

Balancing Conflict and Cost in the Selection of Negotiation Opponents

Steve Munroe, Michael Luck

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

Abstract. Within the context of agent-to-agent purchase negotiations, a problem that has received little attention is that of identifying negotiation opponents in situations where the consequences of conflict and the ability to access resources dynamically vary. Such dynamism poses a number of problems that make it difficult to automate the identification of appropriate opponents. To that end, this paper describes a motivation-based opponent selection mechanism used by a buyer-agent to evaluate and select between an already identified set of seller-agents. Sellers are evaluated in terms of the amount of conflict they are expected to bring to a negotiation and the expected amount of cost a negotiation with them will entail. The mechanism allows trade-offs to be made between conflict and cost minimisation, and experimental results show the effectiveness of the approach.

1 Introduction

Negotiation is a particularly important form of interaction between agents, as it allows conflicts to be resolved in situations of competing interests. There are many existing frameworks for negotiation (e.g., [1, 2]), most of which 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. However, though the steps that are taken within a negotiation are typically the responsibility of an agent to determine, decisions made prior to negotiation are usually taken by the agent's user.

This may be acceptable in certain limited situations such as when an agent negotiates on behalf of a human user, but in persistent, multi-agent systems where agents perform their tasks away from human direction, it is often not possible. Agents in such systems can be faced with goals of differing importance and have varied access to resources, and when attempting to satisfy such goals through negotiation, it is important that the risks that the negotiation might fail be matched to the importance of the goal, and that the costs incurred are similarly matched to current resource levels. To be able to do this, agents must reason about the level of conflict that different negotiation opponents offer, and trade this off against concerns over cost, taking into account current resource availability. Such reasoning and decisions made by an agent hoping to identify suitable negotiation opponents must take place before negotiations begin, and the problems inherent in doing this successfully are important to address if negotiation is to be effectively employed in dynamic domains. In fact, such pre-negotiation factors represent a growing area of research (e.g.[3]).

1.1 Opponent Selection for Negotiation

Negotiation can be a time-consuming activity, and failed negotiations represent a waste of time and resources. A common source of failure can be the existence of conflicting interests between the negotiation participants. Thus, for example, a buyer may require that a service be delivered by a given date that the seller is not willing or able to meet. Such conflicts are therefore an important factor to minimise. Often, an agent requiring a service (the buyer) can be faced with a number of different service providers (the sellers) and the choice about which to negotiate with should be made relative to the probability of conflict and the consequences of any subsequent failure. Now, agents gain value when their goals are satisfied and, if a negotiation fails, the value to be gained can be lost, especially when there is no chance to re-negotiate. The seriousness of this depends upon the importance of the goal, and in general, an agent should be less willing to risk conflict in negotiations over important goals than less important goals. One way to decrease the risk of conflict is to attempt to identify sellers that exhibit more cooperative attitudes towards negotiation. It must be remembered however, that sellers will, in general, try to maintain their own levels of gain from a negotiation, and so it will often be the case that an increase in cooperation on certain factors may be accompanied by a decrease in cooperation in others. For example, by cooperating on a particular issue such as the *time of delivery*, a seller may consequently increase its *price*. Thus often, in order to obtain more cooperation, a buyer may need to pay more as a consequence.

In dynamic domains, access to resources change over time, and an agent with low levels of available resource must prioritise its management and use. Thus, the challenge is to define effective decision mechanisms for an agent in order to balance the need for negotiations exhibiting less conflict with the varying need to manage resources. The ability to make such decisions in the face of dynamically changing environments requires a degree of autonomy from the buyer-agent; in our work we adopt an approach to enabling autonomy based upon the construct of *motivation*, which allows an agent to dynamically evaluate the utility or *worth* of its activities and resources. Such an approach focuses on cognitively plausible ways to enable agents to form valuations on goals and activities, and tries to integrate such methods into the agent architecture. Thus, although motivation performs a similar role to the more traditional utility function approach, it addresses the problem in a different way.

Such dynamism poses considerable problems for both market-based approaches, in which the *public* evaluation of a product or service can be uncertain and time-dependent [4], or in bargaining situations, in which the *private* evaluation of a good or service can depend upon an agent's own concerns and the bargaining context. In this paper we focus on the latter problem, in which an agent's private evaluation of its goals and resources changes over time.

Much work is currently being undertaken that examines opponent selection from the point of view of *trust* and/or *service reputation* (e.g. [5, 6]), 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 agent systems, especially those characterised by openness they are, we argue, only part of the story and must be augmented with the kinds of issues we are investigating in this paper. Specifically the ability to reason

about potential conflict within a negotiation and the impact of resource constraints on an agent's ability to circumvent such conflict.

