Motivation-Based Selection of Negotiation Opponents

Steve Munroe and Michael Luck

Electronics and Computer Science,
University of Southampton, Southampton, UK
{sjm01r, mml}@ecs.soton.ac.uk

Abstract. If we are to enable agents to handle increasingly greater levels of complexity, it is necessary to equip them with mechanisms that support greater degrees of autonomy. This is especially the case when it comes to agent-to-agent interaction which, in systems of selfish agents, often follows the format of negotiation. Within this context, a problem which has hitherto received little attention is that of identifying appropriate negotiation opponents. Furthermore, the problem is particularly difficult in dynamic systems where the need to negotiate over issues and the evaluation of resources may change over time. Such dynamics demand high degrees of autonomy from agents so that such factors can be handled at run-time and without the aid of human controllers. To that end, this paper draws inspiration from biological organisms and theories of motivation, and describes a motivation-based architecture comprising a number of motivation-based classification and selection mechanisms used to evaluate and select between negotiation opponents. Opponents are evaluated in terms of the likely issues they will want to negotiate over and the amount of conflict this might entail. Additionally, the expected cost of a negotiation with an opponent is examined in relation to the agent's current motivational evaluation of its resources. The mechanisms allow prioritisation between each method of evaluation dependent upon motivational needs. Some preliminary evaluation of the model is also presented.

1 Introduction

Negotiation is a particularly important form of interaction between agents. It allows agents to reach agreement over shared concerns, and there are many existing frameworks (e.g. [1, 2, 3]). Most frameworks focus on the problems inherent within the negotiation episode, such as which negotiation strategies and tactics offer the best results, and how best to employ them. They provide a host of techniques and methods that allow agents to autonomously navigate through a negotiation episode. However, though the actual steps taken within a negotiation are often left to the agent to decide, the form and focus of the negotiation is still most often handled by some user of the agent. Thus, though an agent may be used to negotiate the conditions of a holiday package for example, the

M.-P. Gleizes, A. Omicini, and F. Zambonelli (Eds.): ESAW 2004, LNAI 3451, pp. 119–138, 2005.

issues negotiated over, the constraints on what are acceptable outcomes, and the opponents negotiated with, are often not decided by the agent, but are presented as given constraints. This may be fine in such purchase negotiation as this but, in persistent, multi-agent systems where agents must go about their tasks away from human direction, this is clearly not adequate. To address this, agents must be given the ability to decide what they need to negotiate about, what constraints impact upon those needs, and which opponents best meet these needs. In fact, there is a growing realisation that the decisions that must be made prior to negotiation represent key problems that must be addressed if agents are to use negotiation more autonomously than they do at present (e.g.[4]).

1.1 Opponent Selection for Negotiation

When the need for negotiation arises, there may be a number of potential negotiation opponents that an agent can choose from. Each opponent may exhibit different service-related characteristics that distinguish it from other opponents. For negotiation to be successful, the preferences of the agent requiring the service must coincide with the service-related characteristics of the opponent. For example, an agent needing to negotiate on the time of delivery of some service, is better served by an opponent who does not have a preference for this issue. In this way, the issue can be settled quickly and at the value most preferred by the agent. In other words, the agent should attempt to avoid those opponents who have strong preferences regarding service delivery times, as this may risk conflict on the issue, potentially leading to a negotiation that is long, difficult and prone to failure. Thus, an agent should have, in order to increase its chances of engaging in successful negotiations, the ability to reason about the compatibility of an opponent's interests with its own in order to avoid negotiations with a high *conflict* potential.

An additional factor that impacts upon opponent choice is the *relative importance* of the negotiation issues chosen. For example, an agent looking to negotiate the purchase of a service, with only limited monetary resources, will place a greater value on that resource and so will be influenced towards those opponents offering cheaper prices. Good performance on price is, therefore, a more influential selection criterion than other issues, and opponent selection should reflect this. At different times, and with different constraints, however, opponent selection may be influenced by different criteria, such as the *quality* of the offered service.

1.2 Autonomy and Motivation

While existing negotiation frameworks allow for autonomous agent behaviour within negotiation, they mostly omit consideration of it before negotiation begins. However, if agents are to make the kinds of decisions that must be made before negotiation begins, they must be able to display autonomy here as well. Our approach to enable such autonomy is to adopt the construct of *motivation*, which is defined in [5] as

> "*any desire or preference that can lead to the generation and adoption of goals and that affects the outcome of the reasoning or behavioural task intended to satisfy those goals.*"

Motivation influences the decision-making of an agent, leading to the adoption of those activities that best serve its motivational interests. This approach is similar to the more commonly used economic approach that uses the notion of *utility* to guide agent activity. Utility-based agents act under the principle of utility-maximisation, in which activities with higher utility are chosen over those with lower utility. However, whereas utility is an economic abstraction of value or benefit that is overlaid on an agent's choices by the agent designer, motivation is an internally derived measure of value determined both by a set of internal state variables (such as hunger or thirst, for example) and the external environment. For example, in the presence of food, an agent may or may not choose to eat depending on the state of its internal environment (specifically its hunger motivation). In this sense, motivation grounds the generation of measures of value (such as utility) in the agent's internal state, and thus is in a sense prior to, and generative of, such notions. By examining options and weighing up their motivational worth, an agent can be guided in choosing motivationally relevant activities, and it is exactly this behaviour that defines for us the essence of *autonomy*. An agent is autonomous if it makes decisions and selects courses of actions that *further its own interests* based upon *its own assessment* of the situation.

