

Coordination through Joint Intentions in Industrial Multiagent Systems

Nick R. Jennings

My Ph.D. dissertation (Jennings 1992a)¹ develops and implements a new model of multiagent coordination, called JOINT RESPONSIBILITY (Jennings 1992b), based on the notion of joint intentions. The responsibility framework was devised specifically for coordinating behavior in complex, unpredictable, and dynamic environments, such as industrial control. The need for such a principled model became apparent during the development and the application of a general-purpose cooperation framework (GRATE) to two real-world industrial applications. These experiments were successful in that it was possible to instantiate useful cooperation schemes; however, when anything unexpected happened (for example, new information invalidated existing goals, synchronization between actions was disrupted, or agents had misinterpreted the situation), the multiagent community acted incoherently. For example, agents would continue to work on a goal even though one community member knew their processing was obsolete; agents would stop processing requests if a more important task arose, without informing the originator; and agents would wait for the results of a task that had been abandoned. This incoherence occurred because the GRATE agents did not embody sufficient knowledge about the process of team problem solving. Therefore, it was decided to provide

agents with an explicit model of joint problem solving that they could reason about when deciding how to interact with others. JOINT RESPONSIBILITY was then implemented in an enhanced version of GRATE, and a series of comparative experiments were undertaken to assess the qualitative and quantitative benefits of the new approach.

Collaborative Problem Solving

In distributed AI (DAI) systems, problem-solving agents cooperate to achieve the goals of the individuals and of the system as a whole. Each individual is capable of a range of identifiable problem-solving activities, has its own aims and objectives, and can communicate with others. Typically, agents within a given system have problem-solving expertise that is related but distinct and that has to be coordinated when solving problems. Such interactions are needed because of the dependencies between agents' actions and the necessity to meet global constraints and because often, no one individual has sufficient competence to solve the entire problem.

Toward a Knowledge-Rich Cooperation Shell

Building multiagent systems is a complex and time-consuming task.

GRATE simplifies this process by providing a shell that contains built-in generic knowledge related to cooperation and control. The application designer can then build on this preexisting base of knowledge rather than construct the system completely from scratch (as is currently the case). To substantiate the claim that the knowledge is generic, GRATE was used to build two industrial applications: detection and location of faults in an electricity transportation network (Jennings et al. 1992) and cooperative diagnosis of a particle accelerator beam controller (Jennings et al. 1993). In both cases, the designer was able to construct a working multiagent system in a relatively short time span and did not need to augment the built-in knowledge.

Intentions and Joint Intentions

Intentions, such as I intend to enjoy this article, are one of the most popular means of describing the behavior of rational problem solvers (Bratman 1984). They provide objectives to which agents commit themselves, are used to coordinate future actions, and pose problems for means-end analysis. However, they are insufficient for describing collaboration; joint action is more than just the sum of individual actions, even if they are coordinated. Also, group commitment differs from individual commitment because a team can diverge in its beliefs (Cohen and Levesque 1991).

Existing models of joint intentions (Lochbaum, Grosz, and Sidner 1990; Searle 1990) provide only a partial description of the process of collaboration. Most importantly, from the perspective of industrial applications, these models do not describe how joint actions can falter and how individuals and the group should behave in such circumstances. Also, because the existing models were predominantly theoretical, little consideration was given to computational tractability. JOINT RESPONSIBILITY builds on and extends Cohen and Levesque's (1991) work on joint intentions: defining preconditions that must be satisfied before joint problem solving can

commence and extending the notion of joint commitment to plan states. Responsibility specifies that each individual within a team should remain committed to achieving the common objective by the agreed solution until one of the following statements becomes true: The objective has been met, the objective will never be met, the motivation for the action is no longer present, the desired outcome of a plan step is already available, following the agreed action sequence does not achieve the desired outcome, one of the specified actions cannot be carried out, or one of the agreed actions has not been carried out. While in this state, the agent will honor its commitments and carry out its agreed actions. However, if an agent is no longer committed to the joint action or the common solution, it cannot simply abandon its processing because its accomplices might not have been able to detect the problem. For this reason, the responsibility model stipulates that when a team member is no longer jointly committed to the joint action, it must ensure that all its acquaintances are informed of this change of state. Thus, the whole team can reassess the viability of the joint action and, in particular, the actions involving the agent that is no longer committed.

A rule-based interpretation of JOINT RESPONSIBILITY was then used to build agents that had an explicit and principled model of collaboration to guide their individual actions and their social interactions.

