Designing Flexible Automated Negotiators: Concessions, Trade-Offs and Issue Changes

P. Faratin and **C.Sierra**
QMW, University of London,UK.
IIIA (CSIC), Spain.

N.Jennings and **P. Buckle**
QMW, University of London, UK.
Nortel Networks, UK.

Abstract

In this paper we present a multi agent and multi issue negotiation model which can be used for guiding agents during distributed problem solving. This model is composed of protocols which govern and manage agent interactions, and an agent architecture which represents decision and action mechanisms which assist agents during distributed problem solving processes.

1 Introduction

Multi Agent System technology is concerned with deriving cooperative solutions among a set of distributed autonomous agents. Cooperation between agents is needed either because real world problems are often distributed in nature and interdependencies between the sub problems cannot be avoided, or because agents have limited rationality (agents maybe computationally resource bounded or not have the necessary procedural knowledge to perform the required tasks [Russell and Wefald, 1991]). This leads to the need for interactions between agents in order to agree on a certain course of action.

However, cooperation cannot readily be assumed since agents may be autonomous —they decide for themselves what, when and under what conditions their actions should be performed. Since agents have no direct control over others, they must persuade others to act in a particular manner. The type of persuasion we consider in this paper is negotiation which we define as *a process by which a joint decision is made by two or more parties. The parties first verbalise contradictory demands and then move towards agreements by a process of concession making or search for new alternatives* [Pruitt, 1981].

More concretely, we view distributed problem solving as a community of agents which provide and consume *services* to and from one another. A service is a grouping of subtasks which must be performed in order to achieve some portion of the overall problem. This *service-oriented negotiation* involves determining a mutually acceptable contract for certain terms and conditions.

Developing a system of interacting service orientated agents requires the specification of a coordination protocol which dictates agent interactions. Traditionally, formal models of choice achieve coordination through specification of the negotiation space: the issues agents negotiate over, and their possible values that determine the set of alternative solutions. Negotiation is then considered as an optimisation problem, where given the utility function of the agents, the best solution is obtained. This methodology is often adopted in Game Theory.

However, such formal models of choice not only ignore *interactions*, but often involve unrealistic assumptions (such as common knowledge). An alternative coordination methodology is the specification of the rule (or order) of interaction —*who* can say *what* and *when* (absence of any normative rule of behaviour may lead to unmanageable interactions). In this paper we follow this second approach, and ensure coordinated behaviour by defining our protocol as an extension of the normative rules of the Contract Net Protocol [Davis and Smith, 1988], for instance by permiting iterated sequences of offers and counter-offers.

Individual agents, in turn, must be provided with the capability to represent and reason about both their internal and their external world and with the capacity to interact according to the above protocol. It is this individual agent modelling which has been the central focus of the work reported in this paper.

The developed negotiation protocol is first introduced in section 2. The individual agent architecture is then expanded on in section 3 which describe evaluation (section 3.1) and offer generation mechanisms (sections 3.2 and 3.3). The presented model is then compared, in section 4, with other developed models. Finally, section 5 presents the conclusions reached.

2 The Negotiation Protocol

The protocol (see figure 1) starts by a dialogue to establish the conditions for negotiation (deadline, initial issues, etc.). Then, one of the agents makes an offer (transition from state 1 to state 2). After that, the other agent can make a counteroffer or a tradeoff —see section 3.3— (moving to state 3), and the agent that started

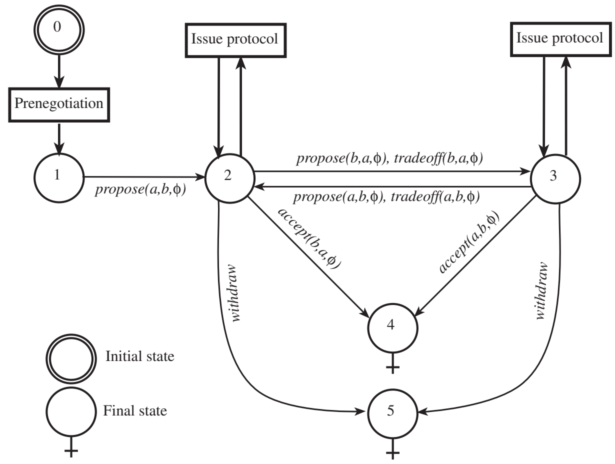

Figure 1: Negotiation protocol.

