Towards a Theory of Negotiation Strategy

(A Preliminary Report)

Iyad Rahwan[*]
Dept. of Information Systems
University of Melbourne
Parkville 3010, Australia

Peter McBurney[†]
Dept. of Computer Science
University of Liverpool
Liverpool L69 7ZF UK

Liz Sonenberg[‡]
Dept. of Information Systems
University of Melbourne
Parkville 3010, Australia

ABSTRACT

This paper takes the first steps towards a generic theory of strategy in negotiation interactions between autonomous computational agents. Previous efforts at defining and studying negotiation strategies have done so within the context of a particular interaction protocol, or class of protocols. In order to develop a protocol-independent theory of strategy, we first identify the factors which may influence the creation of a strategy for a rational agent, and then consider some of these in more detail. As an example, we apply our generic approach to the Zeuthen strategy of the Monotonic Concession Protocol and a particular strategy from the International Trading Agent Competition.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*intelligent agents, multiagent systems*

General Terms

Theory, Design, Economics

Keywords

Negotiation, Strategy, Trading Agent Competition, Monotonic Concession Protocol, Zeuthen Strategy, Interaction Protocols

1. INTRODUCTION

In multi-agent applications, autonomous, self-interested agents often need to interact in order to fulfill their objectives. Such interaction may take numerous forms, and different types of interaction may be best suited for different environments or application domains. Following Walton and Krabbe [33], we take *negotiation* to be a *form of interaction in which a group of agents, with a desire to cooperate but with potentially conflicting interests, seek to reach a mutually-acceptable division of a scarce resource or resources.* These resources may be, for example, the free time or the future actions of the agents concerned. We assume that each agent engages in a negotiation interaction with a particular goal or goals in mind. For instance, an agent may seek to achieve the largest possible share of the resource in question, or it may seek the maximum possible share for itself and some subset of the agents engaged in the negotiation, or it may seek an equitable share for all participants, etc. Such individual agent goals may conflict with the goals of other agents, in the sense that not all goals can be achieved simultaneously. In this sense a negotiation interaction involves, at least potentially, some conflict between the participants. However, the participants are also co-operative, at least to the extent that they are willing to enter into joint interaction to agree on a division of the resource at issue.

To automate such processes, a number of interaction and decision mechanisms have been proposed and studied. These include: game-theoretic and auction-based mechanisms [10, 22, 25]; heuristic-based bilateral offer exchange (i.e., bargaining) mechanisms [2, 3, 9]; and argumentation-based approaches [11, 19, 28]. When we use the term *"negotiation framework"* in this paper, we refer to any of these models.

A central feature of all these mechanisms is that agents have some choice of what they may utter, and possibly when they may make these utterances. Open-cry auction participants, for example, choose both the content of their utterances (within the constraints of the particular auction protocol) and the timing of their utterances; participants in a sealed-bid, single-round auction may only choose the content. In argumentation-based approaches, participants have greater freedom in their choice of content and timing of utterances. A participant in a negotiation framework therefore faces an important question:

What should an agent say, and when, in a particular negotiation interaction?

A *negotiation strategy* may be defined as a rule or algorithm which provides an answer to this question. In this paper we aim at identifying the factors which could guide the design and selection of strategies for agents engaged in negotiation interactions. Of course, a strategy which is effective for one negotiation framework may be completely ineffective for another, and so the specific nature of the framework and its rules is an important factor in the design of strategies. However, despite this particularity, a generic theory of strategy in negotiation interactions may have value, and our work is a step in this direction.

This paper is organized as follows: We first present, in Section 2, a rationale for studying strategies for negotiation interactions in a framework-independent manner. This is followed, in Section 3, by a list of the factors which may guide the design and selection of strategies, and with de-

[*] i.rahwan@pgrad.unimelb.edu.au
[†] p.j.mcburney@csc.liv.ac.uk
[‡] l.sonenberg@dis.unimelb.edu.au

tailed discussion of some of these in Sections 4, 5 and 6. Section 7 considers, as examples, two negotiation strategies from the literature: the Zeuthen strategy for the 2-person Monotonic Concession Protocol; and a strategy used by one of the agents participating in the International Trading Agent Competition in 2002. The paper then concludes in Section 8.