The paper proceeds as follows. In Section 2 we provide some background notions that help us develop our opponent selection model. In Section 3 we discuss in detail the selection mechanisms. Section 4 presents some initial empirical results on the mechanism, and Section 5 discusses some related work. Finally, Section 6 offers some concluding remarks and describes future work.

2 Background Notions for the Selection Mechanism

In this section we briefly discuss some background notions that help to build up the opponent selection mechanisms.

2.1 Negotiation Goals

In earlier work [7] 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 that describes the placing of a box in a store room could be: $In(Box(a), location(StoreRoom))$, where In is the predicate and $Box(a)$ and $location(StoreRoom)$ are the terms.

In our model, goal attributes are classified according to their *status*. Those attributes of a goal whose variables have been instantiated are called *fixed attributes*, while those that do not yet have their variables instantiated are called *potential attributes*. Agents examine the potential attributes of their goals to decide how to instantiate their variables. If the agent can identify more than one way to instantiate a variable that has some value, then the attribute is called a *negotiable attribute*. If non of the ways to instantiate a variable are preferred, or, all possible instantiations are equally preferred, then the attribute is called a *slack attribute*, which means that the agent is indifferent to any instantiation. Negotiable attributes are those attributes that can potentially form the focus of a negotiation, in which case they become *negotiation issues*.

This model of negotiation goals allows us to model the potential amount of conflict in any given negotiation. Negotiations for which there are many issues stand a greater risk of failing due to the increased chance of the existence of incompatible interests between the participants on one or more of the issues. If a buyer-agent can avoid those sellers whose choice of negotiation issues coincides with its own, then the chances of encountering an unreconcilable conflict of interest are decreased and the chances that the negotiation will be successfully concluded are increased.

The other main factor that impacts upon the selection of negotiation opponents is price. It is most generally the case that agents have access to only limited amounts of resources, and as a consequence, the use of those resources must be managed efficiently if the agent is to be effective in its role. In some agent systems the optimal use of resource is a solitary concern that governs all decisions, however, other systems contain agents that must juggle a number of different priorities, and thus, the urge to optimise

a resource may sometimes come second place to other priorities. In the real world, the scarcity of a resource tends to increase its value, for example the price of oil can increase dramatically if production is interrupted by war or natural catastrophe. This can lead to rationing and attempts to optimise its use for only the most important activities. Similarly in agent systems, where resource management is only one of a number of concerns, as resource levels change, agents should begin to change their attitudes towards how they are used. Thus, when resources are low the constraints imposed upon their use should be emphasised and, in situation when resource levels are high, these constraints should be de-emphasised so that other priorities are able to take precedence.

2.2 Overview of the Opponent Selection Mechanism

In Figure 1 we show our opponent selection mechanism, which enables a buyer-agent to successfully identify a seller-agent from a set of seller-agents that best meets its current needs for conflict minimisation and/or price minimisation. The remainder of this paper describes the individual components in detail, but here we provide a high-level overview that provides a reference for the later detail.

When considering the minimisation of conflict, the buyer takes the *negotiation goal* and the set of attributes it identified as negotiable and passes them, along with some information about a seller (described in Section 3) to the *issue analyser*. The issue analyser calculates the expected level of conflict that will entail if a negotiation with this opponent is attempted.

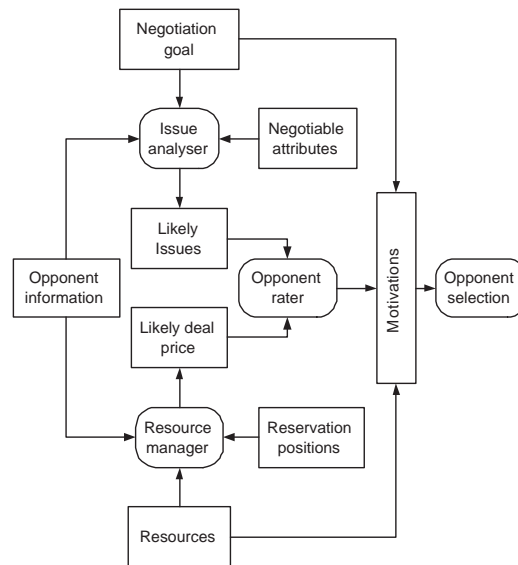


Fig. 1. The opponent selection mechanism

When considering price, the process through the mechanism is as follows. Information about the current resource being used (here a monetary resource), the constraints or *reservation* placed on the use of the resource and information about the seller under consideration are all passed to the *resource manager*. The resource manager calculates the *expected cost* of the negotiation with this seller, and examines if this is under the reservation placed on the use of its monetary resource. If this is so, the resource manager sends the information about the seller to the *opponent rater*, which scores the seller's expected cost based on how far it is under the reservation.