the negotation can in turn make a new counteroffer or a new tradeoff (going back to state 2). In either of these two states, one of the agents may simply accept the last offer being made by the oponent or any of the agents may withdraw negotiations. At any moment agents are permitted to start an ellucidatory dialogue to establish a new set of issues to negotiate over (see section 3.3). This protocol is a natural extension of the contract net protocol permitting iterated offer and counter-offer generation and permitting the modification of the set of issues under negotiation.

3 Agent Negotiation Architecture

The main contribution of the research reported here is the specification of a negotiation architecture that structures the individual agent's reasoning throughout the problem solving. The agent model is founded on notions borrowed from psychological, economic (game theory in particular), decision and social welfare models [Corfman and Gupta, 1993]. Rational behaviour is assumed to consist of maximisation of some value function [Raiffa, 1982]. Given this rationality stance, the decisions faced by agents in negotiation are often a combination of: *offer generation decisions* (what initial offer should be generated, what counter offer should be given in situations where the opponent's offer is unacceptable), and *evaluatory decisions* (when negotiation should be abandoned, and when an agreement is reached). Agents must therefore be capable of representing and reasoning about their goals, actions, plans and knowledge. The solution to these decision problems is captured in the agent architecture. The components (or what is referred to as mechanisms) of the agent architecture which is responsible for generation of offers and counter offers are based on the above given definition of negotiation — *"...a process of concession making or search for new alternatives"* [Pruitt, 1981]. "Concession making" is then modelled as a process which results in progressively decreasing contract value (section 3.2) and "search for new alternatives", as issue trade off (section 3.3) and issue set manipulation (section 3.3). The mechanisms which assist an agent with evaluation of offers[1], is described first, followed by the developed concession and search mechanisms in sections 3.2 and 3.3 respectively.

3.1 Evaluation Mechanisms

In service-oriented negotiation agents can adopt one of two, possibly conflicting, roles: $Buyers$ and $Sellers$. Negotiation between these types of agents ranges over a number of quantitative (e.g. price, duration, volume, etc.) and quali. The set of issues, J, is allowed to change through the negotiation (section 3.3). Quantitative issues are defined over a real domain (represented as vector positions: $x[j] \in \mathcal{D}_j = [min_j, max_j]$). Qualitative issues are defined over a fully ordered domain (i.e. $x[j] \in \mathcal{D}_j = \langle q_1, \ldots, q_n \rangle$).

Choice between contracts is modelled as a *value* decision[2]. When an agent a receives an offer ($x = (x[j_1], \ldots, x[j_n])$ where $j_i \in J$), it rates the overall contract value using the following weighted additive scoring function:

$$V^a(x) = \sum_{1 \le i \le n} w^a_{j_i} V^a_{j_i}(x[j_i])$$

where $w^a_{j_i}$ is the importance (or weight) of issue j_i such that $\sum_{1 \le i \le n} w^a_{j_i} = 1$. Given that the changing of the set of issues during negotiation is permitted, agents will need to dynamically change the values of the weights. The score of value $x[j]$ for agent a, given the domain of acceptable values $\mathcal{D}_j$, is modelled as a scoring function $V^a_j : \mathcal{D}_j \to [0, 1]$. For convenience, scores are bounded to the interval $[0, 1]$, and the scoring functions are monotonous for quantitative issues.

Contract evaluation mechanism has to consider two elements: the contract received and the different mechanisms to propose alternatives to the other agent. These mechanisms are in our case: i) preparing a counter-offer, ii) proposing a trade-off, or iii) trying to chose a new set of issues to negotiate. Each of the mechanisms has a different impact in terms of utility gain or loss. The evaluation function (that will vary depending on each agent) will determine which way to follow: accepting the received contract or proposing an alternative. These different mechanisms are presented in detail in sections below.

