

Automated Negotiation

N.R.Jennings¹, S.Parsons², C.Sierra³ and P.Faratin⁴

¹Dept.ofElectronicsandComputerScience,UniversityofSouthampton,
SouthamptonSO171BJ,UK.
nrj@ecs.soton.ac.uk

²Dept.ofComputerScience,UniversityofLiverpool,
LiverpoolL697ZF,UK.
s.d.parsons@csc.liv.ac.uk

³ArtificialIntelligenceResearchInstitute,SpanishScientificResearchCouncil,
CampusUAB,08193Bellaterra,Barcelona,Spain.
sierra@iia.csic.es

⁴Dept.ofElectronicEngineering,QueenMary&WestfieldCollege,
UniversityofLondon,LondonE14NS,UK.
p.faratin@qmw.ac.uk

1. Introduction

Interactions are a core part of all multi-agent systems. They occur because of the inter-dependencies that inevitably exist between the agents and they manifest themselves in many different forms—including cooperation, coordination, and 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 another 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 validations 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, time, terms and conditions, etc.). Orthogonal to the agreement structure is the issue of the types of operation that 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. 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. [16][23]). 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 that 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. [18][21]). 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 should be 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 that vary widely depending on the negotiation context. The aim of this paper is to briefly examine the space of negotiation opportunities and to identify some of the key techniques in the major areas.

2. A Generic Framework for Automated 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 separated 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 [5], [11], [12].

For a given negotiation, the participants are the active components that 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 agreement 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 negotiating agent poses some part of the agreement space as being acceptable; and (ii) indicating whether it is acceptable. In other words, the minimum negotiating agent is the ability to make and respond to proposals. Now since the context of agents reaching agreements about some joint problem, we have a solution to that joint problem; either a single complete proposed solution, or a group of complete or partial solutions. In terms of the different kinds of proposals become a single point, a region of the space, a region 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 (one 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, the 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

need 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 dislikes¹) or a *counter-proposal* (an alternative proposal generated 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 conditions.

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 that 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 service Z.

A: I propose that I provide you with service Y if you provide me with service X.

B: I propose that I provide you with service X if you provide me with service 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:

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.

- *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 that 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 car's 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 [8][10][21] 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 to increase the likelihood and/or the speed of agreements being reached². 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 ease negotiating over it.

3. Negotiation Techniques

Given this broad space of possibilities, this section seeks to describe three specific approaches to automated negotiation. These approaches are exemplars, with which the authors have been involved, of the full spectrum of opportunities. Each approach is briefly outlined and pointers to more detailed material are provided.

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 [22].

3.1 Game Theoretic Models

This line of work employs techniques and insights from game theory in order to structure and organise negotiations between autonomous agents. In particular, we have shown how one-to-many negotiations can be set up as an English auction for negotiation objects of multiple dimensions [23]. This model is illustrated by applying it to the real-world problem of business process management [6]. The model relies on agents playing dominant strategies and is computationally efficient. Moreover, it can be shown that the developed protocol produces optimal results for the buyer in terms of the amount of revenue it receives.

3.2 Heuristic Approaches

While game theoretic techniques work well in many cases, they also embody a number of assumptions that can be limiting for real-world applications. In particular, these models are often based on notions of perfect rationality (requiring the agent to be computationally unbounded and have full information of both its own and its opponents negotiation options) and they provide limited flexibility in cases where the designer cannot *a priori* impose a negotiation strategy upon the agent. In such cases, heuristic approaches are more suitable [9].

To this end, we have developed a rich suite of negotiation algorithms, based on multi-attribute utility theory, in which an agent has a negotiation strategy (high level objective about *how to* negotiate in a given encounter) and a family of negotiation tactics as a way of fulfilling the strategy [1]. These strategies and tactics need not be fixed at design time, they can be made to evolve during the course of negotiation [13]. Our model allows three broad classes of negotiation behaviour: concession making [19], making trade-offs between negotiation issues [4], and dynamically introducing new negotiation issues into an ongoing encounter [2], [18]. These models have been evaluated empirically [1], [4] and have been applied in a number of real-world scenarios including business process management [7] and telecommunications network management [3].

3.3 Argumentation-Based Approaches

In the majority of cases, heuristic models do not include a meta-level component for argumentation-based negotiation (although, in theory, there is nothing to preclude this). In our argumentation work, however, we adopt a logic-based approach [14]. In particular, we have developed a generic argumentation protocol, along with the necessary languages to support argumentation using this protocol [20]. These languages enable agents to augment their bargaining negotiation proposal with promises of threats or rewards, as well as a variety of forms of appeal. We have also investigated how this protocol can be integrated with mechanisms for making proposals with arguments supporting their acceptability [15]. This framework has been implemented using a multi-context system to represent the internal components of the agent's reasoning model [17].

