Negotiation in Multi-Agent Systems *

Martin Beer[1], Mark d'Inverno[2], Michael Luck[3],
Nick Jennings[4], Chris Preist[5], and Michael Schroeder[6]

1. Department of Computer Science, Liverpool University, UK
2. Cavendish School of Computer Science, University of Westminster, UK
3. Department of Computer Science, University of Warwick, UK
4. Queen Mary and Westfield College, University of London, UK
5. Agent Technology Group, Hewlett Packard Labs, Bristol, UK
6. Department of Computer Science, City University, UK

In systems composed of multiple autonomous agents, negotiation is a key form of interaction that enables groups of agents to arrive at a mutual agreement regarding some belief, goal or plan, for example. Particularly because the agents are autonomous and cannot be assumed to be benevolent, agents must *influence* others to convince them to act in certain ways, and negotiation is thus critical for managing such inter-agent dependencies. The process of negotiation may be of many different forms, such as auctions, protocols in the style of the contract net, and argumentation, but it is unclear just how sophisticated the agents or the protocols for interaction must be for successful negotiation in different contexts. All these issues were raised in the panel session on negotiation.

As a prelude to the discussion, Jennings identified three broad topics for research on negotiation, that serve to organise the issues under consideration. First, *negotiation protocols* are the set of rules that govern the interaction. This covers, the permissible types of participants (e.g., the negotiators and relevant third parties), the negotiation states (e.g., accepting bids, negotiation closed), the events that cause state transitions (e.g., no more bidders, bid accepted), and the valid actions of the participants in particular states (e.g., which can be sent by whom, to whom and at when). Second, *negotiation objects* are the range of issues over which agreement must be reached. These may single issues, such as price, or multiple issues relating to price, quality, timing, etc. Also relevant here are the allowable operations on these objects. In the simplest case, the structure and contents of the agreement are fixed, and negotiation amounts to accepting or rejecting the offer. The next level, however, offers flexibility to change the values of the issues in the negotiation object, through counter-proposals, changing the structure of the negotiation object (by adding guarantees, for example), and so on. Finally, the *agents' reasoning models* provide the decision making apparatus by which participants attempt to achieve their objectives. The sophistication of the model is determined by the protocol used, the nature of the negotiation object, and the range of operations that can be performed on it.

The range of problems and issues in this area is clearly wide, yet they can be considered as fundamental to many forms of interaction. To focus discussion, the panelists were asked seven questions that were felt to be key to negotiation. Each of the questions are presented below along with a summary of the responses of the panelists.

> *Argumentation and negotiation in multi-agent systems can involve sophisticated, high-level reasoning. How relevant to such processes are the lower-level communication lan-*

*This report is the result of a panel discussion at the Workshop of the UK Special Interest Group on Multi-Agent Systems (UKMAS'98). All members of the panel are authors, listed alphabetically.

guages that underpin them? In what way do they constrain the kinds of processes possible to implement with them?

In order to build practical applications, the low-level underpinning infrastructure is clearly very important. For example, any formally-based argumentation/negotiation protocol must deal with the loss of messages or one of the partners of the interaction being down. While the former can be assumed to be handled by a reliable communication protocol (such as TCP/IP), or a standard distributed computing platform (such as CORBA), the latter must be part of the argumentation/negotiation protocol.

Indeed, lower-level communication languages control the conversation classes in which pairs of agents can participate. The conversational components are defined not only by the action specified, but also by the pre-conditions, post-conditions and completion-conditions of that performative. These define the circumstances under which the communication can take place, and include the belief state of the respective agents. To be able to negotiate effectively, the agents must have some beliefs about what the world will be like in the future, rather than an implicit belief in the 'truth' (at least as far as the responding agent is concerned) of a conventional ask-tell conversational class.

For a negotiation to be completed successfully all parties (i.e. agents) must clearly understand the *rules of engagement* or negotiation protocol. For example, in a simple contract-net protocol, in which a manager issues a *call for proposals* and waits for a full set of replies (or timeouts), each bidder must be prepared to honour its bid for the duration of the bid's validity. Otherwise, acceptance of the bid will require a second negotiation, which may itself succeed or fail. If an agent is no longer able to honour its bid, the other agents that originally bid may well have already have accepted other commitments and no longer be in a position to accept that commitment. The whole bidding process might then have to be restarted.