3.2 Concession Mechanisms

Concession mechanisms generate offers which have successively lower scores than previous offers according to a set of environment criteria. Offers are generated by

[1] Much of the design rationale and the supporting arguments for evaluatory and concession mechanisms are given in [Faratin *et al.*, 1998] and the reader is referred to these sources for an in-depth exposition.

[2] Note, there exists a difference between value and utility. See [Keeney and Raiffa, 1976] for an in-depth exposition.

a linear combination of simple functions, called *tactics* [Faratin *et al.*, 1998]. Tactics generate values for issues using only a single criteria (e.g. time, resources, etc.), and are grouped into three classes:

- **Time-dependent tactics** These tactics model increasing levels of concession as the deadline for the negotiation, t^a_{max}, approaches. The offer agent a sends to agent b for issue j at time $0 \le t \le t^a_{max}$, is modelled by a function $\alpha^a_j : T \to [0,1]$, where $\alpha^a(0) = 0$, and
$\alpha^a(t^a_{max}) = 1$. Two possible families of functions (parameterised by β) which model α are:
  - Polynomial function: $\alpha^a(t) = \left(\frac{t}{t^a_{max}}\right)^{\frac{1}{\beta}}$
  - Exponential function: $\alpha^a(t) = \begin{cases} e^{(1-\frac{t}{t^a_{max}})^{\beta} * ln(\frac{t}{t^a_{max}})} & \beta > 1 \\ 1 - e^{(\frac{t}{t^a_{max}})^{1/\beta} * ln(1-\frac{t}{>}t^a_{max})} & \beta < 1 \end{cases}$

  Both families are parameterised by a value $\beta \in I\!R^+$ that determines the convexity of the concession function. After computing $\alpha^a_j(t)$, the offer x agent a sends to agent b for issue j is modelled as:

  $$x^t_{a\to b}[j] = min^a_j + \alpha^a_j(t)(max^a_j - min^a_j)$$

  if V^a_j is decreasing, or

  $$x^t_{a\to b}[j] = min^a_j + (1 - \alpha^a_j(t))(max^a_j - min^a_j)$$

  if V^a_j is increasing.

- **Resource-dependent tactics** Thesetactics model increasing levels of concession with diminishing levels of resources, such as time:

  $$\alpha^a_j(t) = e^{-(t^a_{max} - t)}$$

  The offer x is modelled in a similar fashion to Time-dependents tactics.

- **Behaviour-dependent tactics** Concession here is based on the concessions of the other negotiating party [Axelrod, 1984]. New offers are computed in the following manner:

  $$x^t_{a\to b}[j] = \begin{cases} min^a_j & \text{If } P \text{ is } \le min^a_j \\ max^a_j & \text{If } P \text{ is } > max^a_j \\ P & \text{Otherwise} \end{cases}$$

  A number of P expressions have been designed that vary on which aspect and to what degree they imitate the opponent's behaviour [Faratin *et al.*, 1998]. One such function is *Relative Tit-For-Tat*, which reproduces, in percentage terms, the behaviour that its opponent performed $\delta \ge 1$ steps ago:

  $$P = \frac{x^{t_{n-2\delta}}_{b\to a}[j]}{x^{t_{n-2\delta+2}}_{b\to a}[j]} x^{t_{n-1}}_{a\to b}[j]$$

However, to determine the best course of action agents may need to consider and assess more than just one environmental condition. Each tactic generates a value for an issue using only a single criterion.

There maybe circumstances where multiple criteria are deemed important in computing a value for an issue —multiple tactics maybe needed to perform this computation. The concept of *strategy* is introduced to model this requirement. Strategies are a way of assigning a weight to a set of tactics reflecting their relative importance. Strategies are also responsible for modification, over time, of these weights as the criteria change their relative importance in response to environmental changes.

