

## **Coordination Through Joint Intentions in Industrial Multi-Agent Systems**

Nick R. Jennings

My Ph.D. dissertation (Jennings 1992a)<sup>1</sup> develops and implements a new model of multi-agent coordination, called *Joint Responsibility* (Jennings 1992b), based upon the notion of joint intentions. The Responsibility framework was devised specifically for coordinating behaviour in complex, unpredictable and dynamic environments such as industrial control. The need for such a principled model became apparent through developing and applying 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 (eg new information invalidated existing goals, synchronisation between actions was disrupted or agents had misinterpreted the situation) the multi-agent community acted incoherently. For instance, 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; agents would wait for the results of a task which had been abandoned and so on. 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 about which they could reason 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 which is related, but distinct, and which has to be coordinated when solving problems. Such interactions are needed because of the dependencies between agents' actions, the necessity to meet global constraints and because often no one individual has sufficient competence to solve the entire problem.

## **GRATE: Towards a Knowledge-Rich Cooperation Shell**

Building multi-agent systems is a complex and time consuming task. GRATE simplifies this process by providing a shell which contains inbuilt generic knowledge related to cooperation and control. The application designer can then build upon this preexisting base of knowledge, rather than constructing the system completely from scratch (as is the case at present). 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 multi-agent system in a relatively short space of time and did not need to augment the inbuilt knowledge.

### **Intentions and Joint Intentions**

Intentions, such as “I intend to enjoy this article”, are one of the most popular means of describing the behaviour 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 *et al.* 1990; Searle 1990) provide only a partial description of the process of collaboration. Most importantly, from the perspective of industrial applications, they do not describe how joint actions may falter and how individuals and the group should behave in such circumstances. Also as the existing models were predominantly theoretical, little consideration had been given to computational tractability. Joint responsibility builds upon and extends Cohen and Levesque’s (1991) work on joint intentions - defining preconditions which 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 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. Whilst in this state, the agent will honour 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 may 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. This enables the whole team to reassess the viability of the joint action and in particular the actions involving the agent which is no longer committed.

A rule-based interpretation of Joint Responsibility was then used to build agents which 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 organisation were compared: (i) a responsible community; (ii) an implicit group model in which agents had individual intentions, but did not form explicit collaborating groups; (iii) 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 responsible communities performed significantly more coherently than the other two; this difference being especially noticeable as the domain became more dynamic and unpredictable (i.e. 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 behaviour of collaborating agents. This is especially true in situations where 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 profitably exploited were also highlighted. Finally as a consequence of the insights gained in this work, a proposal for the next generation of multi-agent systems is made. In such “cooperation knowledge level systems” (Jennings 1992b) individuals maintain and reason about explicit and deep representations of social interactions, rather than having an implicit and shallow understanding of these processes.

### **Acknowledgments**

This work benefitted 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 helping to apply GRATE to the real-world applications. This work has been partially supported by the ESPRIT II project ARCHON (P-2256).

### **Notes**

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. Email requests should be addressed to: N.R.Jennings@qmw.ac.uk

### **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).

Jennings, N. R. 1992a. Joint Intentions as a Model of Multi-Agent Cooperation. Ph.D. diss., Dept. of Electronic Engineering, Queen Mary & Westfield College, University of London.

Jennings, N. R. 1992b. Towards a Cooperation Knowledge Level for Collaborative Problem Solving. In Proceedings of the Tenth European Conference on Artificial Intelligence, 224-228. Vienna, Austria.

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. San Jose, Calif.

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. *Int Journal of Engineering Applications of Artificial Intelligence* 6(4).

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. Boston, Mass.

Searle, J. 1990. Collective Intentions and Actions. In *Intentions in Communication*, eds. P. R. Cohen, J. Morgan and M. E. Pollack, 401-416. MIT Press.