Multiple Approaches to Multiple Agent Problem Solving

James Hendler (Chair)
Computer Science Department
University of Maryland
College Park, Md. 20742, USA

Danny Bobrow
Xerox Palo AIto Research Center, USA
bobrow.paflxerox.com

Les Gasser
Ecole des Mines, France
gasser@usc.edu

Carl Hewitt
MIT, USA
hewitt@ai,mit.edu

Marvin Minsky
MIT, USA
minsky@ai.mit.edu

Tick up any modern operating systems textbook, and you will find sections on distributed processing - the sharing of computation among multiple physical processors. Familiarity with this literature brings the reader into contact with a jargon filled with the sorts of terms that computer science thrives on: *load balancing, network topology, routing strategics, circuit switching, collision detection, job migration,* and the like. The discussion centers on how the bits of data that make up a computation can be physically moved among machines.

If however, one picks up a book on distributed AI processing[1], the papers are very different. Instead of solely seeing the technical jargon of computer science, the pages are full of terms borrowed from sociology: *negotiation, interaction, contracts, agreement, organization, cohesion, social order, and collaboration,* to name but a few. This is not to say that serious AI scientists aren't concerned with the issues in the underlying computation, but rather that concentrating on the issues involved in determining how a set of separate agents can give rise to global "intelligent" behavior forces the researcher to go beyond consideration of the computation itself.

It is clear that an understanding of how multiple agents can jointly achieve the solution of complex problems demands insights into the bases of communication and social organization. For example, consider a group of humans working together to build a house. Such an endeavour requires a number of instances of cooperation and organization, in particular, the communication of information between participants and the organization of the activity so that specialists in one area (electricians, carpenters, architects) can cooperate with those working in another, all organized so as to achieve a shared goal. These same features emerge, however, as a swarm of bees work together to establish a hive. The similarities and differences between these situations brings us to the topic of this panel: what .are some of the differences in perspective arising from differing views of multiple agent problem solving? The goal of this panel is to bring together researchers attacking multiagency in different ways and to explore the similarities and differences between their approaches.

The panelists represent a set of very different, approaches to understanding multiagent problem solving. These include:

- *Collaborative. Work:* Danny Bobrow has been studying issues involved in enabling a set of humans to work together in the solving of complex problems. His work involves both the observation of humans attacking complex problems across a wide spectrum of activities and an examination of computer systems aimed at supporting such collaborative work A large part of this work focuses on the establishment of a "common ground" for the agents involved in the problem solving — no agent can fully describe its languages or its assumptions to another, and yel they must work together with shared assumptions if they are to solve problems without conflict. Theories must be developed to explain bow actors can come to establish mutually compatible languages, sets of assumptions (e.g., common defaults), and other common bases for interaction given the absence of global semantics.

Social Conceptions of Knowledge and Actum: Les Gasser's work focuses on "identifying the principles that underly a fundamentally multi-agent (social) conception of action and knowledge for Distributed AI research" (Gasser, 1991). This work investigates foundational issues in understanding and representing the existence and interaction of multiple agents. Important issues include how to integrate the multiple perspectives arising among different participants, how to integrate the situated

for example, *Readings in Distributed Artificial Intelligence,* A. Bond and L. Gasser (eds.), Morgan-Kaufman, 1988.

information arising from an individual's particular problem-solving context into the group's more general knowledge, and how to understand the relationships between the autonomy of an individual agent and an externally imposed social order.

- *Open Information Systems Semantics:* Carl Hewitt's work focuses on the building of large scale open information systems (OIS), those "which are always subject to unanticipated outcomes in their operation and which can receive new information from outside themselves at any time." One major aspect of this work is the development of a semantics for such systems. The work involves the development of a framework in which concepts such as commitment, cooperation, negotiation, conflict, and distributed problem solving can be made rigorous. In addition, his work raises the question of how the important AI concept of *deduction* is changed when it is viewed in a rnultiagent framework.