Also included in the mechanism is a *motivational component* used to influence the strength of the buyer's concerns to minimise either conflict minimisation or price minimisation. In other work [8] we describe a model of motivation for autonomous agents in which motivation takes on the role of a number of *high-level desires* of an agent. Motivations control the overall shape of an agent's behaviour by influencing its decision-making and activity selection so that these higher-level desires are fulfilled. In the motivation model, an agent's goals and resources are linked to its motivations via special attributes, called *motivational cues* that describe features of objects in the agent's environment or mental state. For example, a resource may have an attribute describing its quantity that acts as a motivational cue, or a goal may have an attribute describing its value or worth that also acts as a motivational cue. When an agent considers a motivational cue, either because it senses it in the environment or because it is part of an object it is currently reasoning about, this can cause the motivations it is linked to to change in strength or intensity. As motivations become more intense, they begin to exert more influence on the agent's decision-making. So, for our buyer-agent, if a goal has high worth, this causes the buyer's *achievement motivation* to increase in intensity, thus leading the buyer to try to increase the chances of a successful negotiation by decreasing the amount of conflict. Similarly, as resources dwindle, the buyer's motivation for *conserving scarce resources* increase in intensity thus leading opponent selection towards those sellers who offer cheaper deals.

2.3 Basic Notions

In our model for opponent selection, *attributes* form the basis of all that follows. Attributes represent any perceivable features of either objects in the environment or in an agent's mental state. If \mathcal{A} is the set of all attributes then any object, o in the environment is a subset of the environment, $o \subseteq \mathcal{A}$. Goals and motivations are also objects but only in an agent's inner mental environment. The set of all goals \mathcal{G} is therefore also a subset of \mathcal{A} as is the set of all motivations $\mathcal{M} \subseteq \mathcal{A}$. Goals have value or *worth* for an agent where this worth lies within the interval $[0,1]$. To obtain the worth of a goal we use $gw : \mathcal{G} \rightarrow [0, 1]$.

The set of all negotiation goals, NG , is that subset of \mathcal{G} for which negotiable attributes have been identified $NG \subseteq \mathcal{G}$. A negotiable attribute $na \in \mathcal{A}$ is an attribute that can be instantiated with more than one value. Let \mathcal{V} be the set of all values and for any $na \in \mathcal{A}$, let $attValues : \mathcal{A} \rightarrow \mathcal{V}$ return the set of values, $V \subseteq \mathcal{V}$, that can be used to instantiate na . For example, a negotiation goal may contain the negotiable attribute *time of delivery*, which can be assigned any one of the following set of values $[Monday, Tuesday, Wednesday]$.

Let \mathcal{R} be the set of all resources, $R \subseteq \mathcal{R}$ be a subset of resources, and r be a single resource. *Resources* are finite and represented as countable unit quantities. The quantity of a resource is given by the function: $quantity : \mathcal{R} \rightarrow \mathbb{N}$. For our purposes in this paper, we consider only one such resource, money, which has unit quantities of \$1.

The buyer-agent, b , and the set of seller-agents S form the agent society. The buyer is defined by the tuple $b = \langle A_b \times G_b \times M_b \times R_b \rangle$, where $A_b \subseteq \mathcal{A}$ is the attributes that describe the buyer, $G_b \subseteq \mathcal{G}$ is the buyer's goals, $M_b \subseteq \mathcal{M}$ is the buyer's motivations and $R_b \subseteq \mathcal{R}$ is the buyer's resources. At any time, a subset of the buyer's goals may be negotiable, $NG_b \subseteq G_b$. The buyer generates a negotiation goal using $gengoal : \mathcal{NG} \rightarrow ng$. Once a negotiation goal has been generated, both b and each $s \in S$ must identify which attributes of the negotiation goal they want to negotiate over. To determine this the buyer and each seller uses $getngs : NG \rightarrow A$, where A represents those subset of attributes of an ng that the agent wants to negotiate about.