In this paper, we describe an approach to the selection of negotiation opponents based upon consideration of the likely amount of conflict that might result form a particular selection, and the extent to which an opponent is expected to meet the constraints that impact upon the agent's ability to settle the issue of price. We describe a *motivated agent architecture* and a number of *motivated decision-making mechanisms* that allow an agent to assess the various needs that it has regarding the issues of the forthcoming negotiation, and evaluate the how a particular opponent meets those needs. While the paper is free of formal description, a formal model of the motivated agent architecture exists and can be obtained from the authors. The paper proceeds as follows. In Section 2 related work is described. Section 3 describes our *motivated agent architecture* and our *negotiation goal model* comprising an attribute classification mechanism, and in Section 5 we discusses the kinds of information needed about opponents for selection to work. Section 6 describes the selection mechanisms that work by either *minimising conflict* or minimising *resource use*, and Section 7 presents some preliminary evaluation of the model. Finally, Section 8 offers some concluding remarks.

2 Related Work

Much work is currently been undertaken that examines opponent selection from the point of view of *trust* and/or *service reputation* (e.g. [6, 7]), where both refer to the fidelity of the opponent's behaviour with regard to the negotiated outcome. Whilst trust and reputation are of great importance for designers of agent systems, especially systems characterised by openness they are, we argue, only part of the story and must be supplemented with the kinds of issues we are investigating in this paper such as information relating to the underlying *interests* and *motivations* of the negotiation participants and how this influences the kinds of negotiation encounter they prefer.

Non trust-based opponent selection has been addressed by a number of researchers too. Work by Tesfatsion [8] examines how agents select opponents based upon the amount by which they exceed a fixed performance-based tolerance threshold. Though this work examines similar problems to those in this paper, it does not address the specific problems of the minimisation of conflict through the smart selection of negotiation opponents, and assumes fixed performance thresholds — whereas we deal with dynamically changing performance requirements. In [9], Banerjee *et al.* examine the formation of coalitions, and agents must choose partners based on the expected payoffs gained over a period of time. Although the work considers partner selection, it focuses on cooperative encounters and does not deal with the problems of negotiation. Another approach to opponent selection, using cognition-based strategies, is described in [10], in which several heuristic decision-functions facilitate the selection of optimal opponents. However, the work does not examine the effects of changing evaluations of resources and how this affects selection of opponents, nor does it deal with considerations of conflict, but instead, focuses on the efficacy of the decision heuristics.

Motivation has long been used in psychology [11] and ethology [12], where it is used to explain the *higher-level desires* of an organism. In computational settings, motivation is increasingly being used as a higher-level control mechanism that directs the goal generation, action-selection and decision-making activities of software agents (e.g. [13, 14]) and robots (e.g. [15, 16]). The importance of motivation as an enabler of autonomy in computational agents was perhaps first identified by d'Inverno and Luck [5], who discuss the importance of motivation in allowing an agent to generate its own goals, as opposed to adopting them from others. Further analysis of this view can be found in [17]. More recent efforts have extended Luck and d'Inverno's ideas to consider cooperation [13], planning [18] and norm-based multi-agent systems [19]. The use of motivation within negotiation is a relatively new approach. An early example is described in [20], where motivation is used to enable cooperative negotiations that aim to increase the utility of all participants. Perhaps the closest work to ours is that of Urbig et al. [4], which looks at the links and interdependencies of issue selection and partner choice, as well as their effect on behaviour during negotiation. However, their approach differs from ours by focusing on a formal specification of the possible interdependencies between the three aspects, rather than the development of an agent architecture and accompanying classification and selection mechanisms to enable the autonomy necessary to address these problems.

3 Motivated Agents and Negotiation Goals

Our previous work [21] has involved the development of a *motivated agent architecture*, which enables autonomous decision-making and action selection for computational agents. In the architecture, the external environment, goals, actions and resources are linked to the motivations of the agent through *motivational cues*, which are essentially beliefs which, when true, impact on the strength of the agent's motivations. The goals pursued, and the actions and resources used to satisfy goals, are all determined by the effects they have upon the agent's motivations via the cues to which they are linked. Thus, for example, a warehouse agent may notice that a box has been left lying in a corridor. If boxes in the wrong location act as cues for the agent, it will affect the intensity of the motivation linked to the cue. This may cause the agent to generate a goal to put the box back into its correct location in the warehouse. Whether this goal is adopted depends upon the size of the effect that satisfying this goal will have on motivational intensity levels. Once a goal is adopted, the agent must then select an appropriate plan, which may call for the use of some resources that enable the plan, and the agent must assess the motivational effects of executing the plans using any associated resources.

Figure 1 shows our motivated agent architecture. The agent forms a view of the environment, which is linked, along with the agent's goals, actions and resources, to cues which, if true, affect the agent's motivational intensity levels. If motivational intensities rise above certain threshold levels, goals are generated and the agent's decision-making module then considers the various actions and resources at its disposal, determines what the effect of their use would be on motivational intensity, and selects those that have the most beneficial effects. These are then passed on to the effectors of the agent to take action. In the rest of this

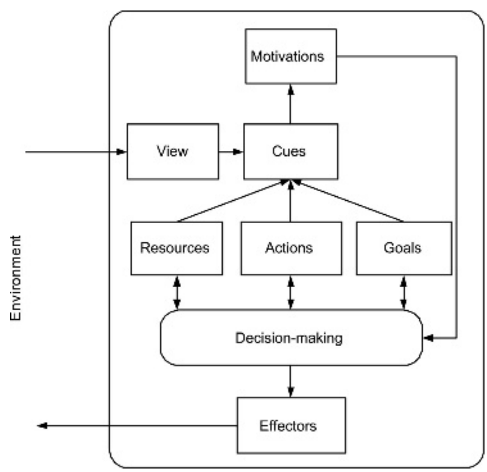

Fig. 1. The motivated agent architecture

Table 1. Possible values for three different negotiation goal attributes

Attribute	Attribute Values				
Price	10	20	30	40	
Delivery	Mon	Tues	Wed	Thurs	Fri
Quality	Low	Med	High		

paper we describe a classification mechanism and two selection mechanisms that lie within the decision-making component of the motivated agent architecture.

