On Argumentation-Based Negotiation

N. R. Jennings[1], S. Parsons[1], P. Noriega[2], and C. Sierra[2]

[1]Dept. Electronic Engineering, Queen Mary & Westfield College,
University of London, London E1 4NS, UK.
{N.R.Jennings, S.D.Parsons}@elec.qmw.ac.uk

[2]Artificial Intelligence Research Institute, Spanish Scientific Research Council,
Campus UAB, 08193 Bellaterra, Barcelona, Spain.
{sierra, pablo}@iiia.csic.es

1. Introduction

Interactions are a core part of all multi-agent systems. They occur because of the inter-dependencies which inevitably exist between the agents and they manifest themselves in many different forms—including cooperation, coordination, collaboration. However perhaps the most fundamental and powerful mechanism for managing these inter-agent dependencies at run-time is *negotiation*—the process by which a group of agents communicate with one to try and come to a mutually acceptable agreement on some matter. Negotiation underpins attempts to cooperate and coordinate (both between artificial and human agents) and is required both when the agents are self interested and when they are cooperative. It is so central precisely because the agents are autonomous. For an agent to influence an acquaintance, the acquaintance needs to be convinced that it should act in a particular way. The means of achieving this state are to make proposals, trade options, offer concessions, and (hopefully) come to a mutually acceptable agreement. In short, to negotiate.

Given its ubiquity and importance in many different contexts, negotiation theory covers a broad range of phenomena and encompasses multifarious approaches (e.g. from Artificial Intelligence, Social Psychology, and Game Theory). Despite this variety, however, negotiation research can be considered to deal with three broad topics:

- *Negotiation Protocols*: the set of rules which govern the interaction. This covers the permissible types of participants (e.g. the negotiators and any relevant third parties), the negotiation states (e.g. accepting bids, negotiation closed), the events which cause negotiation states to change (e.g. no more bidders, bid accepted), and the valid actions of the participants in particular states (e.g. which messages can be sent by whom, to whom, at what stage).

- *Negotiation Objects*: the range of issues over which agreement must be reached. At one extreme, the object may contain a single issue (such as price), while on the other hand it may cover hundreds of issues (related to price, quality, timings, penalties, terms and conditions, etc.). Orthogonal to the agreement structure is the issue of the types of operation which can be performed on it as dictated by the negotiation protocol. In the simplest case, the structure and the contents of the agreement are fixed and participants can either accept or reject it (i.e. a take it or leave it offer). At the next level, participants have the flexibility to change the values of the issues in the negotiation object (i.e. they can make counter-proposals to ensure the agreement better fits their negotiation objectives). Finally, participants might be allowed to dynamically extend the structure of the negotiation object (e.g. a car salesman may add one year's free insurance into a negotiation in order to clinch the deal).

- *Agents' Decision Making Models*: the decision making apparatus the participants employ to act in line with the negotiation protocol in order to achieve their negotiation objectives. The sophistication of the model, as well as the range of decisions which have to be made, are influenced by the protocol in place, by the nature of the negotiation object, and by the range of operations which can be performed on it.

The relative importance of these three topics varies according to the negotiation and environmental context. Thus, in some circumstances the negotiation protocol is the dominant concern (e.g. [8] [13]). For example, the system designer may determine that the negotiation is best organised using a particular form of auction (e.g. English, Dutch, Vickrey, First-Price Sealed Bid). This mechanism design choice constrains the types of operations which can be performed on the negotiation object (no counter-proposals or issue extensions) and prescribes the behaviour of the agents' decision making models (e.g. strategic behaviour is pointless and agents should simply bid their true reservation value). In other cases, however, the agent's decision making model is the dominant concern (e.g. [1] [9] [11]). Here, the protocol does not prescribe an agent's behaviour and there is scope for strategic reasoning to determine the best course of action. In such cases, the relative success of two agents is determined by the effectiveness of their reasoning model—the better the model, the greater the agent's reward.