3.3 Search Mechanisms

The concession mechanism discussed above models utility decay in a manner that can be varied according to a unique, or a combination of, environmental condition/s. However, negotiation does not always entail loss of utility —negotiation can also involve search for new alternatives. Indeed, if an agent behaves according to the principle of maximisation of value then it should clearly prefer to offer only the best possible contract it deems appropriate and not change its offer throughout the negotiation process. Based on empirical results (see [Faratin *et al.*, 1998]) and real world negotiations (see [Raiffa, 1982]), such a strategy often terminates in unsuccessful outcomes. Whereas concession mechanisms, through losing utility, allow negotiating parties to find agreements, there are other mechanisms which are extensively used in real life negotiations which allow agents to construct new offers *without* necessarily losing utility. Two such mechanisms are proposed: *trade offs* and *issue set manipulations*.

Trade Off Mechanisms

Successful real world negotiation often involves recognition of how to manipulate acceptance levels on a set of issues during negotiation in a manner that is acceptable to both parties. Concession over issues is one method. The other is trade off, where one party lowers it's acceptance level on some issues and simultaneously demands more on other issues. The rationale behind trade off mechanisms is the design of a decision mechanism which models an agent's desire to increase the likelihood of a successful outcome while at the same time maximising it's own contract value. To this end, this paper introduces a decision mechanism which involves searching all possible contracts with the same score as the previously offered contract (hence there is no loss in contract utility) and selection of the contract which is the "closest" to the opponent's last contract offer.

Trade off mechanisms generate new contracts that lie on what is called the iso-value (or indifference) curves [Raiffa, 1982]. Because all newly generated contracts lie on the same iso-value curve then agents are indifferent between any two given contracts on this curve. An iso-curve is then defined as:

Definition 1 *Given a scoring value θ, the iso-curve set at degree θ for agent a is defined as:*

$$iso_a(\theta) = \{x \mid V^a(x) = \theta\}$$

Theory of fuzzy similarity is used in order to model "closeness". The best tradeoff then would be the most similar contract on the iso-curve. More formally, tradeoff is defined as:

Definition 2 *Given two consecutive offers $x^t_{a\to b}$ and $x^{t+1}_{b\to a}$, with $\theta = V^a(x^t_{a\to b})$, trade off for agent a with respect to $x^{t+1}_{b\to a}$ is defined as:*

$$tradeoff_a(x^{t+1}_{b\to a}) = \arg_x \max_{x\in iso_a(\theta)} \{Sim(x, x^{t+1}_{b\to a})\}$$

Similarity between two contracts is defined as the weighted combination of the similarity of the issues:

Definition 3 *The similarity between two contracts x and y over the set of issues J is defined as:*

$$Sim(x,y) = \sum_{j\in J} w^a_j Sim_j(x[j], y[j])$$

With, $\sum_{j\in J} w^a_j = 1$. Sim_j is the similarity for issue j.

Following the results from [Valverde, 1985], a similarity function can be defined which satisfies the axioms of reflexivity, symmetry, and t-norm transitivity as a conjunction (modelled as the infimum) of appropriate fuzzy equivalence relations induced by a set of criteria functions h_i. Hence the similarity between two values for issue j, $Sim_j(x,y)$ is defined as:

Definition 4 *Given a domain of values D_j, the similarity between two values $x, y \in D_j$ is defined as:*

$$Sim_j(x,y) = \bigwedge_{1\le i\le m} (h_i(x) \leftrightarrow h_i(y))$$

where $1 \le i \le m$ are a set of comparison criteria with $h_i : D_j \to [0,1]$, and $\leftrightarrow$ being an equivalence operator.

Simple examples of the $\leftrightarrow$ operator are $h(x) \leftrightarrow h(y) = 1- \mid h(x) - h(y) \mid$ or $h(x) \leftrightarrow h(y) = min(h(y)/h(x), h(x)/h(y))$. For example, a function that models the criteria of whether a price is low, $lowprice : Price \to [0,1]$, could be:

$$lowprice(x) = \begin{cases} 1 & \text{x} < \text{£}10 \\ \frac{\text{£}20-x}{\text{£}10} & \text{£}10<\text{x}<\text{£}20 \\ 0 & x \ge \text{£}20 \end{cases}$$

Issue Set Mechanisms

Negotiation processes are directed and centred around the resolution of conflicts over a set of issues J. This set may consist of just one or more issues (distributed and integrative bargaining respectively). If more than one issue is involved in negotiation, then the negotiation protocol can either transform this integrative bargaining environment into distributed bargaining by forcing the agents to reach agreements through a process which "builds up a package"—agents make agreements on the elements of J one by one[3]. Alternatively the negotiation protocol may leave the negotiation set J unaltered and agents are then forced to make agreements on "packages". Choice of which protocol to use is important because both have different properties which directly influence the quality (in terms of joint as well as individual gains) and the convergence speed to a solution.