3.1 Negotiation Goals

Though the model above provides a general view of motivated agents, we need to refine it to describe how an agent can generate and reason about *negotiation goals*. In our work we have developed a model of negotiation goals that allows agents to autonomously decide what they want to negotiate about, what they do not want to negotiate about and what they do not care about. The components of a goal are its *attributes*, which represent the traditional AI notion of an *atom* composed of a predicate and a sequence of terms. So, for example, an attribute of a goal to place a box in a store room could be represented as: $In(boxA, storeRoom)$.

Our model is unique, however, in that we classify attributes according to their status in a forthcoming negotiation. Those attributes of a negotiation goal that describe what must be achieved are called *fixed attributes*, while those that *may* form the focus of the negotiation are called *potential attributes*. Potential attributes allow us to model negotiation issues such as *price*, *time* or *quality*, and they are composed of a *predicate* and a *set of values* to instantiate the predicate. So, for example, $Price(boxA, X)$ must have a range of values that can instantiate the variable, X. Table 1 shows the three potential attributes of *Price*, *Time* and *Quality* along with a set of values for each that can be used for instantiation.

Thus, a negotiation goal initially comprises a set of fixed attributes for describing what must be achieved, and a set of potential attributes that may or may not be negotiated over. Table 2 shows the initial structure of such a goal. The example goal here is the goal to place a box in a particular location. The identity of the box and the name of the location therefore represent fixed attributes of the goal. The potential attributes refer to the price the agent is willing to pay to have this goal satisfied, the time at which the goal is to be satisfied and the quality of the method used to move the box (imagine there are three possible ways to move the box, each with a different chance of damaging the contents of the box). Initially, the goal consists of just these types of attributes (i.e. fixed and potential). The task of the agent upon the generation of this goal is to decide which of the potential attributes are to be made *negotiable*, which should be made *fixed* and therefore, not negotiable, and which are of no importance (called *slack* attributes) and therefore, can be omitted from the negotiation altogether.

Table 2. The initial structure of a negotiation goal

Partial Goal Template		
Attributes	Fixed	Potential
Box id	✓	
Destination	✓	
Price		✓
Time		✓
Quality		✓

The decision about the status of a potential attribute depends on two factors. First, the preferences of the agent towards how the attribute will be instantiated must be considered. Constructing these preferences is achieved by assessing the effects of each instantiation on the agent's current activities. Second, the designer must supply a set of *classification rules* that are applied to the agent's preferences to determine whether the attribute is fixed, negotiable or slack. The form of the classification rules depends on the designer's needs for the domain in question, but we offer some example rules here:

1. An attribute is classified as *fixed* if the preferences of the agent contains *at most* one value that has positive motivational worth and all the rest as having negative motivational worth.
2. An attribute is classified as *negotiable* if the preferences of the agent contains *more than* one value that has positive motivational worth.
3. An attribute is classified as *slack* if *all* the values contained in the agent's preferences have the *same* motivational worth (both positive or negative).

The rationale for these rules is as follows. If only one value of an attribute has positive motivational worth, it is preferable for the agent to demand that this value be met in any forthcoming negotiations as, if any other value is used to instantiate the attribute, the agent will be in a worse state than previously. This means that the agent should include the attribute with the value as a fixed attribute of the negotiation goal. If, however, more than one value is identified that gives the agent positive motivational worth, then the agent can be afford to be flexible with regard to the attribute. Through this flexibility the agent increases its chances of reaching a successful settlement with an opponent, as there is greater room for agreement. Finally, if all the values for a given attribute have the same level of worth, there is no point in negotiating over the attribute, as the agent is indifferent to any instantiation. This allows the agent to prune irrelevant issues, thus making the negotiation more efficient and provides an incentive to opponents to enter into negotiation as, if they know that a preferred issue is irrelevant to the other agent, they can instantiate the attribute using their own preferred value.

In Figure 2, we show our *attribute preference construction and classification mechanism*. The potential attributes from a negotiation goal are passed to the *preference generator* that examines each possible value that can be used to instantiate

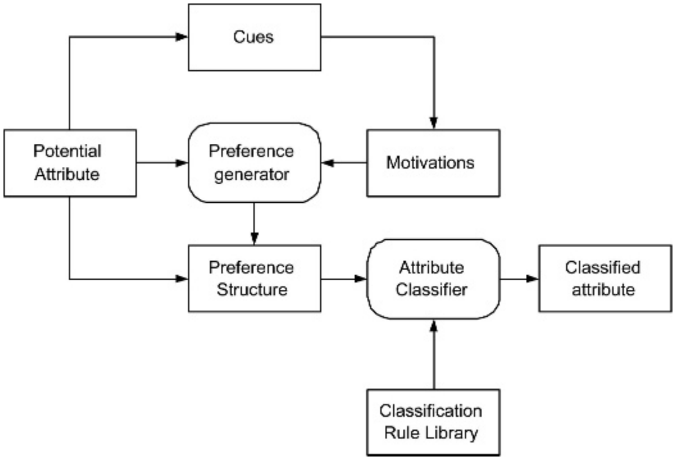

Fig. 2. The attribute preference construction and classification mechanism

Table 3. The final structure of a negotiation goal

Negotiation Goal Template				
Attributes	Fixed	Potential	Negotiable	Slack
Box id	✓			
Destination	✓			
Price			✓	
Time				✓
Quality	✓			

the attribute, and assesses their motivational worth by examining how the use of a value affects motivations via cues linking the attribute to motivations. By doing this for all the values of an attribute, a *preference attribute* is formed, which is simply a potential attribute with an associated preference ordering over its attribute-values. At this point, the preference attribute is passed to the *attribute classifier* that applies the attribute classification rules stored in the *attribute classification rule library.* After the application of the rules, each preference attribute of the negotiation goal is classified as either fixed, negotiable or slack. The final form of a negotiation goal for an agent wanting to engage in negotiation will be like that of Table 3. In the table the *price* attribute has been made negotiable, the *time* attribute has been made slack, and the *quality* attribute has been made fixed, along with the other two previously identified fixed attributes of *box id* and *destination.*