Given the wide variety of possibilities, it is clear that there is no universally best approach or technique for inter-agent negotiation. Rather, there is an eclectic bag of methods with properties and performance characteristics which vary widely depending on the negotiation context. Against this background, our work concentrates on a particular class of negotiation which we term *argumentation-based*. In this form of negotiation, agents generate and exchange arguments to back up or justify their negotiation stance. In more detail, this paper seeks to: (i) motivate the importance of an argumentation-based approach (section 2); (ii) describe the architecture of an agent's decision making model as it pertains to supporting argumentation-based negotiation (section 3); and (iii) indicate the techniques we have developed to date for formally specifying the mechanisms of one such model (section 4).

2. A Generic Framework for Argumentation-Based Negotiation

Negotiation can be viewed as a distributed search through a space of potential agreements (figure 1). The dimensionality and topology of this space is determined by the structure of the negotiation object. Indeed, one could consider each attribute of the negotiation object to have a separate dimension associated with it; clearly, in this view, the space of figure 1 concerns two attributes. Thus, when new issues are added (or old ones removed) during the course of a negotiation, then extra dimensions are added (or removed) and the number of points of agreement may increase (or decrease). Similarly, if an agent changes one of the values of one of the attributes within an offer, it is moving from one point in the agreement space to another. For more on this metaphor for viewing the agreement space see [4] [5].

For a given negotiation, the participants are the active components who determine the direction of the search. At the start of this process, each agent has a portion of the space in which it is willing to make agreements. Typically, it also has some means of rating the points in its space and some means of using this rating to determine the actual agreements it makes. Negotiation proceeds by the participants suggesting specific points (or regions) in the agreement space as potentially acceptable. During the negotiation process, the participants' agreement spaces (as well as their rating functions) may change: they may expand, contract, or shift, for instance because their environment changes, or because they are persuaded to change their views. The search terminates when the required number of participants find a mutually acceptable point in the agreement space or when there are insufficient negotiators left to reach an agreement.

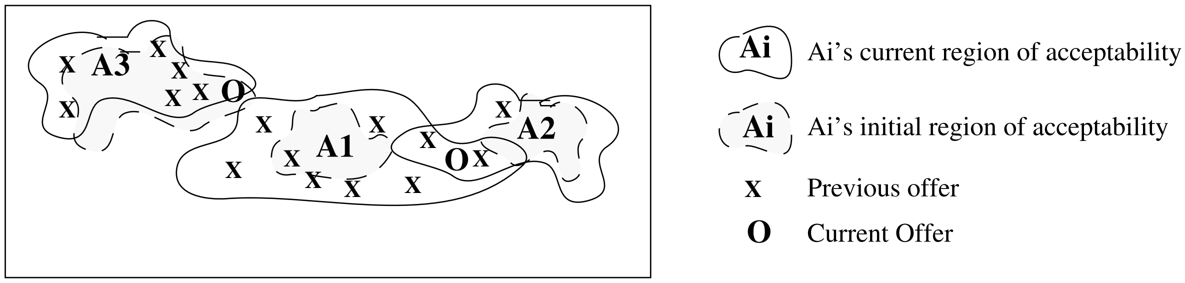

Figure 1: The Space of Negotiation Agreements

From this representation, it can be seen that the minimal negotiation capabilities are: (i) to propose some part of the agreement space as being acceptable; and (ii) to respond to such a proposal indicating whether it is acceptable. In other words, the minimum capability required of a negotiating agent is the ability to make and respond to proposals. Now since our work is set with in the context of agents reaching agreements about some joint problem, we consider a proposal to be a solution to that joint problem; either a single complete proposed solution, a single partial solution, or a group of complete or partial solutions. In terms of the agreement space, these different kinds of proposals become a single point, a region of the space, a set of points, or a set of regions of the space (for example a partial solution would be any region of the space in which the quality was above some level and the price below a certain threshold). We allow a proposal to be made either independently of other agents' proposals, or based on previous comments made by other agents.

The most minimal kind of negotiation we can imagine is that which takes place in a Dutch auction. The auctioneer (one agent in the negotiation) calls out prices (negotiation objects with a single attribute). When there is no signal of acceptance from the other parties in the auction (other agents in the negotiation) the auctioneer makes a new offer which it believes will be more acceptable (by reducing the price). Here, because of the convention (protocol) under which the auction operates, a lack of response is sufficient feedback for the auctioneer to infer a lack of acceptance. However in anything more complex than this rather special case, the minimal requirement for the "other agents" is that they are able to indicate dissatisfaction with proposals that they find unacceptable.