It is clear that the design of a communications language can restrict or enable different forms of high level reasoning among the agents involved in negotiation. For example, KQML (Labrou and Finin, 1997) places unnecessary restrictions on an agent community - an agent sending a query is expected to decide how the recipient will handle it, rather than the recipient (and possibly other members of the community) making that decision. The sender must decide whether the recipient will respond directly, broker the query, recommend an agent to send the query to, or recruit an agent who will send the response directly to the original agent. Based on this decision, the agent must choose an appropriate KQML performative to send. However, in everyday negotiation, if one is asked "Can you fix this bug?", it is more likely to be you who decides how to involve others. You might reply, "No, but Julie knows about that sort of thing. Ask her.", or you might say yes, but subcontract it to Julie without involving me. Hence the decision of whether to act as a broker, a recommender, or to respond directly is left to the recipient of the request, not the requester. In KQML, this is not possible.

*The contract net is described as negotiation though it is simply a very straightforward protocol: task announcement, make bid, award contract etc. Does this count as negotiation? If not, what *does* constitute negotiation in terms both of the participating autonomous agents and the nature of the communication itself?*

If negotiation is defined as any communicative process that results in mutually acceptable agreement then clearly the contract net protocol does constitute negotiation albeit in a rather simple form. Similarly, the *sealed-price second bid auction* is a means of negotiation, which involves even less communication than the contract net. One of the tasks of a protocol designer is to create a protocol that accomplishes the desired result with the minimum amount of communication. Simplicity is best.

Negotiation in a multi-agent systems context will always be specified by a protocol — be it simple or sophisticated, deterministic or non-deterministic — and thus will not be different in its nature from the contract net protocol. In fact, even the Encyclopedia Britannica defines *negotiation* quite generally as, "the action or process of negotiating or being negotiated." This definition looks a little like a circular argument, but there is more in the definition of *negotiate* as, "to carry on business, ..., to confer with another so as to arrive at the settlement of some matter ...."

These definitions easily cover such protocols as the contract net and other "simple" protocols.

Can all aspects of negotiation be incorporated into predefined communication protocols?

Most members of the panel suggested that this is not possible and that it is the agent's reasoning model that is the key to successful negotiation. Preist offers as an example a second hand electronic components market, in which an important consideration is whether one agent trusts the other to deliver when he says he will. Trust is established partly as a result of the way in which the negotiation is carried out. It is unlikely that this could be represented in protocols.

Beer, however, takes a slightly different position in arguing that the use of special purpose negotiation performatives allow different belief states in the future to be clearly identified, thus avoiding the possibility of confusion that might arise with simple ask-tell performatives. He argues that it is the *conversation class* that defines the reason for the communication, not the performatives used.

Are agent communication languages merely stock workhorses with the majority of current examples (ACL, KQML, etc) roughly equivalent? Does it matter which is chosen for a particular application or problem area?

Since communication languages can restrict what is possible in different ways, clearly the choice of language may make a difference in some cases. However, ACL and KQML are roughly the same and, for many applications, they are fine, as KQML already proved. Interestingly, the KQML specification (Finin *et al.*, 1993) uses implicit and explicit negation. For example, KQML defines the speech act $untell(a)$ to mean that a cannot be proved, i.e. $not\ a$ holds, whereas $tell(\neg a)$ means that there is an explicit proof for $\neg a$. Unfortunately, KQML uses this distinction only implicitly and, consequently, many speech acts may be mis-used by practitioners leading to problems when agents inter-operate.

According to Schroeder, while KQML has already been around for a while and has been tested in many applications, FIPA ACL [1] is relatively new and error-prone. Wagner points out several shortcomings of FIPA ACL (Wagner, 1997). For example, is the distinction between $inform$, $tell$, $assert$ only relevant for an external observer; for an agent's subjective view they are equivalent. Also, an $inform$ may be used to answer queries which should, however, be done only by a $reply$.

Although these points may seem minor, Schroeder argues that they will prevent the use of ACL for large-scale, open multi-agent systems. At the moment, ACL is used only in small or medium-sized applications so that problems are not apparent. Once different ACL implementations become available, ACL may cause inter-operability problems and will need to be revised. Another far more important point is that KQML and ACL are too *academic* for many applications. When TMN agents communicate, they use their simple vocabulary to set and get variable values. When human or artificial agents negotiate in auctions they do not necessarily require speech acts and logic to carry out sophisticated transactions. In fact, most of them don't use them at all.