2. RATIONALE FOR STUDYING STRATEGIES

In this section, we argue for the need for a generic characterization of strategy in negotiation. On a very abstract level, a strategy specifies what the agent should utter and when, in a negotiation interaction. We make the standard distinction between strategies, which govern an entire interaction or large parts of it, and *tactics*, which govern just a small number of utterances in an interaction. Posing a direct question to another participant at a particular point in a negotiation may be viewed as a tactic, for example, which implements a strategy to gather information about that participant's preferences.

In the automated negotiation literature, analysis of negotiation strategies has been conducted using methods based on game-theoretic analysis, heuristic experimentation, and argumentation-based studies. In game-theoretic analysis, researchers usually attempt to determine the optimal strategy by analyzing the interaction as a game between identical participants, and seeking its equilibrium [8, 22, 29]. The strategy determined by these methods is optimal for a participant, given the game rules, the assumed payoffs, and the goals of the participants, and assuming that the participants have no knowledge of one another not provided by introspection. On a further assumption that participants behave according to the assumptions of rational-choice theory, then this approach can guide the design of the interaction mechanism itself, and thus force such agents to behave in certain ways [32]. Of course, in the real world, agents may be resource-constrained, malicious or whimsical, or simply badly-coded, so that participant behaviour may not conform to the assumptions of rational choice theory. Recently, research on bounded-rationality in game theory has started to address some of these issues [23].

In cases where it is not possible to reach the optimal outcome (for example due to resource limitations), some heuristics have been devised. Heuristics are rules of thumb that produce 'good enough' outcomes, and are mainly based on empirical testing and evaluation. In heuristic-based frameworks, strategies have been proposed which are based on the underlying utility model, or on decay functions of factors such as utility and time [2, 10]. Finally, argument-based frameworks are negotiation frameworks that allow agents to exchange, in addition to proposals and indications of their acceptance or rejection, meta-information about them, such as the reasons for their proposal, and for accepting or rejecting them. Strategies have been presented for some of these frameworks also; for example, Sierra and colleagues [28] consider the execution of an authoritarian strategy in which an agent makes an appeal to its authority over others in order to exert pressure on its negotiating counterpart. Similarly, Sadri and colleagues [24] describe agents negotiating over scarce resources who are always willing to share resources which are not currently needed for their own goals, a rule which partly determines the utterances each agent may make in the interaction.

However, most existing frameworks in the literature focus on particular strategies under specific negotiation frameworks. The frameworks are typically defined as specific interaction protocols and communication languages; the decision making mechanisms of the participants may also be specified. Analysis of strategies may then be undertaken in controlled settings under various conditions, for example, under time constraints or with incomplete information. Very specific bargaining strategies have been analyzed, for example, by investigating their optimality [4], their performance in multiple negotiation rounds [2], or the resulting social welfare [8]. There has been little work, to our knowledge, which looks at strategies in negotiation interactions in a generic way.

There may be good reasons for this. One reason is that, as we mentioned above, effective strategies are likely to be protocol-dependent. However, there is still no formal theory of interaction protocols covering all types of mechanism. It is to be expected therefore that such work will focus first on defining the protocols and exploring their properties, before devising strategies for their participants. Another reason for the absence of a generic theory of strategy is that the study of negotiation mechanisms, broadly construed, is still immature. In particular, there is no broad understanding of the factors that could influence the design of strategies. This paper takes a step towards characterising such factors.

One might argue that the game-theoretic notion of strategy may be sufficient for a generic theory of strategy. However, while mathematical game theory has been applied to auction mechanisms for half a century, the articulation of other types of negotiation protocols is a recent phenomenon. Several researchers have argued that game-theoretic negotiation mechanisms fail to provide systematic means that allow participants to persuade one another to change their preferences and beliefs, or to reshape the negotiation object itself, and so on. As a result, new negotiation protocols are emerging, building on notions of persuasion and logical argumentation [17, 19, 11, 31, 30]. It is not clear yet whether the game-theoretic characterisation of strategy is suitable for these frameworks. It is also not obvious that strategies which are optimal under assumptions of participant rationality and common knowledge of participant rationality will also be optimal when these assumptions are violated.

Why would a protocol-independent theory of strategy be useful? In addition to potentially providing an understanding of the nature of strategy, a generic theory of strategy may also provide a common framework for comparison of strategies across negotiation frameworks. This should then provide assistance to agents (or their designers) in selecting between strategies. For example, why should an agent adopt a particular decay function for conceding on its utility rather than a linear concession strategy? Why should an agent choose to include time, trust, or the history of the interaction in the design of strategies? Why should an argumentative agent resort to making threats rather than seeking to learn more about its negotiating counterpart? Why should an agent always share resources it does not currently need? It is hoped that a general theory of strategy design, assessment and selection would provide a means to answer these questions.