Experimental Evaluation

A series of comparative experiments were undertaken to assess the performance characteristics of the responsibility model (Jennings and Mamdani 1992). Three types of problem-solving organization were compared: (1) a responsible community; (2) an implicit group model in which agents had individual intentions but did not form explicit collaborating groups; and (3) groups of problem solvers who set up joint intentions but, when the joint action became unsustainable, behaved selfishly and simply abandoned their local processing without informing

their fellow team members.

These experiments showed that the responsible community performed significantly more coherently than the other two, this difference being especially noticeable as the domain became more dynamic and unpredictable (that is, the chance of joint action unsustainability increased). This gain in performance was achieved with negligible extra processing requirements for the coordination mechanisms.

Conclusions

This work shows, through empirical evaluation on a real-world problem, that a suitably formulated model of joint intentions is a powerful mechanism for coordinating the behavior of collaborating agents. The model is especially useful when agents have to make decisions using partial and imprecise information and when the environment itself is evolving and unpredictable. It also indicates how theoretical models of coordination can be used as a basis for implementation-level systems. Two new domains in which DAI techniques can be exploited profitably were also highlighted. Finally, as a consequence of the insights gained in this work, a proposal for the next generation of multiagent systems was made. In such *cooperation knowledge-level systems* (Jennings 1992b), individuals maintain and reason about explicit and deep representations of social interactions rather than have an implicit and shallow understanding of these processes.

Acknowledgments

This work benefited greatly from the guidance received from Abe Mamdani, Erick Gaussens, and Thies Wittig. Jose Corera, Inaki Laresgoiti, Juan Perez, Rob Aarnts, Paul Skarek, Laszlo Varga, and Joachim Fuchs provided valuable help by applying GRATE to the real-world applications. This work has been supported in part by the ESPRIT II Project ARCHON (P-2256).

Note

1. Copies of this dissertation (KEAG-TR-92-18) can be obtained by writing to the author at the Department of Electronic

Engineering, Queen Mary and Westfield College, University of London, Mile End Road, London E1 4NS, UK. E-mail requests should be addressed to N.R.Jennings@qmw.ac.uk. An updated version of the thesis, entitled "Cooperation in Industrial Multi-Agent Systems," will be published by World Scientific Press in 1994.

References

- Bratman, M. E. 1984. Two Faces of Intention. *Philosophical Review* 93:375-405.
- Cohen, P. R., and Levesque, H. J. 1991. Teamwork. *Nous* 25(4): 487-512.
- Jennings, N. R. 1992a. Joint Intentions as a Model of Multi-Agent Cooperation. Ph.D. diss., Dept. of Electronic Engineering, Queen Mary and Westfield College, University of London.
- Jennings, N. R. 1992b. Toward a Cooperation Knowledge Level for Collaborative Problem Solving. In Proceedings of the Tenth European Conference on Artificial Intelligence, 224-228. New York: John Wiley and Sons.
- Jennings, N. R., and Mamdani, E. H. 1992. Using Joint Responsibility to Coordinate Collaborative Problem Solving in Dynamic Environments. In Proceedings of the Tenth National Conference on Artificial Intelligence, 269-275. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Jennings, N. R.; Mamdani, E. H.; Laresgoiti, I.; Perez, J.; and Corera, J. 1992. GRATE: A General Framework for Cooperative Problem Solving. *Journal of Intelligent Systems Engineering* 1(2): 102-114.
- Jennings, N. R.; Varga, L. Z.; Aarnts, R. P.; Fuchs, J.; and Skarek, P. 1993. Transforming Standalone Expert Systems into a Community of Cooperating Agents. *International Journal of Engineering Applications of Artificial Intelligence* 6(4): 317-331.
- Lochbaum, K. E.; Grosz, B. J.; and Sidner, C. L. 1990. Models of Plans to Support Communication. In Proceedings of the Eighth National Conference on Artificial Intelligence, 485-490. Menlo Park, Calif.: AAAI.
- Searle, J. 1990. Collective Intentions and Actions. In *Intentions in Communication*, eds. P. R. Cohen, J. Morgan, and M. E. Pollack, 401-416. Cambridge, Mass.: MIT Press.

Nick R. Jennings is a lecturer in the Department of Electronic Engineering at Queen Mary and Westfield College (University of London). He obtained his Ph.D. in AI from Queen Mary and Westfield College in 1992. His current research interests include coordination in multiagent systems, conflict resolution, joint intentions, and metalevel control architectures.