If agents can only accept or reject others' proposals, then negotiation can be very time consuming and inefficient since the proposer has no means of ascertaining why the proposal is unacceptable, nor whether the agents are close to an agreement, nor in which direction of the agreement space it should move next. Hence the proposer is essentially picking points in the agreement space based only on its own imperatives and hoping that it will eventually stumble upon something acceptable. To improve the efficiency of the negotiation process, the recipient needs to be able to provide more useful feedback on the proposals it receives than just whether or not it agrees to them. This feedback can take the form of a *critique* (comments on which parts of the proposal the agent likes or dislikes[1]) or a *counter-proposal* (an alternative proposal gen-

1. To avoid introducing an unnecessarily large number of different types of statement, we consider simple accept/reject statements to be special cases of critiques.

erated in response to a proposal). From such feedback, the proposer should be able to generate a proposal which is more likely to lead to an agreement (if it chooses to do so).

Consider the concept of a critique first. A critique provides two forms of feedback: (i) it suggests constraints on particular negotiation issues and (ii) it indicates acceptance/rejection of particular parts of the proposal (or indeed of the whole proposal). To illustrate these points, consider the following short dialogues which are examples of proposals followed by critiques:

```
A: I propose that you provide me with service X under the following condi-
     tions.
B: I am happy with the price of X, but the delivery date is too late.

A: I propose that I will provide you with service Y if you provide me with
     service X.
B: I don't want service Y.
```

In the first case, the critique indicates those aspects of the proposal which are acceptable and those which need to be modified and it also suggests a constraint on one of the issues (delivery date earlier than the current suggestion). In the second case, the critique indicates outright rejection of part of the proposal. Generally speaking, the more information placed in the critique, the easier it is for the original agent to determine the boundaries of its opponent's agreement space.

Counter proposals are the second feedback mechanism. A counter-proposal is simply a proposal, which is more favourable to the sender, made in response to a previous proposal. The following are examples of proposals followed by counter-proposals:

```
A: I propose that you provide me with service X.
B: I propose that I provide you with service X if you provide me with serv-
     ice Z.

A: I propose that I provide you with service Y if you provide me with serv-
     ice X.
B: I propose that I provide you with service X if you provide me with serv-
     ice Z.
```

In the first case, the counter-proposal extends the initial proposal, and in the second case it amends part of the initial proposal. Counter-proposals differ from critiques in that the feedback is less explicit (the recipient of a counter-proposal has to infer the constraints and preferences from the way the proposal is re-constituted), but generally more detailed (since specific regions of the opponent's agreement space are identified).

On their own, proposals, critiques and counter-proposals are bald statements of what agents want. Thus, their scope is confined solely to the structure of the negotiation object. While it is perfectly possible to base negotiations on just these object-level constructs (indeed this is precisely what most extant models do), doing so diminishes some of the potential of negotiation technology. For example, it means that agents cannot:

- *justify* their negotiation stance;

  An agent might have a compelling reason for adopting a particular negotiation stance. For example, a company may not be legally entitled to sell a particular type of product to a particular type of consumer or a particular item may be out of stock and the next delivery might not be until the following month. In such cases, the ability to provide the justification for its attitude towards a particular issue can allow the opponent to more fully appreciate an agent's constraints and behaviour.

- *persuade* one another to change their negotiation stance;

  Agents sometimes need to actively change their opponents agreement space, or its rating over that space, in order for a deal to be possible. In such cases, agents seek to construct arguments which they believe will make their opponent look more favourably upon their proposal. Thus, arguments seek to identify opportunities for such change (e.g. a car salesman throws in a stereo with a car to increase the value of the good), create new opportunities for change (e.g. a car salesman adds a new dimension to the rating function by highlighting the cars novel security features) or modify existing assessment criteria (e.g. car salesman gets buyer to change evaluation function by convincing him that security is more important than mileage).