Another potential value of a generic theory of strategy is related to the types of generic conclusions and analytical results we can make about strategic interaction. In particular, it may be easier to produce general conclusions about the design and selection of strategies by generic approaches, in a top-down manner, than by generalizing the conclusions of protocol-specific investigations, in a bottom-up manner.

Finally, it is not obvious *before* an investigation of a generic theory of strategy that nothing useful will be discovered by it. This, we believe, is a sufficient motivation to move forward in our attempt.

3. STRATEGY DESIGN

What factors may influence the design of strategies for an agent in a negotiation interaction? In asking this question, we are not seeking to be descriptive — *what factors do influence strategy?* — nor seeking to be prescriptive — *what factors should influence strategy?* Rather, we seek to identify the possible factors which may inform strategy design for an agent engaged in purposeful behaviour.[1] The following factors may influence the design of strategies for a computational agent engaged in a negotiation interaction with other agents:

Goals: What objectives or goals the agent wishes to achieve from undertaking a negotiation interaction over these resources with these other agents at this particular time. As mentioned earlier, an agent's goals may pertain only to itself, or may have an altruistic element. Moreover, an agent may have multiple goals, some of which may be made explicit and some not. Indeed, an agent may enter into an interaction with no intention of seeking a division of the resources in question, but merely to confuse, distract, or otherwise delude the other participants, or even non-participants.

Domain: Strategies may differ according to the nature of the resources under negotiation. Jeffrey Rosenschein and Gilad Zlotkin [22], for example, distinguish between task-oriented and state-oriented domains.

Protocol: The nature of the interaction protocol used for the negotiation. A strategy of providing supporting information, for example, makes no sense for a protocol which does not allow such information to be presented in the negotiation interaction.

Capabilities: The capabilities of the agent within the interaction. Actions or utterances may be legally valid under a given protocol, but not be in the repertoire of the agent concerned.

Values: The values of the agent. Some actions may be permitted by the protocol, and within the capabilities of the agent, but not be actions which the agent would ever execute. An agent may refuse to tell lies, for example. Mark Young [34] has recently argued persuasively that most human negotiators seek to maintain a coherent personality profile through their business life, and thus do not even consider actions or utterances in a negotiation which are inconsistent with their understanding of this profile.

[1]We are therefore excluding strategy design for agents engaged in random behaviours or whimsy, except insofar as such behaviours have a higher purpose.

Counterparts: The nature of the other participants in the interaction, as perceived by the agent. Clearly, a strategy appropriate for one type of counterpart may be inappropriate for another, and so the agent's models of its counterparts may well influence its strategy design. Experiences of past interactions with the same or similar agents may inform these counterparty models, and thus also strategy design.

Resources: The time and resources available to the agent, including computational, memory, and other resources, such as expert advice. For example, in legal disputes, participants may delay resolution in order to force an opponent with fewer resources to settle the dispute or to withdraw from the interaction [1].

Alternatives: The nature of any alternatives to resolution available to the agent. One popular book on negotiation [6] advises participants to develop their *"Best Alternative to a Negotiated Agreement"*, so as to guide their negotiation strategy.

In the next three sections, we discuss some of these influencing factors in more detail.

4. NEGOTIATION GOALS

Since we are dealing only with purposeful agents, then each party to a negotiation may be assumed to have some goal or goals which leads it to enter into the negotiation interaction itself. These goals may be at the highest level of an agent's stack of goals or they may not be, in which case they may support some other, even higher, goals. However, for the purposes of analysis of negotiation strategy we may assume this goal is at the top-most level of the agent's hierarchy of goals. Entering into a particular negotiation interaction over certain resources with particular agents at a particular time, will, the agent believes, assist it in seeking to achieve these goals.

An agent's negotiation goals may be a particular agreement to divide the scarce resources under discussion. Such an outcome is not the only goal an agent may have. For example, an agent may engage in a negotiation interaction to acquire information about a new domain, as when potential house-buyers participate in auctions in order, not to purchase a house, but to learn about prevailing house prices, or even to learn about the auction process itself. Similarly, an agent may enter into a particular negotiation interaction in order to establish or maintain a larger relationship with the other agents concerned, or to gain knowledge about such agents and their negotiating behaviours. As an example, John Lukacs [15] argues that, during May 1940, Prime Minister Winston Churchill of Britain pretended to entertain the possibility of a negotiated peace deal with Nazi-led Germany in order to strengthen his base of political support with key members of his own Conservative Party; Lukacs argues that Churchill was not serious about these possibilities, but in the early days of his Premiership he needed political support from people who were.