For simplification the ontology of the set of possible negotiation issues J, is assumed to be a shared knowledge amongst all the agents. It is further assumed that agents begin negotiation with a prespecified set of "core" issues, J^{core}, and possibly other mutually agreed none core set members, $J^{\neg core}$. Alterations to J^{core} is not permitted, but elements of $J^{\neg core}$ negotiation set can be altered dynamically. Agents can add or remove issues into $J^{\neg core}$ as search for new possible and up to now unconsidered solutions.

If J^t is the set of issues being used at time t (where $J^t = \{j_1, \ldots, j_n\}$), and $J - J^t$ is the set of issues not being used at time t, and $x^t_{a\to b} = (x_{a\to b}[j_1], \ldots, x_{a\to b}[j_n])$ is $a's$ current offer to b at time t, then issue set manipulations is defined throught two operators: *add* and *remove* which agents can apply to the set J^t. The *add* operator assists the agent in selecting an issue j' from $J - J^t$, and an associated value $x[j']$, which gives the highest score to the agent.

Definition 5 *The best issue to add to the set J^t is defined as:*

$$add(J^t) = \arg_j \max_{j\in J-J^t} \{ \max_{x[j]\in D_j} V^a(x^t_{a\to b}.x[j])\}$$

where . stands for concatenation.

An issue's score evaluation is also used to define the *remove* operator in a similar fashion to the *add* operator. Thus, this operator assists the agent in selecting the best issue, to remove from the current negotiation set J^t with the highest score.

Definition 6 *The best issue to remove from the set J^t (from $a's$ perspective), is defined as:*

$$remove(J^t) = \arg_{j_i} \max_{j_i\in J^t - J^{core}} \{V^a(x^{t+2}_{a\to b})\}$$

with $x^{t+2}_{a\to b} = (x^t_{a\to b}[j_1], \ldots, x^t_{a\to b}[j_{i-1}], x^t_{a\to b}[j_{i+1}], x^t_{a\to b}[j_n])$

Note, that unlike tradeoff mechanisms the value-based *add* operator can result in contracts that have varying scoring levels. Therefore, these mechanisms are not score maintainers but rather search mechanisms for new possible solutions.

[3] Modelling is further complicated by the introduction of commitment levels. For example, agreements made earlier can be tentative and subject to revision in the course of negotiation. Alternatively, all agreements are absolute. Throughout this paper it is assumed that commitments are absolute in all cases.

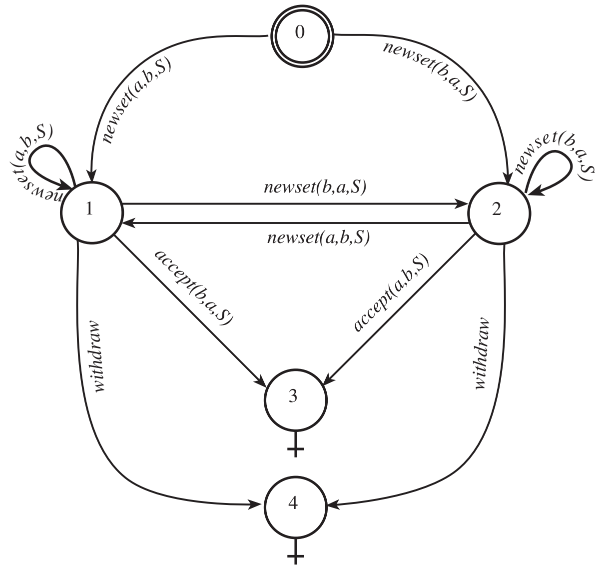

Figure 2: Pr otocol for issue manipulation.