In both cases, negotiators are providing *arguments* to support their stance (hence *argumentation-based negotiation*). Thus, in addition to generating proposals, counter-proposals and critiques, the negotiator is seeking to make the proposal more attractive (acceptable) by providing additional meta-level information in the form of arguments for its position. The nature and types of the arguments can vary enormously (see [2] [3] [11] for more details), however common categories include: threats (failure to accept this proposal means something negative will happen to you), rewards (acceptance of this proposal means something positive will happen to you), and appeals (you should prefer this option over that alternative for some reason). Whatever its precise form, the role of the supporting argument is either to modify the recipient's region of acceptability or its rating function over this region. In so doing, arguments have the potential[2] to increase the likelihood and/or the speed of agreements being reached[3]. In the former case, by persuading agents to accept deals that they may not originally have countenanced. In the latter case, by convincing agents to accept their position on a given issue and to cease negotiating over it.

3. Building Argumentation-Based Negotiators

To design and build an agent capable of effective argumentation-based negotiation, it requires the following:

1. Mechanisms for passing proposals and their supporting arguments in a way that other agents understand.

2. Techniques for *generating proposals* (counter-proposals or critiques) and for providing the supporting arguments;

   Proposal generation involves two main activities: (i) instantiating the negotiation

2. Poorly designed argumentation systems also have the potential to increase the length of the negotiation as the various merits of arguments and counter-arguments are debated. However, poor design of the other aspects of the negotiation technology can have similarly adverse effects, and so it is not something specific to argumentation-based negotiation.

3. For example, if arguments are preferred if they are more likely to lead to an agreement (which requires some metric on the agreement space) it is possible to prove that argumentation leads to quicker agreement [12].

object in accordance with the agent's acceptability region and its rating function; (ii) determining which argument(s) should accompany the agreement (if any) in order to maximise the likelihood of it being accepted. The complexity of the former point is determined by the nature of the strategic reasoning which is appropriate for the given negotiation protocol. This may vary from little reasoning, to maintaining complex models of negotiation opponents and trying to make predictions from them. In terms of the latter point, in the majority of cases there will be many types of argument which can be made in support of a proposal (varying from explanations to threats). In determining which ones to send, the agent needs to pick those arguments which are most likely to be effective, but *within the constraints of the agent's negotiation objectives*. Thus, for example, continually issuing threats may provoke short-term gains, but may not be a good long-term strategy if the agent has to interact frequently with the same group.

3. Techniques for *assessing proposals* (counter-proposals or critiques) and their associated supporting arguments;

   Received proposals need to be evaluated to determine how the agent should respond. This evaluation involves two main facets: (i) assessing the desirability of the proposal contained in the negotiation object; (ii) assessing the likely impact of the supporting arguments. From this, a number of potential outcomes are possible: the negotiation object is acceptable as it stands, the negotiation object alone is unacceptable but the supporting arguments overcome this and make the proposal acceptable, or the negotiation object is unacceptable and the supporting arguments are insufficient to warrant proposal acceptance[4]. Having assessed the proposal, the agent may decide to update its acceptability region or rating function to reflect the incoming proposal's arguments.

4. Techniques for *responding to proposals* (counter-proposals or critiques) and their associated supporting arguments;

   Having assessed a proposal, the agent can respond by accepting it, by rejecting it, by generating a critique, or by returning a counter-proposal. So the first functional requirement is to determine which of these courses of action should be taken. In the case of a critique, the agent has to determine what components it wants to accept and which it wants to reject, which issues it intends to provide constraints on, and what such constraints should be. It must then decide what arguments (if any) it will offer in support of this stance, and how it should respond to any arguments which accompanied the incoming proposal (varying from ignoring them to trying to undermine them). Counter-proposals are handled in a broadly similar manner, except that rather than giving feedback and constraints the agent has to instantiate the negotiation object with particular values.

While these argumentation specific capabilities undoubtedly increase the complexity of the agent, we feel such efforts are justified by the increased rewards which argumentation-based negotiation promises.

---

[4]. The fourth case of the proposal being acceptable and the argument making it unacceptable is highly unlikely to occur in practise.

4. Progress to Date

The categories given in the previous section can be used as a framework within which to assess our work to date on argumentation-based negotiation.

We have identified the basic mechanisms for passing proposals and their supporting arguments [10], and have an implementation of these mechanisms which allows simple negotiating agents to be easily constructed [1]. This work satisfies requirement 1 above.

