of DCS in general, and place demands on the underlying DCS for supporting dynamic coordination.

Amit Sheth of Bellcore, Piscataway, NJ, discusses the issue of cooperative problem-solving in the context of distributed information and database systems. Thus, in many complex and distributed applications, cooperative tasks (agents) that achieve cooperative objectives through controlled access to shared objects, can be used as a modeling tool. The discussion focuses on and draws examples from multimedia systems, a model that allows a uniform representation of application dependent and independent correctness criteria, and an interactive transaction model that achieves cooperative objectives by controlled manipulation of objects in a shared object database.

Mike Papazoglou of Australian National Univ., Canberra, discusses distributed computing environments in which work is executed by "intelligent" agents acting autonomously, cooperatively, or collaboratively, depending on the resources required to complete a common task. A goal of this vision is to be able to efficiently and transparently use computing resources such as data and knowledge that are available on all computers in large computer / communications network. The panelist discusses the implication that this type of collaborative work has on integrating distributed information (i.e., Data + Knowledge) systems and emphasizes how it is related to DAI.

11TH INTERNATIONAL CONFERENCE ON DISTRIBUTED COMPUTING SYSTEMS

Arlington, Texas, USA
May 20 - 24, 1991

Sponsored by the IEEE Computer Society Technical Committee
on
Distributed Computing
Fort Worth, Texas Section

IEEE Computer Society Press
Los Alamitos, California

Washington • Brussels • Tokyo