Agents may even enter into a negotiation interaction with one counterpart in order to have a stronger negotiation position relative to another counterpart in a separate interaction; business-to-business negotiations often involve such parallel, competitive negotiations [14, Chapter 3]. These objectives are all valid — and, by any definition, rational

— objectives from a negotiation interaction. We make no judgment on their wisdom, feasibility, or ethical content.

5. CAPABILITIES

A key influence on strategy design will be the interaction capabilities of the agent negotiator — what is the agent capable of doing in the interaction.[2] We first present a set of essential capabilities, then discuss some constraints that may limit the agent's ability to execute its capabilities.

At the bottom level, an agent engaged in a negotiation interaction must be able to make utterances which are legal according to the rules of the protocol. Above this level are some higher-order capabilities, which may, depending on the specific protocol, require utterance of a sequence of locutions to be effected:

1. **Making proposed deals.**
2. **Accepting proposed deals.**
3. **Rejecting proposed deals.**
4. **Presenting information proactively to a counterparty.** An agent may present information in order to influence a counterparty's beliefs or intentions; for example, describing the alternative deals in such a way that a particular proposed deal appears as the most attractive.
5. **Seeking information from a counterparty.** Participants may have varying capabilities to extract information from their counterparties, for example, due to differing levels of authority in a social structure.
6. **Providing information reactively to a counterparty.** Agents may have differing capabilities to provide information to one another; some agents may not be able to lie, or to answer evasively, for example.
7. **Seeking to exert pressure on a counterparty.** Depending on some factors, such as the social structure between the participants, an agent might have the ability to threaten or reward other participants for adopting particular intentions or behaviours. Interaction protocols allowing such capabilities have been proposed in [11, 28].
8. **Retracting Commitments.** Depending on the rules of the protocol, agents may have ability to retract commitments or proposals they have made previously. Retraction has been studied, for example, in argumentation theory [33] and in game-theoretic studies of bargaining [26].
9. **Withdrawal.** Under an assumption of true agent autonomy, agents have the ability to withdraw from any interaction at any stage. Depending on the protocol, they may also have the ability to threaten to withdraw.

Of course, an agent may be said to also have capabilities which are complex combinations of these. For example, an ability to prevaricate may be constructed from abilities to: request irrelevant information; provide irrelevant, misleading or confusing information; or repeat previous questions or statements.

In discussion of the influence of an agent's capabilities on strategy design, we also need to take account of constraints on the exercise of any potential capabilities. Such constraints include:

Interaction Protocol: The rules of the negotiation interaction protocol may preclude or require certain utterances or certain types of utterances by agents at particular times in an interaction. The FIPA Agent Communications Language, FIPA ACL, for example [5], requires agent sincerity: only statements which are believed by an agent may be uttered using the *inform* locution. In principle, such a condition must severely limit the use of FIPA ACL for negotiations.

Values: As mentioned earlier, the agent's values may preclude or require certain behaviours, and so constrain the potential capabilities of the agent.

Resource Constraints: Time, memory or processing limitations on an agent may limit its capabilities in a negotiation interaction.

6. TACTICS

As stated earlier, we assume an agent enters a particular negotiation interaction over particular resources using a particular interaction protocol with particular counterpart agents at a particular time, in order to achieve its negotiation goals. Having decided to so enter a particular negotiation, an agent may adopt sub-goals for portions of the interaction, with the belief that such sub-goals assist in realizing the overall negotiation goals. Sub-goals themselves may be further decomposed into lower-level goals, and so on.

For example, a potential buyer entering into a negotiation with a car-dealer aiming to buy a car may seek to achieve this negotiation goal by realizing each of the following subgoals (in sequence):

A. Learning about the alternative models available from the dealer;

B. Establishing a preference ordering over some or all of these models; and

C. Getting the cheapest price for the most-preferred model.

The buyer might achieve the first sub-goal by posing a series of questions to the car dealer. The second sub-goal may be achieved by introspection, perhaps involving a process of comparison of the expected utilities of different models [21]. To achieve the third sub-goal, the buyer may seek to achieve two lower-level goals:

C.1 Informing the dealer about an offer made by a competing dealer; and

C.2 Bargaining with the dealer through an exchange of offers.

