Learning Coordination Plans in Distributed Problem-Solving Environments

Toshiharu Sugawara
NTT Basic Research Laboratories
3-1 Wakamiya, Morinosato
Atsugi, Kanagawa 243-01
Japan
sugawara@ntt-20.ntt.jp

Victor Lesser*
Department of Computer Science
University of Massachusetts
Amherst, MA 01003
USA
lesser@cs.umass.edu

To achieve globally coherent activity in a cooperative, distributed multi-agent system, a complete view of other agents' activities is often required. For specific problem-solving situations, however, due to the cost of meta-level processing, it may not be worthwhile to acquire the complete view, and thus some level of non-coherent activity may be the optimal coordination strategy (Lesser 1991).

For example, coordination to avoid redundant activities may be unnecessary if processing resources are not overloaded and if the communication channel is neither expensive nor overloaded. In this case, local problem solving is done more efficiently where there is no additional overhead for coordination. If a coordination strategy can be developed whose costs can be varied depending upon the amount and type of non-local information it uses to make coordination decisions, then it seems that only a selected, possibly situation-specific view is necessary. In this research, we propose integrating into each agent a distributed learning component that agents can use to acquire through experience which information is most effective for a specific situation and how to exploit this information.

Our approach to learning coordination rules can be described in the coordination model proposed by Decker and Lesser (Decker & Lesser 1993). In this model, each agent makes scheduling decisions based on a subjective view of its own and other agents' task structures and the relationships among these tasks. A subjective agent view does not include an important relationship or the quantitative character of the relationship is incorrectly inferred because of the use of a default assumption or out-of-date/incorrect data.

Knowledge used by our learning framework includes a heuristic collection of rules and procedures for 1) recognizing situations where there is costly incoherent behavior, 2) identifying control decisions that lead to this behavior, and 3) modifying these control decisions or replacing them with new decision processes that rectify the inappropriate control. It is also assumed that an agent locally records an abstracted trace of its recent problem-solving actions which can be reviewed on-line by the learning component. When an undesirable situation is observed, the learning component first identifies the tasks and messages which contributed to achieving the final goal. It then locates in the trace which tasks induced the observed undesirable situation. Next, agents rebuild the situation model they had locally in the diagnostic process. These models are exchanged and agents generate a more comprehensive model about the situation. Agents then reproduce the inference process based on this model. Two cases are possible. The first case can be solved in a very local way, based on the comprehensive model, by just choosing an alternative task to satisfy the given diagnostic goals. In this case, the learning component determines what aspects of the task and resource structure, which if taken into account in the scheduling of local activities would lead to more coordinated behavior; it then creates the rule to acquire or calculate this data (i.e., the agent's subjective view of the task relations). The second case is not solved solely by making existing rules more context sensitive with respect to the external agent environment. In this case, control rules for a new situation-specific scheduling strategy are added to increase the ratings of appropriate tasks or alter the coordination strategy in this specific case. The examples are described in (Sugawara & Lesser 1993)

*This research was supported in part by a grant from ARPA contract N00014-92-J-1698, an ONR contract N00014-92-J-1450, and a grant from Network General Corporation. The content of the information does not necessarily reflect the position or the policy of the organizations supporting this research, and no official endorsement should be inferred.

References

Decker K. and Lesser V. R. 1993. Quantitative Modeling of Complex Environments, *Int. J. of Intelligent Systems in Accounting, Finance and Management.*

Lesser V. R. 1991. A Retrospective View of FA/C Distributed Problem Solving, *IEEE Trans. on System., Man, and Cybernetics*, 21(6):1347-1362.

Sugawara T. and Lesser V. R. 1993. On-Line Learning of Coordination Plans, COINS Technical Report, 93-27, Univ. of Massachusetts.