[1] `http://www.fipa.org/`

Can we ever envisage a standard set of negotiation protocols which all autonomous agents can adhere to, or will negotiation protocols have to be designed afresh for different applications?

What role do standards have here? Similarly, is it likely that there will be a single universal standard for inter-agent communication on which to base negotiation, or is that an unattainable ideal?

Standard negotiation protocols such as on-line auctions already exist, and more are being developed all the time. In principle, it should be possible to develop some kind of *uniform protocol theory*, which embodies all kinds of protocols — ranging from auctions to protocols for co-operation to coalition formation. However, to expect every agent (irrespective of domain) to be able to conform to any of these protocols would result in seriously heavyweight agents. Fortunately, though, it does not necessarily follow that every application will need a different protocol. There are many different tasks requiring coalition formation, and Preist expects that a (possibly small) set of such protocols would apply to most of them, just as a small set of auction protocols satisfies most applications involving one-to-many negotiation over a commodity.

Beer argues that the mechanisms by which negotiation can be undertaken are fairly clearly defined (for example by FIPA). It is the negotiation *currency* that is very much application dependent. Most work reported is based on monetary values of some form, but we can easily conceive of situations where speed of response or some measure of 'closeness' would be more appropriate.

Schroeder puts forward the alternative view that it is unlikely that any such standard will emerge for two reasons. First, the applications involving argumentation/negotiation are very different in their nature ranging from transactions in auctions to argumentation in the philosophical sense. Second, if argumentation/negotiation is of importance to the participants, they will want to use the best protocol available making a standard solution unlikely.

Clearly, it is still a matter of much debate as to whether there will ever be a set of inter-agent negotiation protocols that all agents will adhere to.

What sorts of capabilities are required of agents in order to engage in effective argumentation and negotiation? With a view to implementing argumentation and negotiation, what particular characteristics might be desired of such lower-level infrastructure?

In responding to this question, Schroeder illuminated the discussion by referring to an historical example. The first scientist who linked argumentation as a discipline in philosophy to mathematics was Leibniz. His aim was the definition of an automatic calculus for discourse (Herring, 1992). His ideas and solutions can teach us a lot about the pitfalls of arguing agents. Leibniz's motivation for a calculus of reasoning was nothing less than world peace and the ability to settle questions such as the existence of god. Leibniz believed that arguments derived from his "*calculus raciocinator*" had to be accepted by everybody due to their formal nature. His idea for the calculus was simple and brilliant. He mapped basic concepts to prime numbers and complex numbers to non-primes such that the basic concepts defining the complex concepts are the prime factors of the non-prime. As an example he assigns the primes 2 and 3 to the basic concepts "animal" and "rational", respectively and derives that the complex concept "human" is therefore 2×3, i.e. 6. Then he concludes that a "monkey" with number 10 is not a "human" because neither can we divide 10 by 6 nor vice versa. We can deduce, however, that both have got the concept "animal", i.e. the factor 2, in common.

This idea is quite nice and was revolutionary at its time, because it treated human reasoning mathematically for the first time. But was it useful? Leibniz wants to show this by "assuming all these wonderful numbers as given". And this is the problem of Leibniz's calculus and the arguing/negotiating

agents we are currently implementing: the developed algorithms and protocols are only useful if they are filled with "life". Leibniz failed to do so, and so may we. The most crucial part in any practical argumentation/negotiation is not the formal definition of algorithms and protocols but the ability to understand an argument. Even with projects like CYC and a lot of work on ontologies, we are nowhere near this. And, as long as computers cannot understand, simple argumentation/negotiating protocols will be the most effective tools.

This should not put us off researching more complex issues, but we should be aware of the limits. When implementing such systems, one soon realises that quite a sophisticated infrastructure is needed. Many applications will require non-monotonic reasoning and two kinds of negation — explicit negation by which there is an explicit proof in support of the negation of a, and implicit negation which expresses the absence of evidence for a. This latter form can be used to express things like the reasonable doubt standard where the absence of evidence for someone's guilt leads to the conclusion that they are not guilty.