4 Constructing Preferences over Potential Attributes

Imagine a situation in which an agent is considering a goal regarding the movement of a box from one location to another. The fixed attributes of the goal state

Table 4. Three example plans

Plan name	Description	Quality
planA	Move the box with a trolley	High
planB	Move the box by carrying it manually	Medium
planC	Move the box by dragging it	Low

that a particular box, *boxA* is to be moved to a particular location, *roomB*, i.e $In(boxA, roomB)$. One of the potential attributes of the goal concern the *quality* of any plans used to satisfy the goal, where there exist three possible plans each with a different *quality* rating $\{high, med, low\}$ (imagine that each plan has a different chance of causing damage to the box, with the plan with the *lowest* chance has the *highest* quality rating). Three example plans are given in Table 4.

Now, in general, higher quality plans are preferred, hence we may have the preference ordering: $planA > planB > planC$ in which case the client agent chooses *planA*. However, there may be times when it is better to use *planB* or even *planC*. For example, using trolley required for *planA* might mean that another goal that needs the trolley is affected. However, this information may only be available at runtime, and so it is important for the agent to be able to re-order its preferences to meet the demands of the current situation. In the remainder of this section we show how this can be achieved.

4.1 A Motivated Preference Construction and Classification Mechanism

When considering the different ways in which a potential attribute can be instantiated, it is necessary to determine what effects it has on the agent's current activities. In order to do this we take the current activities of an agent, represented by the agent's *intentions*, each of which comprises a *goal* and a set of *plans* used to satisfy the goal. Plans represent the actions and *subgoals* that must be achieved before the goal can be satisfied and, in order to check whether any given instantiation of a potential attribute hinders or facilitates an intention, we must examine the plans associated with the intention to see what effects the instantiation has.

First we get the subgoals of an intention denoted by the term *subgoals*, which takes an intention and returns the set of subgoals encoded in the plans used to satisfy that intention. Then, we take a specific instantiation and determine those subgoals that are hindered *hindered* and those that are facilitated *facilitated*.

4.2 Scoring the Effect of an Instantiation

Having identified which intentions are hindered and which are facilitated, we next must provide a means to score each instantiation as a function of the *degree* of hinderance and facilitation. Then, we use this score to determine the preference ordering of the various different possible instantiations. The two functions, *hinderscore* and *facilitatescore*, defined below provide this functionality. The functions take an intention, *int*, and obtain its *worth* (derived from the strength of the motivation responsible for the goal's generation), which is then divided by the number of unhindered subgoals in the case of *hinderscore* or the number of facilitated subgoals in the case of *facilitatescore* (+1 is added to the denominator of *hinderedscore* to avoid division by zero).

$$hinderscore = -1 \times \frac{int.worth}{1 + (\#\{int.subgoals\} - \#\{int.hindered\})} \quad (1)$$

$$facilitatescore = \frac{int.worth}{\#\{int.facilitated\}} \quad (2)$$

We do this for all of the intentions of the agent and, once this is done, we combine both the hinder and facilitate scores into an overall score.

$$overallscore = \sum_{i=0}^{n} facilitatescore(int_i) + \sum_{i=0}^{j} hinderscore(int_j) \quad (3)$$

where i is the number of goals that are facilitated, and j is the number of goals that are hindered.

4.3 Potential Attribute Preference Orderings

Once we have calculated the overall score for each of the values associated with a potential attribute we can use this to order the values. This provides us with a *preference ordering* over the values for a potential attribute. Once a preference ordering has been established over the attribute-values associated with a potential attribute we call such an attribute a *preference attribute*. As an example, consider a goal to have a box moved from one room in a warehouse to another. The *time* when the box can be moved is a potential attribute, and the values that can instantiate the attribute are drawn from the set *Days*:

$$Days = \{Mon, Tues, Wed, Thur, Fri, Sat, Sun\}$$

If a preference ordering has been established over this set, then *time* is a preference attribute and its *pref* relation might look like : $\{Wed > Tues > Mon > (Thur, Fri, Sat, Sun)\}$, meaning that Wednesday is preferred over Tuesday, which is preferred over Monday, which is preferred over Thursday, Friday, Saturday and Sunday, all of which are equally preferred.

4.4 Resource Dependent Attributes

In the goal of Table 3, the attribute of price is negotiable. This is the usual situation for such an attribute, as the values that can instantiate the attribute constitute the use of a *resource* and, in general, a preference over the use of different amounts of resource can always be identified. For example, an agent buying a service off another agent always prefers cheaper prices over more expensive ones, and vice versa for an agent selling the service. Such attributes of a negotiation goal are, therefore, classed as *resource dependent*, and the problem here is not one of identifying the structure of the preference, but rather the limits on what values are acceptable.