The *remove* operator can also be defined in terms of the similiarity function defined above in section 3.3. This type of similarity based *remove* operator selects from two given threads $x^t_{a\to b}$ and $x^{t+1}_{b\to a}$, which issue to remove from $x^t_{a\to b}$ that maximises the similarity with respect to $x^{t+1}_{b\to a}$. Therefore, this mechanism can be considered as more cooperative. We define this similarity based *remove* operator as:

Definition 7 *The best issue to remove from b's perspective from the set J^t is defined as:*

$$remove(J^t) = \arg_{j_i} \max_{j_i \in J^t - J^{core}} \{sim((x^{t+1}_{b\to a}[j_1], \ldots, x^{t+1}_{b\to a}[j_{i-1}], x^{t+1}_{b\to a}[j_{i+1}], \ldots, x^{t+1}_{b\to a}[j_n]), (x^t_{a\to b}[j_1], \ldots, x^t_{a\to b}[j_{i-1}], x^t_{a\to b}[j_{i+1}], \ldots, x^t_{a\to b}[j_n]))\}$$

It is not possible to define a similarity based *add* operator since the introduction of an issue does not permit an agent to make comparisons with $x^{t+1}_{b\to a}$ simply because there is no value offered over that issue. Agents deliberate over how to combine these *add* and *remove* operators in a manner which maximises some measure —such as the contract score. However, it is acknowledged that the production of such a search tree of possible operator combinations for all issues can become combinatorially expensive. Another computational requirement by issue set mechanisms is the need for an agent to be able to dynamically recompute the issue weights since addition and substraction of issues indirectly influence an issue's relative weighting with other issues.

The protocol for establishing a new set of negotiating issues is straighforward. Each negotiating agent can start this dialogue, and from then on, each can either propose a new set, accept the other's proposed set or withdraw.

3.4 Mechanism Selection

The above three mechanisms are independent components of the agent architecture which can be used to represent and reason over internal and external states throughout problem solving. The resultant behaviour of application of these mechanisms is a rational and coherent behaviour within the rules of the protocol. For example, if the time to reach an agreement is fast diminishing, then concession mechanisms are the most rational mechanism choice. Alternatively, issues may be removed or added to escape deadlocks in negotiation. Mechanisms can be selected based on a number of states: the attitude of agents towards risk (concessionary mechanisms have a higher probability of success for example); proximity of the contract score to reservation value (trade off or issue set manipulation can become candidate mechanisms when the value of all issues being offered has reached reservation value); past cases (adding an issue has resulted in quicker agreements), etc. In general, the selection of which mechanism to use is based on a number of elements in the mental state of the agent.

4 Related Work

Negotiation has been studied in a number of related disciplines. However, the central focus of the work reported here, has been the design of a negotiation agent architecture for structured interactions in real environments over services. The theoretical primitives of the developed model is based on formal models of choice. However, the developed model does not adopt the approach of formal models either because there is often no actual mechanism solution (cooperative models) or because the assumptions are often inadequate (beliefs are not often common knowledge, and individuals are computationally bounded). Formal theories are often also a theory of behaviourism and excludes from the models any deliberative intervention over both the agent's internal state and it's interactions. In general, formal theory has failed to generate a general model governing rational choice in interdependent situations and has instead produced a number of highly special models applicable to specific types of inter-dependent decision making [Zeng and Sycara, 1997].

Contract Net protocol [Davis and Smith, 1988], however, is closer to the work reported here, where a protocol is used for modelling interactions. CNP is a distributed coordination architecture which models commitments, bounded rationality and incomplete knowledge. However, the CNP protocol does not allow agent interaction from a conflicting start point when in fact negotiation in many domains is often characterised by initial presence of conflict, which is then resolved through a negotiation process. The problem domain of this research is non-cooperative in that agents can be selfish. Furthermore, in non-cooperative domains the search for acceptable solutions maybe more elaborate than the CNP's two messages —negotiation is an iterative process. In addition to this, CNP is a theory of system architecture and is silent with respect to the individual agent architecture and consequently, like formal theories, is inadequate for agent design since any agent architecture is as good as another as long as they obey the CNP protocol. The proposed model in this paper not only specifies a negotiation protocol used for iterative interaction modelling but also provides mechanisms which agents can implement and execute according to their own requirements (for example, an agent may wish to use only tradeoff mechanisms throughout the negotiation.).