Determining the Buyer's Reservation Price When attempting to satisfy an issue such as price in a negotiation, b must ensure that the amount it pays does not incur a cost greater than the benefit it gains in having the goal satisfied. Thus, in order to proceed, b must determine how much it can afford to pay to have the goal satisfied through negotiation, where this amount is called the *reservation price*.

For b , the reservation price is determined by two factors: the worth to b of the negotiation goal, and b 's current evaluation of its monetary resources. The worth of the goal is used to determine the maximum amount of negative worth or cost that b can bear in paying for the satisfaction of the goal, and the current worth evaluation of the resource must be used to discover how much resource (at the current evaluation) must be paid out to equal the worth of the goal.

The worth of the goal is determined *a priori* by the designer using the function gw defined above, but the evaluation of the monetary resource dynamically changes with the amount available as discussed in Section 2.2, and so we enable b to determine this at runtime.

We achieve this by assigning to b 's monetary resource a base unit-worth, $bw : \mathcal{R} \rightarrow [0, 1]$, which can then be modified by the function, $worthmodifier : \mathbb{N} \rightarrow RAT$, which takes into account how much of the resource b has and changes the worth placed on the resource as a consequence. For example, a base worth is assigned to \$1 by the designer, say 0.5, then during runtime we modify this number by finding out how much money b has and applying the $worthmodifier$ function to increase or decrease the base worth. Figure 2 shows an example $worthmodifier$ function. If the b has \$100 (represented by the origin of the graph), then the base unit-worth of \$1 is unchanged, i.e. 0.5. However, if cash levels fall to \$50, then the base unit-worth of \$1 is increased by 0.05 to 0.55. Conversely, if cash increases to \$150 then the worth is lowered by decreasing the base unit-worth by 0.05 to 0.45. Clearly, many such functions could be defined to describe the relationship between resource levels and worth, and this is just one example.

Once the unit-worth of the monetary resource has been updated to take into account current resource availability, b 's reservation price, $bres$, is calculated in the following manner, $bres = \frac{gw(g)}{(bw(money) + worthmodifier(quantity(r)))}$, which returns the quantity of money that, if paid out, has negative worth equal to the positive worth of the goal.

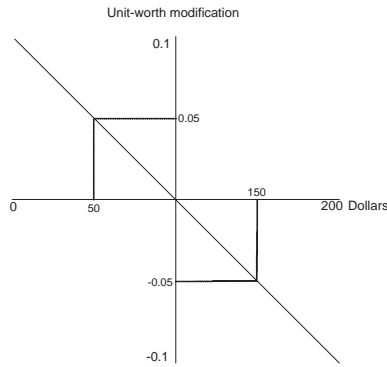


Fig. 2. An example unit-worth modifier function

3 Opponent Selection Mechanisms

Having described the basic underlying notions of our model and how the buyer determines its reservation price for a forthcoming negotiation, we now go on to describe how the selection of opponents is achieved. Recall that the buyer's selection of a negotiation opponent depends on two things: first, sellers who offer lower prices are, all else being equal, preferred over those offering higher prices; second, sellers whose selection of negotiation goal attributes do not conflict with the buyer's selection are, all else being equal, preferred over those whose selections do conflict, since the latter may lead to longer negotiations, higher prices and a greater chance of failure. Both of these factors must be considered in a complete opponent selection mechanism, which we describe below.

3.1 Seller's Past Issue Choices

As discussed in Section 2.1, negotiation issues are those attributes of the negotiation goal that have been selected for negotiation by *both* the buyer and the seller. By contrast, attributes that are identified as important by only one participant are *uncontested*, and can be instantiated at the value most preferred by the agent that selected it for negotiation. When the buyer agent seeks to negotiate for a goal with high worth it should, therefore, attempt to minimise the number of issues in the negotiation by identifying the sellers who are likely to offer negotiations with the least amount of issues. In order to be able to making good opponent-selection decisions, a buyer requires historical information about sellers. (In what follows, we use the information from prior negotiations with the seller, but this might equally be obtained through third-party assessments if there is no direct negotiation history.) To be useful, this historical information must provide price profiles for individual agents (which we consider later) and the frequency with which a seller has previously selected particular issues for the current negotiation goal. This is given by examining the frequency that an attribute has been selected in past negotiations, given by $selfreq = \frac{chosen(a)}{available(a)}$, where $chosen(a)$ returns the number of times

attribute a , has been chosen, and $available(a)$ returns the number of times the attribute a has been available. Thus, for a negotiation goal with three attributes and for which there are three different sellers, this historical information available to the buyer can be summarised as in Table 1. In the the first column are three potential sellers labelled 1, 2, and 3. The three attribute columns contain the selection frequencies for attributes A, B and C, for each seller, calculated as specified above.