One way to do this is to simply ensure that the amount that an agent pays [1] must not incur more costs than the benefit the agent gains from the satisfaction of its goal. However, in most negotiation frameworks, the utility to be gained from a goal is fixed, and the amount of acceptable cost is calculated from a position of zero cost. For example, imagine an agent with a goal of utility 10, where utility is measured in dollar units. In determining how much an agent can spend in order to satisfy the goal, it is easy to see that, in order to be efficient, the agent must not spend more that \$10. However, we might assume that the agent is only happy to spend this much as long as it has \$100 in reserve, but that if this reserve falls to \$50 then spending tha that amount might not seem so attractive, and the agent may lower its limit to, say, \$8 [2]. By doing so, the agent begins to focus on *optimising* the use of its monetary resources as the resource dwindles, by changing the value placed on the resource in response to changes in the quantity available. Such dynamic evaluation of resources is often overlooked in negotiation frameworks, and we argue that this limits their flexibility to deal with dynamic domains in which access to resources changes over time. Thus, the overall worth of a particular resource cannot be merely represented by its *objective worth*, but must also be supplemented by a *subjective worth* derived by the current *need* an agent has for it. To model this subjective need we use *resource-based motivations* that track the levels of resources available to an agent (via appropriate cues), increasing in intensity when such resources diminish, and decreasing in intensity when they are renewed. The intensity of such motivations is used to identify the *unit value* of a given resource. Only when the unit value of a resource is determined, is it possible to calculate the maximum amount, or the *reservation*, of the resource that can be used to settle an issue. The reservation is thus calculated to be that quantity of resource, the use of which has a cost equal to the benefit gained from the negotiation goal being satisfied.

In order to deal with resource-dependent attributes, we must make some changes to the mechanism shown in Figure 2. Figure 3 shows the amended mechanism, in which the dashed lines indicate the new connections, and there is a decision point where the type of attribute being dealt with is determined. If the attribute is resource dependent, then a *default preference structure* is used that simply describes a, monotonic preference between the different values associated

[1] We focus on the perspective of a buyer in the description of the evaluation of resources, though the analysis is similar (albeit reversed) for the seller's perspective

[2] This behaviour is simply that identified by economists when making the point that utility does not equate to monetary worth — in other words, the richer a person, the less utility he gains by increasing his wealth and conversely, as shown above, the less wealthy a person, the more utility is lost by a decrease in wealth.

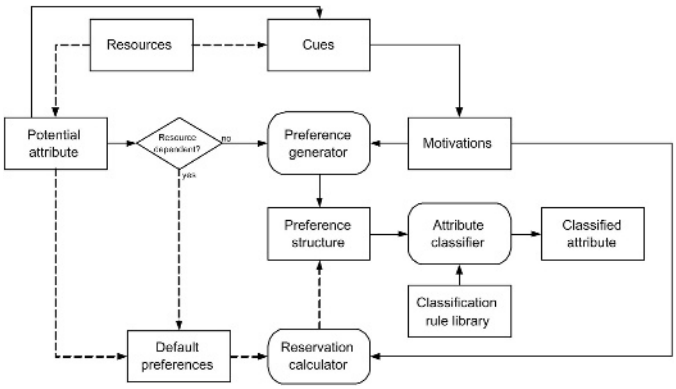

Fig. 3. The attribute preference construction and classification mechanism with resource-dependent attributes

with the attribute. Next, the associated resource is evaluated by the agent's *resource-based motivations* to determine its current unit worth. This is then passed, along with the default preference structure, to the *reservation calculator* that determines the maximum amount of resource that can be used, i.e. that unique quantity of resource whose total worth is equal to the worth of the goal. The amended preference structure is then passed to the attribute classifier as per other, non-resource-dependent attributes, and is classified according to the rules in the classification rule library.

5 Knowledge About Negotiation Opponents

In order for selection to work the agent must have some information about prospective opponents. In this section we describe such information.

5.1 An Opponent's Past Issue Choices for a Negotiation Goal

In previous negotiations, an opponent will have made choices over what potential attributes it wanted to negotiate over. Assuming that a record of these choices has been stored and that there is some regularity to the choices, it it possible to predict the opponent's choice of attributes for a future negotiation over the same goal. If the opponent's choices are variable, then the information about the choice of attributes might be based on a probability distribution over the attributes. To avoid unnecessary complications at this time, we assume that the record of an opponent's attribute choice for a given negotiation goal is of the form of a frequency distribution. Thus, for a given issue, we get the probability of it being chosen by

$$choiceProb = \frac{chosen(a)}{available(a)} \tag{4}$$

Table 5. Table showing the frequency of attribute selection for three negotiation opponents

Opponent	Attributes		
	A	B	C
1	0.5	0.7	0.2
2	0.4	0.4	0.9
3	0.3	0.8	0.4

where $chosen(a)$ returns the number of times the attribute, a, has been chosen by the opponent and $available(a)$ returns the amount of times the attribute a has been available as a potential choice. Thus, for a given negotiation goal with three attributes and three different negotiation opponents we may end up with the information shown in Table 5.

In the table, we can see three potential opponents in the first column labelled 1, 2, and 3. In the second, third and fourth columns, we have probabilities attached to each of the attributes (labelled A, B, C) for each opponent, calculated as the frequency that the attribute has been chosen in past negotiations for the current negotiation goal. A selecting agent then uses these frequency scores in combination with its own set of attribute choices to determine which opponents are likely to select issues that do not conflict with its own choices. We discuss how this can be achieved in Section 6.

5.2 Price Profiles of Negotiation Opponents

We must also provide our agent with information regarding opponent *price ranges* if our agents are to be able to assess the suitability of opponents with regard to current resource levels and their associated value. In negotiation, the participants normally announce *initial ask* prices (which can be considered the prices at which an opponent advertises its services) that then become adjusted through a process of concession-making in order to discover a price that all participants can accept. Initial ask prices tend to exaggerate the reservation price in order to increase the chance that a deal can be struck that is better for the agent than the reservation point.

Now, different agents may have different rules for setting initial ask prices given a reservation price but, for a given agent, the distance between the reservation and initial ask prices will tend to be stable. This means that any observed change in the initial ask price of an agent for a negotiation goal reflects a change in the subjective evaluation of the resource by the agent, where this change gives rise to a new (but private) reservation price and a corresponding publicly observed change in the initial ask price.

The *deal price* of a negotiation is the price at which agreement is found, and it depends upon a number of factors. First, it depends upon the initial ask prices of each participant of the negotiation, second, their reservation prices and third, the sequence of offers and counter offers on price made by each participant, i.e.