The iterative nature of negotiation, over multiple issues and agents, is modelled by the PERSUADER system through the concepts of argumentation and mediation [Sycara, 1989]. However, negotiation, as defined in this paper, is a *mutual* selection of outcome and precludes any intervention by outside parties. Furthermore, persuasion mechanisms operate on the beliefs of agents with the aim of changing one or both parties beliefs. This is not the case for negotiation —it is not necessary for the agents to have similar beliefs at the end of negotiation.

Other systems such as KASBASH is an attempt to actually engineer a real world application [Chavez and Maes, 1996]. The system models time, actions and strategies involved in negotiation. However, negotiation in KASBAH is over a single issue and agents are semi-autonomous —the system models only a subset of the decision making which is involved in negotiation and the user makes all the other decisions. Furthermore, the decisions that are delegated to the agents (called strategies in Kasbah) is severely limited to only three and even their selection is not autonomous. The model presented in this paper handles multiple issues and is designed for fully autonomous agents.

5 Conclusions

This paper has presented a distributed negotiation model which coordinates both agent interactions and individual agent decisions. Protocols have been defined which structure interactions and model the iterated nature of search for solutions. Mechanisms have been proposed for finding solutions which are based on realistic assumptions, are practical and model the complex nature of negotiation. Other approaches to modelling negotiation have often tended to view negotiation in a much simpler manner, whereas the number and behaviour of the proposed decision mechanisms reflect the complexity of what is usually referred to by the term negotiation.

The direction for future research will be primarily focused at empirical evaluation of the developed model to determine its properties. An application on telecommunication networks in collaboration with Nortel Network is currently being developed using the proposed model.

References

[Axelrod, 1984] R. Axelrod. *The Evolution of Cooperation.* Basic Books, Inc., Publishers, New York, USA., 1984.

[Chavez and Maes, 1996] A. Chavez and P. Maes. Kasbah: An agent marketplace for buying and selling goods. *Conference on Practical Applications of Intelligent Agents and Multi-Agent Technology*, 1996.

[Corfman and Gupta, 1993] K. P. Corfman and S. Gupta. Mathematical models of group choice and negotiations. *Handbooks in Operation Research and Management Sciences*, 5:83–142, 1993.

[Davis and Smith, 1988] R. Davis and R. Smith. Negotiation as a metaphor for distributed problem solving. In A. Bond and L. Gasser, editors, *Readings in Distributed Artificial Intelligence*, pages 333–357. Morgan Kaufmann Publishers, Inc., 1988.

[Faratin *et al.*, 1998] P. Faratin, C.Sierra, and N.Jennings. Negotiation decision functions for autonomous agents. *Robotics andAutonomous Systems*, (24):159–182, 1998.

[Keeney and Raiffa, 1976] R. L. Keeney and H. Raiffa. *Decisions with Multiple Objectives.* John Wiley and Sons, New York, 1976.

[Pruitt, 1981] D. G. Pruitt. *Negotiation Behavior.* Academic Press, 1981.

[Raiffa, 1982] H. Raiffa. *The Art and Science of Negotiation.* Harvard University Press, Cambridge, USA, 1982.

[Russell and Wefald, 1991] S. Russell and E. Wefald. *Do the Right Thing.* The MIT Press, 1991.

[Sycara, 1989] K. Sycara. Multi-agent compromise via negotiation. In L. Gasser and M. Huhns, editors, *Distributed Artificial Intelligence Volume II*, pages 119–139, San Mateo, California, 1989. Morgan Kaufmann.

[Valverde, 1985] L. Valverde. On the structure of F-indistinguishability. *Fuzzy Sets and Systems*, 17:313–328, 1985.

[Zeng and Sycara, 1997] D. Zeng and K. Sycara. How can an agent learn to negotiate. In J. Mueller, M. Wooldridge, and N. Jennings, editors, *Intelligent Agents III. Agent Theories, Architectures, and Languages*, number 1193 in LNAI, pages 233–244. Springer Verlag, 1997.