4. Conclusions and Future Work

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. This framework has been used to identify three key methods of approach that the authors have been involved in developing. For the future, additional work is needed on techniques that allow designers to make informed choices about which negotiation models are appropriate in which circumstances and in allowing the agents themselves to alter the negotiation mechanism at run-time to better suit their prevailing circumstances. The underlying information in both of these cases will be a mixture of analytical and empirical data and it can only be obtained by following a broad-based research agenda that incorporates the full range of negotiation methods.

References

- [1] P. Faratin, C. Sierra, and N. R. Jennings (1998) "Negotiation Decision Functions for Autonomous Agents" *Int. Journal of Robotics and Autonomous Systems* 24(3-4)159-182.
- [2] P. Faratin, C. Sierra, N. R. Jennings and P. Buckle (1999) "Designing Responsive and Deliberative Automated Negotiators" *Proc. AAAI Workshop on Negotiation: Settling Conflicts and Identifying Opportunities*, Orlando, FL, 12-18.
- [3] P. Faratin, N. R. Jennings, P. Buckle and C. Sierra (2000) "Automated Negotiation for Provisioning Virtual Private Networks using FIPA-Compliant Agents" *Proc. 5th Int Conf. on Practical Application of Intelligent Agents and Multi-Agent Systems (PAAM-2000)*, Manchester, UK.
- [4] P. Faratin, C. Sierra and N. R. Jennings (2000) "Using Similarity Criteria to Make Negotiation Trade-Offs" *Proc. 4th Int. Conf. on Multi-Agent Systems (ICMAS-2000)*, Boston, USA.
- [5] N. R. Jennings, S. Parsons, P. Norrie, and C. Sierra (1998) "On Argumentation-Based Negotiation" *Proc. Int Workshop on Multi-Agent Systems*, Boston, USA.
- [6] N. R. Jennings, P. Faratin, T. J. Norman, P. O'Brien and B. Odgers (2000) "Autonomous Agents for Business Process Management" *Int. Journal of Applied Artificial Intelligence*, 14(2).
- [7] N. R. Jennings, P. Faratin, T. J. Norman, P. O'Brien, B. Odgers and J. L. Alty (2000) "Implementing a Business Process Management System using ADEPT: A Real-World Case Study" *Int. Journal of Applied Artificial Intelligence*, 14(3).
- [8] M. Karlins and H. I. Abelson (1970) "*Persuasion*" Crosby Lockwood and Son.
- [9] S. Kraus (1997) "Negotiation and Cooperation in Multi-Agent Environments" *Artificial Intelligence* 94:79-97.
- [10] S. Kraus, K. Sycara and A. Evenchik (1998) "Reaching agreements through argumentation: a logical model and implementation" *Artificial Intelligence* 104:1-69.
- [11] B. Laasri, H. Laasri, S. Lander and V. Lesser (1992) "A generic model for negotiating agents" *Int. Journal of Intelligent and Cooperative Information Systems*, 1(2), 291-317.
- [12] R. Loui and D. Moore (1998) "Dialogue and Deliberation" *Negotiation Journal* (submitted).
- [13] N. Matos, C. Sierra and N. R. Jennings (1998) "Determining successful negotiation strategies: an evolutionary approach" *Proc. 3rd Int. Conf. on Multi-Agent Systems (ICMAS-98)*, Paris, France, 182-189.
- [14] 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.

- [15] S.Parsons,C.SierraandN.R.Jennings(1998)“Agents that reason and negotiate by arguing” *Journal of Logic and Computation* 8(3)261-292.
- [16] J.S.Rosenschein and G.Zlotkin(1994)“ *Rules of Encounter* ”MIT Press.
- [17] J.Sabater,C.Sierra,S.Parsons and N.R.Jennings(1999)“Using multi-context systems to engineer executable agents” *Proc.6th Int. Workshop on Agent Theories Architectures and Languages (ATAL-99)* Orlando,FL.131-148.
- [18] C.Sierra,P.Faratin and N.R.Jennings(1999)“Deliberative Automated Negotiators Using Fuzzy Similarities” *Proc EUSFLAT-ESTYLF Joint Conference on Fuzzy Logic, Palma de Mallorca, Spain*,155-158.
- [19] 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.
- [20] 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.
- [21] K.Sycara(1989)“Argumentation: Planning other Agents’ Plans” *Proc 11th Int. Joint. Conf on AI* ,517-523.
- [22] F. Tohme (1997)“Negotiation and defeasible reasons for choice”, *Proc AAAI Spring Symposium on Qualitative preferences in deliberation and practical reasoning* ,95-102.
- [23] N. Vulkan and N.R.Jennings(2000)“Efficient Mechanisms for the Supply of Services in Multi-Agent Environments” *Int Journal of Decision Support Systems* .(to appear)