Table 6. Ask prices, deal prices and concessionary flexibility for an opponent

Neg Instance	Price Profile For Negotiation Opponent		
	Ask price	Deal price	Concessionary flexibility
1	10	9	0.1
2	9	7	0.2
3	11	8	0.3

their *concession strategy.* Though in general, agents will be able to select from a set of different concession strategies (each of which will affect the deal price), for simplicity we assume the agents only have one concession strategy, and we leave to later work the added complexity that different concession strategies bring to negotiation.

Given a fixed concession strategy for the agents, the distance between the initial ask price and any subsequent deal price will follow a predictable pattern that can be analysed by the other participants in the negotiation to predict what, given a particular initial ask price, the final deal price will be. The information about an opponent's initial ask price and deal price (what we call the *price profile*) must be available if an agent is to make judgements on the quality of an opponent on price.

Thus, over a number of negotiations it becomes possible to predict the deal price that will result in a negotiation for a particular goal with a particular negotiation partner by considering the amount of *concessionary flexibility* an opponent has shown during previous negotiations.

Table 6 shows the the price profile of a seller agent over three separate negotiations for the same negotiation goal. The first column indicates the negotiation instance and the second and third columns represent the initial ask price and the deal prices obtained. The fourth column shows the concessionary flexibility exhibited by the agent, which takes the difference of the initial ask price and the deal price and maps this difference to an interval of [0,1] where 0 means no flexibility and 1 represents maximum flexibility (i.e. the agent has accepted a price of zero. Not a likely situation!). The average concessionary flexibility can then be calculated as

$$aveCF = \frac{\Sigma_{i=0}^{n} cf_i}{n} \tag{5}$$

where cf_i is the i^{th} concessionary flexibility score and n is the total number of negotiations for which the information is available. So, for any new negotiation over the same goal with this opponent, even with different initial ask prices, we can use the average concessionary flexibility of the opponent to make a prediction about what the deal price will be. The expected deal price can then be compared against the agent's own reservation price to predict whether a negotiation with an opponent advertising an initial ask price is likely to reach a deal on price that is acceptable to the agent.

6 Negotiation Opponent Selection Mechanisms

The process of opponent selection now consists of an analysis of the possible issue selections of the opponent and their expected deal price. To make an optimal selection, the agent must minimise two measures. First, when considering issue selection, the agent should select the opponent that offers a *minimal* amount of *conflict*, and second, when considering price the agent should attempt to *minimise* the use of its *monetary resource*. We describe each of these in detail in the following subsections.

6.1 Selecting to Minimise Conflict

Negotiation issues are those negotiable attributes that have been identified by *more than one* of the participants. Those attributes that have been identified by only one participant are *uncontested*, as none of the other participants in the negotiation have classified the attribute as negotiable and, therefore, do not have a preference for its instantiation. A clear criterion on which selection can therefore be based is the minimisation of the amount of contested attributes in the forthcoming negotiation through the identification of those opponents who contest the *least* number of negotiable attributes.

In Figure 4, two agents have each identified which of a negotiation goal's attributes they want to negotiate over (shown by the two circles covering the attribute set of the negotiation goal). The intersection of these two subsets of attributes are those attributes identified by *both* participants and are therefore contested and thus potentially in conflict. These attributes thus represent the *issues* of the negotiation. When considering different opponents, the agent should

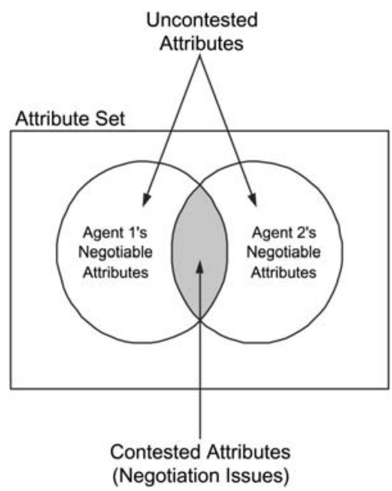

Fig. 4. Identifying Negotiation Issues

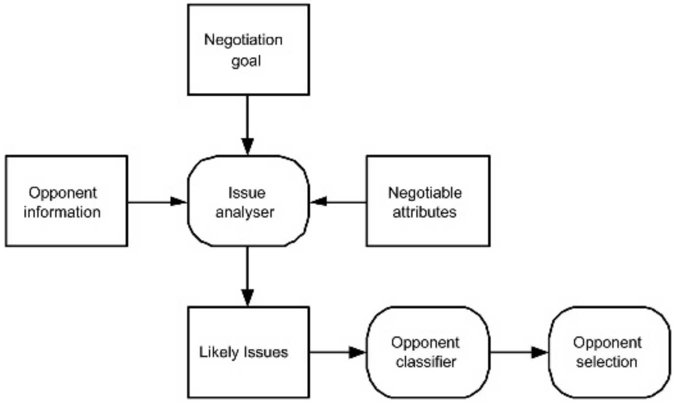

Fig. 5. The conflict minimisation selection mechanism

look for smaller intersections indicating less potential conflict and thus fewer issues to negotiate over. As the number of issues increases, the potential conflict increases and the negotiation becomes more difficult to resolve successfully.

Figure 5 shows the selection mechanism based on minimisation of conflict. The *negotiation goal* and the set of attributes identified as negotiable by the agent is passed, along with information about an opponent, to the *issue analyser*. The issue analyser examines the issues that the opponent is expected to select, and calculates the amount of conflict that can be expected in terms of the number of contested issues identified. This results in a *conflict score* which we calculate as

$$conflictScore = -1 \times \#\{negatt_{agent} \cap negatt_{opponent}\} \qquad (6)$$

where $negatt_{agent}$ are the attributes identified as negotiable by the selecting agent and $negatt_{opponent}$ are the expected negotiable attributes of the opponent. The opponent is then rated according to its conflict score and is passed to the *opponent selection* component.