Table 1. Frequency of attribute selection

Sellers	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

3.2 Seller Price Profiles

We also provide the buyer with information regarding seller *price ranges*, called seller *price profiles*, used to assess the suitability of sellers on price. Sellers announce an *initial ask price* (the price at which they advertise their service), which in general, will be a price that is higher then the actual *deal price* obtained in the negotiation. In past negotiations between the buyer and a seller, the difference between the seller's initial ask price and the resultant deal price will have been revealed. Assuming some regularity in the negotiation performance of both the buyer and seller, analysing this difference over a number of different negotiations can help the buyer to predict the deal prices of future negotiations over the same goal once the seller's ask price is known. This is done by calculating the amount of *movement* from ask to deal price exhibited by a seller in a negotiation, where we call this movement the *concessionary flexibility* of a seller agent.

In order to calculate the concessionary flexibility we assume that the full range of prices that a seller can demand for the satisfaction of a goal through negotiation is known to the buyer and falls within the interval $[0,u]$, where u is some upper limit. The concessionary flexibility, cf , exhibited by a seller, s , for the attribute of *price* in a negotiation for a goal, ng , is then given as $cf(s, price, ng) = \frac{ap-dp}{u}$, where ap is the initial ask price of the seller and dp is the deal price obtained.

Table 2 shows the information that a buyer might possess about the price profile of a seller 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 seller calculated as described above.

Over a number of different negotiations, the average concessionary flexibility, $avecf$, exhibited by a seller, s , is given as $avecf(s, price, ng) = \frac{\sum_{i=1}^n cf_i}{n}$, where cf_i is the concessionary flexibility score obtained from the i^{th} negotiation out of a total of n negotiations about ng with s . Given a concessionary flexibility score and an initial ask-price for a

Table 2. Seller price profiles

Neg Instance	A Seller's Price Profile		
	AP	DP	CF
1	10	9	0.1
2	9	7	0.22
3	11	6	0.45

seller-agent, it is thus possible to estimate the deal price for a new negotiation. This is done by taking the initial ask-price announced by the seller for a new instance of negotiation and subtracting the amount given by the average concessionary flexibility score, $expdp = ap - (avecf \times ap)$. Though this method cannot guarantee a completely accurate estimation of the deal price it can help the buyer to make a rough estimate of the deal price of a forthcoming negotiation.

3.3 Rating Sellers on Conflict

Given the above information about the likely issue selection and deal price of different sellers, it is now possible to provide a rating of the different sellers in terms of both expected conflict and expected cost. First, we can determine the likely number of issues in the negotiation that will result from a given opponent selection. Second, we can also make a prediction about the likely deal price obtainable in a negotiation from a given seller. We rate a seller on conflict in the following manner. First we attempt to estimate the number of issues that will result in a negotiation with a given seller. This is done by taking the buyer's own negotiable attribute selection, $natts_b$, and making a prediction of what attributes the seller will select. However, since we have only probabilistic information regarding a seller's attribute choice, we must establish for the buyer a selection risk threshold, srt , that fixes a probability beyond which the buyer believes an issue will be selected by a seller. Then we form the issue set, $issues$, that contains only those attributes chosen by the buyer for which the probability, p , that they will *also* be selected by the seller is greater than srt . So, given $natts_b$, we take an attribute a belonging to the negotiation goal, and if $a \in natts_b$, has probability, $p > srt$ of being selected by the seller-agent, then $issues = issues \cup \{a\}$. Second, we give the seller a conflict score, cs , by taking the difference between the cardinality of $natts_b$ from the cardinality of $issues$, i.e. $cs = \#natts_b - \#issues$. Note that the higher the conflict score the *less* issues there are and thus the better the chance of reaching a deal.

3.4 Rating Sellers on Price

The costs that a buyer can bear depends on the worth of the goal and the current worth of the resource used as payment. When resources are low, the buyer must focus on selecting sellers that offer cheaper deals. Rating sellers on price is simply a matter of taking a seller's expected deal price, checking to see if it is under the buyer's current reservation price and, if so, giving it a price score, ps , using the following equation, $ps = bres - expdp_s$, where $bres$ is the reservation price of the buyer and $expdp_s$ is the deal price expected from the seller.