- *Abstraction and Reaction:* James Hendler has been investigating planning systems which can react to dynamic change in the world. His system is based on the use of a number of asynchronously processing "monitors" which observe conditions in the world and report their observations to effectors and to other monitors. The monitors, arranged in an abstraction hierarchy, thus form a distributed network of cooperating entities, in essence a social order is imposed on a network of agents with very limited autonomy. One of the key questions under consideration is how this sort of multiagency can achieve complex problem solving tasks.

- *The Society of Mind:* Marvin Minsky's work looks at rnultiagent problem solving from a very differ ent perspective. In his model "mind" (or thought) arises from the interaction of a myriad of simple agents, none of which by themselves can be characterized as having thoughts. He raises questions as to how these agents function, how they are embodied, how they communicate, what happens when they "disagree," how they are differentiated, and, most importantly of all, how such a group of agents can have a competence far beyond that of any of the single agents.

The panelists will address basic questions as to how their different perspectives on the nature of the agents, the means of communication, and the social organization of agents affect their views of rnultiagent problem solving. Issues to be addressed include:

1. What implications result from the differing views of the "granularity" of agents (from individual humans through mindless agents)?
2. What are the tradeoffs between autonomy and order in the different models?
3. What does each approach have to say about the basic nature of multi-agent communication?
4. In each model, what information does an agent need to know about itself? About others?
5. How does each account for the arising of globally desirable behavior from interactions of between independent individual agents?

References

Bobrow, D, Fogelsong, D. and Miller M, "Definition Groups: making sources into first-class objects," in Wegner, P and Shriver, B (editors) *Research directions in object-oriented programming,* MIT Press 1987

Bond, A. and Gasser,L. eds. *Readings in Distributed Artificial Intelligence.* Morgan Kaufman Publishers, San Mateo, CA., 1988.

Gasser, L. "Social Conceptions of Knowledge and Action: Distributed Artificial Intelligence and Open Systems Semantics," *Artificial Intelligence,* January/February 1991.

Gasser, L. "The Integration of Computing and Routine Work," *ACM Transactions on Office Information Systems,* Vol 4:3, July 1986, pp. 225-250.

Hendler, J. *Integrating Marker-Passing and Problem Solving,* Lawrence Erlbaum Associates, Hillsdale, NJ. 1987.

Hendler, J., Tate, A., and Drummond, M. "AI Planning:Systems and Techniques" i, *AI Magazine* 11(2), May, 1990. ii. Reprinted (Japanese Translation) in *Nikkei Artificial Intelligence,* Nikkei Business Publications, Japan, Oct., 1990.

Hewitt, C. "Open Information Systems Semantics for Distributed Artificial Intelligence" *Journal of Artificial Intelligence ,* Vol. 14, No, 1, January 1991.

Hewitt, C. "Towards Open Information Systems Semantics," *Towards Open Information Systems Science,* Hewitt, C. E. , Manning, C. R. , Inman, 1. T. , and Agha, G. (eds.), MIT Press, Spring 1991.

Hewitt, C. "Scientific Communities and Open Information Systems," *Towards Open Information Systems Science,* Hewitt, C. E, , Manning, C. R. , Inman, I. T. , and Agha, G. (eds,), MIT Press, Spring 1991.

Hewitt, C. and Inman, J. "DAI Betwixt and Between: From Intelligent Agents to Open Systems Science'[1] *IEEE Transactions on Systems, Man, and Cybernetics* Special Issue on Distributed Artificial Intelligence, 1991.

Minksy, M. *The Society of Mind,* Simon and Schuster, New York, 1987.

Minsky, M. Preface, in Connectionist Models and Their Implications: Readings from Cognitive Science, eds. D. L. waltz and J, Feldman vii-xvi. Norwood, N.J.: Ablex, 1988

Minsky, M. "Logical vs. analogical or symbolic vs. connectionist or neat vs. scruffy." In Patrick H, Winston with Sara A. Shellard, editor, *AI al MIT: Expanding Frontiers,* MTT Press, 1990.

Sanborn, J. and Hendler, J. "Monitoring and Reacting: Planning in dynamic domains'[1] *International Journal of AI and Engineering,* 3(2), April, 1988.

Tatar, D., Foster, G,, and Bobrow, D. "Reflections on a face-to-face Multiuser System," IJMMS, Special Issue, 1991