6.2 Price-Based Selection

When considering resources and their use in negotiation, an agent must try to select those opponents that are likely to accept a price that falls within the agent's own current reservation price. Given information about an opponent's concessionary flexibility, the agent takes the opponent's current ask price and estimates what the deal price will be, and then compares this to its own reservation price. If the expected deal price for an opponent exceeds the reservation price, then the opponent is omitted from further analysis, otherwise it is passed to the *opponent selection module*. Figure 6 shows the selection mechanism based on price considerations. Information about the current resource being used (here a monetary resource), the reservation of the resource and information on the current opponent are all passed to the *resource manager*.

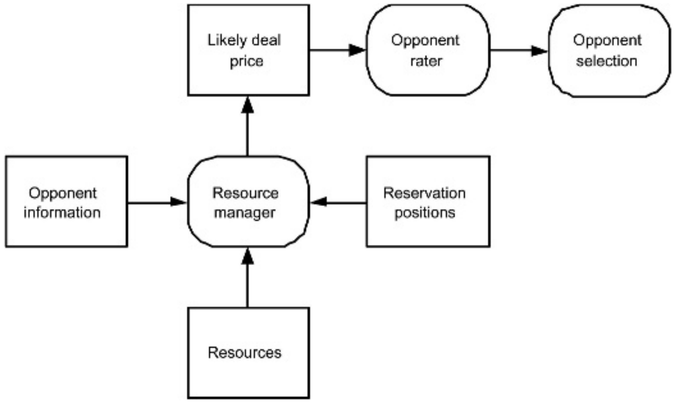

Fig. 6. The price-based selection mechanism

The resource manager calculates the expected deal price, given the initial ask price of the opponent and its concessionary flexibility score, and examines if this is under the reservation price. If this is so, the resource manager sends the information about the opponent to the *opponent rater* which ranks the opponent in relation to the other opponents under consideration. Once all opponents have

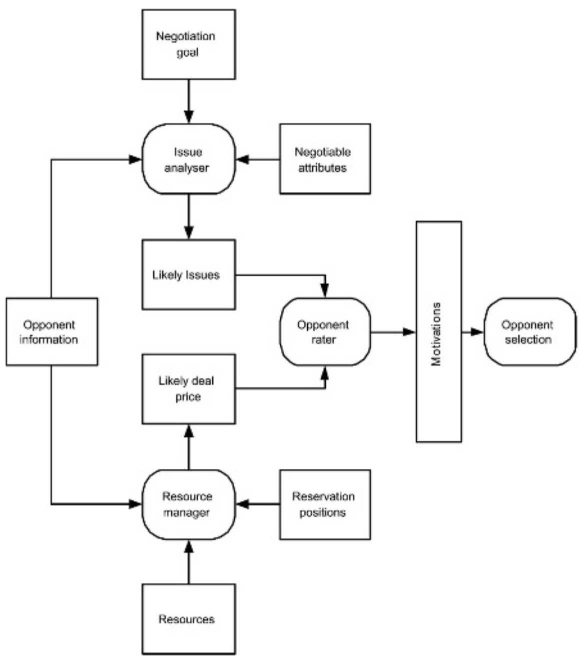

Fig. 7. The opponent selection mechanism

been so ranked, they are sent to the opponent selection module that performs the selection.

6.3 Combining Conflict Minimisation and Price-Based Selection Mechanisms

We now combine both selection processes into one mechanism; the *opponent selection mechanism* shown in Figure 7. In addition to the components in the conflict and resource selection mechanisms, the opponent selection mechanism shows a *motivational component* that can be used to weigh one selection process over another. This allows an agent to change the focus of selection from one criterion to another as required. Thus, for example, at times of extremely low resource levels, the agent can prioritise those opponents who offer extremely cheap services while paying less attention to the number of issues that might need negotiating over. At other times, for example when a negotiation goal has a large amount of value, the agent can prioritise the minimisation of conflict so as to increase the chance of engaging in a successful negotiation while de-emphasising the importance of cost. The mechanism is simply the two selection mechanisms discussed above combined into one, and thus we omit further discussion of the individual components.

7 Evaluating the Model

The work presented in this paper is still in development, but we have performed a limited evaluation to examine if a buyer agent using only the *price-based selection mechanism* is able to select the best negotiation opponent in terms of optimising monetary resources. We tested this in the following way. First we cal-