3.5 Combining Conflict Minimisation and Cost Minimisation

The *opponent selection mechanism* combines both of the above approaches and allows the buyer to prioritise the minimisation of conflict or the minimisation of cost, or, to make an opponent selection that is a compromise of the two. This works in the following way. As discussed in Section 2.1, a buyer's goals are linked to its achievement motivation, and goals with high worth affect this motivation by increasing its intensity. Similarly, the buyer's resources are linked to the conserve scarce resources motivation, and as resource levels decrease this motivation also increases in intensity. In order to allow the agent to make opponent selections that observe the state of its motivations, we take these intensities and use them to weight the scores used to rate the different sellers.

Thus, if a goal is of high worth we can influence the selection decision to favour those sellers that offer less conflicting negotiations by calculating the weighted conflict score (wcs) as $wcs = gw(ng) \cdot cs$, where $gw(ng)$ gives the worth of the negotiation goal, ng , and cs is the conflict score. Similarly, if resources are low then we can use the intensity of the *conserve scarce resources motivation* to ensure that opponent selection favours sellers offering cheaper deals by calculating the weighted price score, (wps), as $wps = worth(r) \cdot ps$, where $worth(r)$ gives the current intensity of the conserve scarce resources motivation and ps is the price score. Both scores are then normalised and combined to produce an overall *rating* for a seller-agent: $rating(s) = wcs + wps$. Once all candidate opponents are rated, the agent selects the seller with the highest rating to be the negotiation opponent.

The opponent selection algorithm is shown in Algorithm 1, which takes the buyer agent's motivations, a new goal and the existing goals of the buyer as arguments, and returns the selected opponent. First, the issues for the new goal are determined (line 1), and then the set of candidate negotiation partners are identified after the goal is announced (line 2). The buyer then obtains the profile for each candidate (line 4), and determines the expected issue choices of each for the goal (line 5). Next, the conflict score is calculated (line 6), as the number of issues that both the buyer and the candidate select, and the expected deal price of each candidate is determined (line 8) and scored (line 9). Once both scores have been calculated, they are weighted using the value representing the worth of the goal and the strength of the *conserve resource* motivation, and then combined to produce an overall score for each candidate (line 10). The candidates are then ranked (line 12), and the highest scoring candidate is selected (line 13) and returned (line 14).

4 Empirical Evaluation

In this section we report on some initial empirical evaluation of the selection mechanisms. For the experiments the following parameters were fixed. The buyer agent attempted to find opponents for four negotiation goals. Goal one had a worth of 0.4, goal two had a worth of 0.6, goal three had a worth of 0.8 and goal four had a worth of 1. Each goal was negotiated over 20 times totaling 100 negotiations in each run. A total of 100 iterations of each run was performed whereupon averages were obtained for each measure presented. Thus, each experimental condition contained 1000 negotiations. There

Algorithm 1 The Opponent Selection Algorithm

Inputs:*motivations* — The buyer agent's motivations*goals* — The buyer agent's goals*exstnggoals* — The buyer agent's existing goals**outputs:***o* — The selected opponent**Algorithm:**

```
1: issues = determine issues(exstnggoals, g, motivations)
2: candidates = announceGoal(g)
3: for all c such that  $c \in \textit{candidates}$  do
4:   profile = getprofile(c)
5:   exptissues = getexpectedissueset(profile, g)
6:   conflict = getconflictset(issues, exptissues)
7:   cscore = scoreconflict(conflict)
8:   dp = getexpecteddealprice(profile)
9:   pscore = scoreexpecteddealprice(dp)
10:  overallscore = weightscores(cs, ps, M)
11: end for
12: rankcandidates(candidates)
13: o = selectopponent(candidates)
14: return o
```

was a population of nine seller agents all of which were available as potential opponents for all goals. The seller agents were split into three groups of three. In the first group the sellers selected each attribute of the negotiation goal for negotiation with a probability of 0.8, and had reservation values in the interval [1,5] dollars. The second group selected each attribute of the negotiation goal for negotiation with a probability of 0.5, and had reservation values in the interval [5,10] dollars. The third group selected each attribute of the negotiation goal for negotiation with a probability of 0.1, and had reservation values in the interval [10,15] dollars. This meant that group one offered high conflict but low cost negotiations, group two offered medium conflict and medium cost negotiations and group three offered low conflict but expensive negotiations.