We have experimented with various mechanisms for generating proposals [1] [9] (albeit not in situations in which arguments are taken into account) and with some simple mechanisms for providing arguments (where these are built by creating plans for joint action and using the plans as arguments for their adoption) [6] [7]. Thus we have done some work on requirements 2(i) and 2(ii), though the generation of supporting arguments is an area of continuing research.

We have looked at the use of various mechanisms for assessing the worth of proposals [1] [9]. We have also considered a number of ways of assessing the worth of arguments which support proposals, both those based on domain specific information [10], and more general mechanisms [6] [7], which rate arguments by considering what counter-arguments can be built against them. This work addresses requirements 3(i) and 3(ii).

Finally, we have considered ways to respond to proposals [1] [9] when just trading proposals, and have considered how to respond to arguments when response is guided by trying to find arguments which will defeat those already made by the other agents in the negotiation [6] [7]. This work addresses requirements 4(i) and 4(ii).

As can be seen from this, our work to date has looked at all the areas we consider important from the point of view of developing agents capable of argument-based negotiation. However, sadly, this does not mean that we are close to being able to build such agents! Before we can, we need to (a) extend our work on building, assessing and responding to arguments so that we can carry out such activities when proposals are more realistic than the simple joint plans we have worked on so far and (b) extend our implementation to be able to automate this behaviour.

5. Conclusions

This paper has argued for the centrality of negotiation in multi-agent systems research and has provided an informal framework for describing its key features. A case has been made for negotiation to include arguments to justify positions and to persuade opponents to change their stance. The key components which need to be addressed when building an argumentation-based negotiator have been outlined and our progress in this endeavour have been presented.

References

[1] A. Ayodele and S, Parsons (1998) “A framework for building negotiating agents”, Technical Report, Department of Electronic Engineering, Queen Mary and Westfield College.

[1] S. Bussmann and H. J. Müller (1992) “A Negotiation Framework for Cooperating Agents” Proc. CKBS-SIG, Keele, UK.

[2] M. Karlins and H. I. Abelson (1970) “Persuasion” Crosby Lockwood and Son.

[3] S. Kraus, M. Nirkhe and K. Sycara (1993) “Reaching agreements through argumentation: a logical model (Preliminary report)” Proceedings of the DAI Workshop, 233-247.

[4] B. Laasri, H. Laasri, S. Lander and V. Lesser (1992) “A generic model for negotiating agents” International Journal of Intelligent and Cooperative Information Systems, 1(2), 291-317.

[5] R. Loui and D. Moore (1998) “Dialogue and Deliberation” Negotiation Journal (submitted).

[6] S. Parsons and N. R. Jennings (1996) “Negotiation Through Argumentation—A Preliminary Report” Proc. 2nd Int. Conf. on Multi-Agent Systems, Kyoto, Japan, 267-274.

[7] S. Parsons, C. Sierra and N. R. Jennings (1998) “Agents that reason and negotiate by arguing” Journal of Logic and Computation 8 (3). (to appear)

[8] J. S. Rosenschein and G. Zlotkin (1994) “Rules of Encounter” MIT Press.

[9] C. Sierra, P. Faratin and N. R. Jennings (1997) “A Service-Oriented Negotiation Model between Autonomous Agents” Proc. 8th European Workshop on Modelling Autonomous Agents in a Multi-Agent World (MAAMAW-97), Ronneby, Sweden, 17-35.

[10] C. Sierra, N. R. Jennings, P. Noriega, and S. Parsons (1997) “A Framework for Argumentation-Based Negotiation” Proc. 4th Int. Workshop on Agent Theories, Architectures and Languages (ATAL-97), LNAI 1365, Rhode Island, USA 177-192, Springer Verlag.

[11] K. Sycara (1989) “Argumentation: Planning other Agents’ Plans” Proc 11th Int. Joint. Conf on AI, 517-523.

[12] F. Tohme (1997) “Negotiation and defeasible reasons for choice”, Proc AAAI Spring Symposium on Qualitative preferences in deliberation and practical reasoning, 95-102.

[13] N. Vulkan and N. R. Jennings “Efficient Mechanisms for the Supply of Services in Multi-Agent Environments” Submitted for publication.