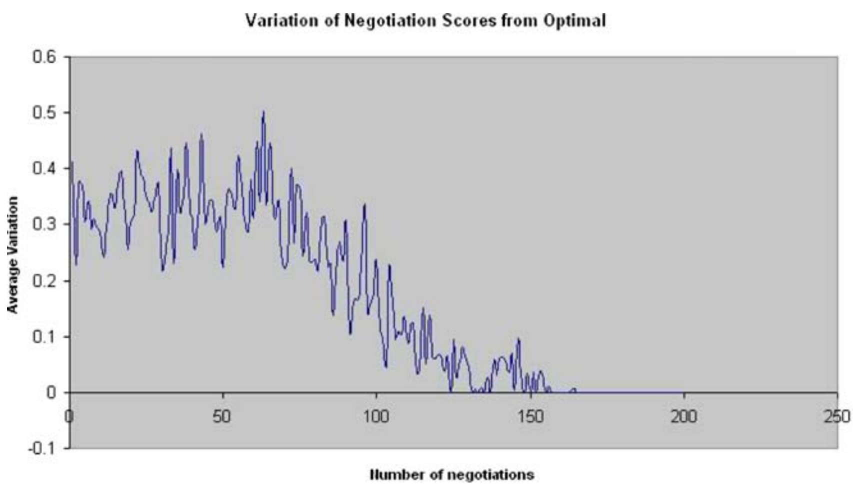

Fig. 8. Performance of price-based opponent selection

culated which opponents, out of a set of available opponents, offered the optimal deal price for the agent given its current resource evaluation, then we allowed the agent itself to choose an opponent and we compared this choice with the previously identified optimal opponents to obtain a measure of variation of the agent's selection choice from optimal. Each run of the experiment lasted for 200 negotiation rounds. In each run there was 1 buyer-agent, a pool of 100 potential opponents, and we performed 20 such runs. The buyer agent had a *conserve money* motivation that determined the current reservation value on the agent's monetary resource. The graph in Figure 8 shows the average variation between the buyer-agent's obtained deal prices resulting from its opponent selection and the optimal deal prices over the 20 experimental runs. The buyer-agent eventually learns to find the best opponents given its current motivational state, shown by the variation line falling to 0 after negotiation number 150. Future work will involve more extensive experimentation on various aspects of the model to ascertain its benefits and limitations.

8 Conclusion

By linking negotiation issues to motivations, agents are able to evaluate prospective negotiation opponents in terms of a) the likelihood that a conflict over issues will exist and b) the expected performance of the opponents on resource-based issues. Such evaluation is important for systems that need agents to act in an autonomous manner. As an agent's circumstances change, the need to negotiate over issues may also change, and this must be considered when making selections over which opponent to negotiate with. In this paper, we have described a motivated agent architecture comprising a classification mechanism and two selection mechanisms that enable an agent to successfully identify those opponents with whom the chances of conducting a negotiation with both minimal conflict and suitable outcomes is possible. In terms of future work, an important factor when investigating negotiation is to consider the impact that protocols and strategies may have on the outcome. Though our work currently ignores these considerations, our aim is to show how an autonomous approach to opponent selection, might be constructed. We expect future work to address the problems involved in opponent selection in more demanding negotiation scenarios such as those that include multiple strategies and protocols.

References

1. David, E., Azoulay-Schwartz, R., Kraus, S.: Protocols and strategies for automated multi-attribute auctions. In: 1st International Conference on Autonomous Agents and MultiAgent Systems (AAMAS 2002), Bologna, Italy, ACM Press (2002) 77–85
2. Sierra, C., Jennings, N.R., Noriega, P., Parsons, S.: A framework for argumentation-based negotiation. In Singh, M.P., Rao, A.S., Wooldridge, M., eds.: Intelligent Agents IV. Volume 1365 of LNCS. Springer-Verlag (1998) 177–192
3. Faratin, P., Sierra, C., Jennings, N.R.: Negotiation decision functions for autonomous agents. Journal of Robotics and Autonomous Systems **24** (1998) 159–182

4. Urbig, D., Schroter, K.: C-IPS approach to negotiating agents: Specifying dynamic interdependencies between issue, partner, and step. In: 3rd International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS 2004). (2004)
5. d'Inverno, M., Luck, M.: Understanding Agent Systems. Springer-Verlag (2001)
6. Sabater, J., Sierra, C.: Social regret, a reputation model based on social relations. SIGecom Exch. **3** (2002) 44–56
7. Ramchurn, S.D., Sierra, C., Godo, L., Jennings, N.R.: Devising a trust model for multi-agent interactions using confidence and reputation. Applied Artificial Intelligence **18** (2004) 833–852
8. Tesfatsion, L.: A trade network game with endogenous partner selection. In: Computational Approaches to Economic Problems. Kluwer (1997) 249–269
9. Banerjee, B., Sen, S.: Selecting partners. In: 4th International Conference on Autonomous Agents (AGENTS 2000), Barcelona, Spain, ACM Press (2000) 261–262
10. Dutta, P.S., Moreau, L., Jennings, N.R.: Finding interaction partners using cognition-based decision strategies. In: IJCAI 2003 Workshop on Cognitive Modeling of Agents and Multi-Agent Interactions. (2003) 46–55
11. Maslow, A.: The farther reaches of human nature. Penguin Books, New York (1971)
12. Tinbergen, N.: The study of instinct. Oxford University Press, Oxford (1989)
13. Griffiths, N.: Motivated Cooperation. PhD thesis, University of Warwick (2000)
14. Coddington, A., Luck, M.: Towards motivation-based plan evaluation. In Russell, I., Haller, S., eds.: 16th International FLAIRS Conference, St. Augustine, FL, USA (2003) 298–302
15. Morignot, P., Hayes-Roth, B.: Adaptable motivational profiles for autonomous agents. Technical report, Knowledge Systems Laboratory, Stanford University (1995)
16. Balkenius, C.: The roots of motivation. In J.A. Mayer, H.L. Roitblat, S.W. Wilson, editors, *From Animals to Animats 2* (Cambridge, MA: MIT Press, 1993)
17. Luck, M., Munroe, S., d'Inverno, M.: Autonomy: Variable and generative. In Hexmoor, H., Castelfranchi, C., Falcone, R., eds.: Agent Autonomy, Kluwer (2003) 9–22
18. Coddington, A.: Self-motivated Planning in Autonomous Agents. PhD thesis, University of London, London (2001)
19. López y López, F., Luck, M., d'Inverno, M.: Constraining autonomy through norms. In: 1st International Conference on Autonomous Agents and MultiAgent Systems (AAMAS 2002), Bologna, Italy, ACM Press (2002) 674–681
20. Zhang, X., Lesser, V., Wagner, T.: A proposed approach to sophisticated negotiation. In AAAI Fall Symposium on Negotiation Methods for Autonomous Cooperative Systems (2001)
21. Munroe, S., Luck, M., d'Inverno, M.: Towards motivation-based decisions for worth goals. In Marik, V., Mueller, J., Pechoucek, M., eds.: 3rd International Central and Eastern European Conference on Multi-Agent Systems (CEEMAS 2003). (2003) 17–28