For a negotiation to be successful, an agreement had to be found for all the issues. As the number of issues increased however, the chance that one of the issues caused the negotiation to fail increased. This was achieved by assigning a probability, p , for each issue that reaching an agreement on that issue would fail. For example, let $\rho(a)$ be the probability that $a \in \textit{issues}$ would lead to failure. If there was, say, n different a 's in *issues*, each with ρ chance of causing failure, the chance that at least one out of the three issues would do so is given by the following term, $1 - (1 - \rho)^n$. So, if $\rho(a) = 0.3$ and there are three a 's in *issues* we have $1 - (1 - 0.3)^3 = 0.65$ chance that one of the issues will lead to failure.

To determine the outcome of a negotiation in terms of price, we first checked to see if there was a *zone of agreement* (za), defined as the difference between the reservation prices of the buyer, $bres$, and seller, $sres$, i.e. $za = bres - sres$. If $za < 0$ then there was

no zone of agreement and the negotiation failed. If $z_a \geq 0$, the deal price of the negotiation was calculated as the mid point between the two agents' reservation positions, i.e. $dealprice = \frac{bres-sres}{2}$.

Our initial experiments sought to determine if the selection mechanisms performed as they should. Their performance was compared against a default strategy of randomly selecting an opponent for each negotiation. Figure 3 shows the results of opponent selection using different selection strategies in terms of the cost of negotiation.




Fig. 3. Average cost of negotiations using different selection criteria

In the figure, the x-axis shows the number of negotiations, while the y-axis shows the cost of the negotiation. All values shown are for successful negotiations. The bottom of the graph shows the worth of the goal under negotiation (the small dotted line), which increases over time. The line labelled *cost* indicates the selection of opponents of a buyer whose sole concern is to minimise cost, the line labelled *conflict* indicates a buyer whose sole concern is to minimise conflict, the line labelled *conflict and cost* indicates a buyer who attempts to minimise both cost and conflict. The graph shows that the *cost* focused buyer succeeds in minimising cost even when the worth of the goal increases showing an average cost over all negotiations of 3.9. The *conflict* focused buyer however does not manage this and, as the goal increases in worth, is led into selecting more costly negotiation opponents showing an average cost over all negotiations of 8.8. The *cost* and *conflict* focused buyer manages keep costs lower than the *conflict* focused agent but not the *cost* focused agent and shows an average cost of 4.5. The buyer using the random strategy displays wild behaviour jumping from high to low cost negotiations without a pattern, giving an average overall negotiation cost of 7.7.

Figure 4 shows the same buyers' performances when considering the frequency of successfully completed negotiations. It can be seen that, although the *cost* focused buyer manages to keep costs low as shown in Figure 3, the frequency of successful negotiations is low, displaying an average of 46% success rate over all negotiations conducted which remains relatively constant. This is expected since, as the agent is not concerned with minimising conflict, the success rate depends on the chance that incompatible issue choices will arise, which will, over all negotiations, be constant in this configuration. The *conflict* focused buyer however, performs much better, with the frequency of successful negotiation increasing steadily as the buyer obtains more information about the sellers to almost a 90% success rate for high worth goals and an average success rate of 71%, though the cost of these negotiations increases also (see Figure 3). The *conflict* and *cost* focused buyer again manages to make a compromise, keeping costs low and producing a success rate average of 58%. The buyer selecting randomly performs the same as the *conflict* and *cost* focused agent, successfully concluding 58% of negotiations, but combining success rate with cost, shows that the *conflict* and *cost* focused buyer easily outperforms the random selection strategy.

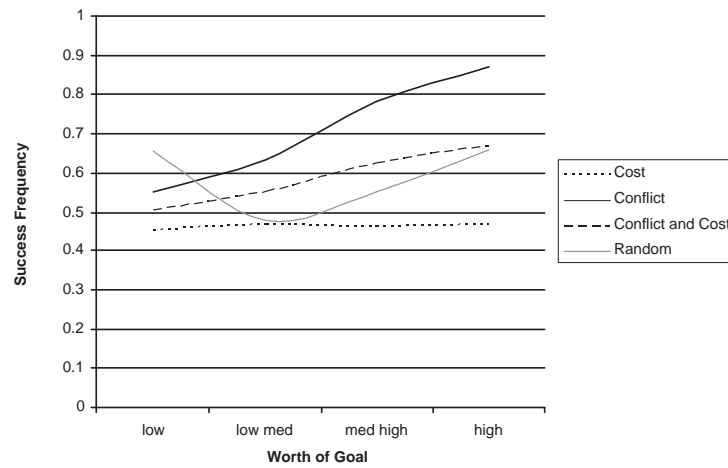


Fig. 4. Success rate for negotiations using different selection criteria

5 Related Work

Opponent selection has been investigated by a number of researchers. Work by Tesfatsion [9] examines how agents select opponents based on the amount by which they exceed fixed performance requirements on the issues of the negotiation. 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 targeted selection of negotiation opponents, and assumes fixed performance expectations, whereas we deal with dynamically changing performance requirements on conflict and price. In [10], 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 [11], 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 it affects selection of opponents, nor does it deal with considerations of conflict, but instead focuses on the efficacy of the decision heuristics.