Given such an expressive knowledge representation, we may need agents with different credibility. When arguing about the credit-worthiness of a client, the agent taking the final decision should be sceptical; when distributing work to other agents, the distributor should be credulous rather than arguing with helpers who stepped forward to do the job. Technically, the definition of such agents requires a flexible definition of an acceptable argument. When dealing with undercuts and rebuts, a sceptical agent allows counter-arguments of both types whereas a credulous agent admits only undercuts (Schroeder, 1999).

Jennings addressed the question of which capabilities are required for effective negotiation and argumentation with a firm and brief enumeration of the minimal requirements for negotiation and argumentation. Agents must be able to make proposals, counter-proposals, accept and reject proposals, and generate arguments in support of their adopted stance on some subject.

Preist extended this with some points relating specifically to competitive negotiation. First, all parties should be clear about the belief sets of each of the agents involved in the negotiation process. At the very least a shared ontology of the domain being negotiated over, or a way of establishing one, needs to exist. Then, they must be able to make contact with potential negotiators and recognise the type of 'game' in progress. Questions such as the number of negotiators, the number of potential trading partners, and the market power of individuals, for example, need to be considered and an appropriate strategy based on this information selected. During the negotiation, agents require both a means to evaluate the relative value to the agent of different offers made during negotiation and a means to evaluate the relative value to other agents of offers the agent may potentially make. Finally, in some situations, the ability to exchange constraints on acceptable offers and to reason with these constraints may make negotiation more efficient.

How is argumentation related to negotiation?

Jennings suggested that while proposals, counter-proposals and the ability to accept and reject proposals are minimal requirements of negotiation, they are inadequate if agents seek to *justify* their stance or to *persuade* others to change their negotiation stance. If there are compelling reasons for adopting a particular position, agents may use *arguments* to try to change their opponents view, and to make their own proposal more attractive by providing additional meta-level information in its support. The nature and types of arguments are varied, but include threats, rewards, appeals, and so on. Thus, argumentation is a *type* of negotiation that includes meta-level exchanges to support a negotiator's position.

Preist countered that argumentation is an approach to defining the semantics of negotiation, but that there are other ways of doing so, and that argumentation can also be used to define semantics of

other systems. Similarly, Schroeder believes that while argumentation may be useful for negotiation it is not clear that argumentation is necessary for negotiation at all. Argumentation and negotiation are disciplines of their own, and have been developed independent of any restrictions imposed by a computer science concept. Ultimately, however, it can be argued that value arises from application and, while core aspects of argumentation and negotiation may be distinct, using argumentation as a means of attaining success through effective negotiation offers a particular kind of vision for this particular line of research.

What is clear from the panel discussion is that the issues surrounding negotiation offer many avenues to explore in agent research, covering a range of varied aspects including some of those touched on above, such as communication languages, protocols, standards, strategies, and models of reasoning. Despite the centrality (or perhaps even because of the centrality) of negotiation to multi-agent interaction, much work remains to be done here. Indeed, as the applications arena continues to broaden, with the fields of web-based systems and electronic commerce, for example, surging forward, the need for effective mechanisms for agent interaction in these areas suggests that the importance of negotiation will remain high on the agenda of researchers both in this field and others.

References

Labrou, Y and Finin, T, 1997. Proposal for a new KQML Specification TR CS-97-03, February 1997, Computer Science and Electrical Engineering Department, University of Maryland Baltimore County, Baltimore, MD.

Finin, T, Weber, J, Wiederhold, G, Genesereth, M, Fritzson, R, McKay, D, McGuire, J, Pelavin, R, Shapiro, S and Beck, C, 1993. Specification of the KQML agent communication language, Technical report, The DARPA Knowledge Sharing Initiative, External Interfaces Working Group, Baltimore, MD.

Herring, H (ed.), 1992. *Gottfried Wilhelm Leibniz. Schriften zur Logik und zur Philosophischen Grundlegung von Mathematik und Naturwissenschaften*, volume VII. Wissenschaftliche Buchgesellschaft, Darmstadt.

Schroeder, M, 1999. "An efficient argumentation framework for negotiating autonomous agents" in *Proceedings of Modelling Autonomous Agents in a Multi-Agent World MAAMAW99*. Springer-Verlag.

Wagner, G, 1997 "Artificial agents and logic programming" in *Proceedings of the Workshop on Logic Programming and Multiagent Systems*, `www.inf.fu-berlin.de/ wagnerg/fipa2.html`