The importance of motivation as an enabler of autonomy in computational agents was perhaps first identified by d’Inverno and Luck [12]. More recent efforts have extended Luck and d’Inverno’s ideas to consider planning [13] and norm-based multi-agent systems [14]. The use of motivation within negotiation, however, is a relatively new approach. One example is described in [15], where motivation is used to enable cooperative negotiations that aim to increase the utility of all participants. For pre-negotiation issues such as those addressed in this paper, one recent study is that of Urbig *et al.* [3], which examines 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 selection mechanisms to enable the selection of opponents in dynamic domains.

6 Conclusion and Future Work

By explicitly reasoning about the potential for conflict in negotiation and enabling agents to dynamically evaluate the constraints placed on the use of their resources we have shown how opponent selection in dynamic domains can be effectively conducted. Initial empirical results show that the basic approach works and that agents can be made to dynamically change the focus of their selections depending on the changing needs for conflict minimisation or cost minimisation. Future work will involve looking at how the amount of dynamism in the environment can be used to help the buyer to make decisions on how constrained it makes its negotiation goals. In highly dynamic environments the buyer may achieve more if it leaves more of the attributes of a negotiation goal open for the seller to determine or, to maintain some influence, sets of negotiable attribute-values can be identified from which the seller can choose one to instantiate the attribute. Limitations of our model include the lack of a proper treatment of negotiation strategies. The actual process of negotiation is abstracted away in order to simplify the model. Future work should address this by explicitly including the effects that the different strategies of both buyer and seller agents can have on the resolution of conflicting issues.

References

1. Sierra, C., Jennings, N.R., Noriega, P., Parsons, S.: A framework for argumentation-based negotiation. *Lecture Notes in Computer Science* **1365** (1998) 1774
2. Faratin, P., Sierra, C., Jennings, N.R.: Negotiation decision functions for autonomous agents. *Journal of Robotics and Autonomous Systems* **24** (1998) 159–182
3. 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)
4. Dash, R., Jennings, N., Parks, D.: Computational-mechanism design: A call to arms. *IEEE Intelligent Systems* **18** (2003) 40–47
5. Sabater, J., Sierra, C.: Social regret, a reputation model based on social relations. *SIGecom Exch.* **3** (2002) 44–56
6. Ramchurn, S.D., Sierra, C., Godo, L., Jennings, N.: Devising a trust model for multi-agent interactions using confidence and reputation. *Int. J. of Applied Artificial Intelligence* **18** (2004) to appear.
7. Munroe, S., Luck, M.: Motivation-based selection of negotiation opponents. In: LNAI 3451: Engineering Societies in the Agents World. Springer (2005) to appear.
8. Munroe, S., Luck, M., d’Inverno, M.: Towards motivation-based decisions for worth goals. In: Proceedings of the 3rd International Central and Eastern European Conference on Multi-Agent Systems. (2003)
9. Tesfatsion, L.: A trade network game with endogenous partner selection. In: Computational Approaches to Economic Problems. Kluwer (1997) 249–269
10. Banerjee, B., Sen, S.: Selecting partners. In: AGENTS 2000, Proceedings of the fourth international conference on Autonomous agents, ACM (2000)
11. Dutta, P.S., Moreau, L., Jennings, N.R.: Finding interaction partners using cognition-based decision strategies. In: Proceedings of The IJCAI-2003 workshop on Cognitive Modeling of Agents and Multi-Agent Interactions. (2003) 46–55
12. d’Inverno, M., Luck, M.: Understanding Agent Systems. Springer (2001)
13. Coddington, A.: Self-motivated Planning in Autonomous Agents. PhD thesis, University of London, London (2001)
14. Lopez y Lopez, F., Luck, M., d’Inverno, M.: Constraining autonomy through norms. In: Proceedings of the First International Conference on Autonomous Agents and Multi-Agent Systems. (2002)
15. Zhang, X., Lesser, V., Wagner, T.: A proposed approach to sophisticated negotiation. In AAAI Fall Symposium on Negotiation Methods for Autonomous Cooperative Systems (2001)