Each of these sub-sub-goals may be achieved directly by making a series of utterances, or through decomposition into

[2]Here we are not referring to the agent's internal capabilities, such as its ability to evaluate an offer.

further sub-goals, and so on. The hierarchical goal structure we have outlined here has a structural similarity to the landmarks theory of conversation protocols of Sanjeev Kumar and colleagues [12]. However, our approach concerns only the goals of an individual agent and not the joint goals of all participants to an interaction.

We now list a number of possible higher-level sub-goals that may contribute to an agent's achievement of its negotiation goals. Following standard English usage, we call these sub-goals, *tactics*. To our knowledge, no comprehensive list of all possible applicable negotiation tactics is available in the multi-agent literature or in the literature on human negotiation. Therefore, we list those tactics inspired by our experience with the multi-agent negotiation literature as well as informal advice to negotiation participants [6, 34, 13].

1. **Seek to change a counterpart's beliefs.** One participant in a negotiation may judge it to be in its interests to have other participants believe certain propositions about the beliefs, intentions, preferences, or constraints of the first participant, or about the domain in question. These propositions may be true or false. Providing information to the counterparts may enable an agent to explain the reasons for its beliefs, preferences, etc. It has been argued [6, 20] that agreement is more likely in negotiation interactions when participants understand each others' interests (desires, preferences, goals, etc.) rather than their current positions.

2. **Gain a better understanding of a counterpart.** Counterparts may be seeking to mislead a participant about their beliefs, intentions, preferences, constraints, etc., or about the domain. An agent may then seek to gain a better understanding of its counterparts' true mental states or constraints.

3. **Seek to discuss a particular issue.** By moving the interaction towards particular issues, a participant may be able to frame the problem in certain ways, and thus influence the mental states of its counterparts. A seller of a particular make of car, for example, may seek to turn the topic of discussions with potential buyers towards attributes on which this make of car scores highly.

4. **Seek to avoid discussion.** For the same reasons, a participant may wish to steer discussion away from particular issues.

5. **Seek fast termination.** An agent with time or processing resource constraints might seek a fast resolution or termination of the negotiation.

6. **Seek to delay.** An agent who believes it has greater time or other resources than other participants may seek to delay resolution of the interaction beyond the perceived resource limits of its counter-party.

7. **Resist a counterpart.** An agent may resist attempts by a counterpart to achieve one of the above tactics.

Note that an agent may change its tactics and even its goals in the course of negotiation. An agent might abandon a goal or a tactic if the agent perceives that this goal or tactic is not currently achievable or is counter-productive, for example. As with any other intentions, the defeasibility of goals in a computational agent requires some structure of intention-reconsideration, e.g., [27].

7. EXAMPLES

In this section, we present two examples from the multi-agent negotiation literature and analyse them from the perspective presented above. Our aim is to demonstrate that the strategic criteria adopted by agent designers in these frameworks fall within our proposed sketch. This, we hope, would support our claim that a general theory of strategy may by applicable in a top-down fashion to inform the design and improvement of strategies in specific settings.

7.1 Monotonic Concession Protocol

We now consider the Monotonic Concession Protocol [22, pp. 40–49] for negotiation between two agents. Under this protocol, the two agents make proposed deals to one another in a sequence of rounds, where the deal space can be described by some single real- or integer-valued dimension (call it *utility*). The protocol begins with each participant simultaneously proposing a deal to the other. If these proposed deals overlap, the interaction terminates successfully with a deal. If not, the interaction continues to a subsequent round. At each subsequent round, participants may repeat the offer they have just made, or make a new offer which is closer to the opponent's end of the single dimension (i.e., they may concede). Again, if the proposed deals at any round overlap, the interaction ends with a deal. If neither agent concedes at a given round, the interaction terminates without a deal, a situation called the *conflict deal*.

When should a participant concede, and, if so, by how much? In [35], Frederick Zeuthen proposed a strategy, now commonly called the Zeuthen strategy, for an agent to decide these questions. Subsequently, John Harsanyi [8] showed that the Monotonic Concession Protocol ends in a Pareto-optimal deal if both agents use the Zeuthen strategy. The Zeuthen strategy assumes that each agent knows its own utility of each proposed deal and the utility of the conflict deal, and that each agent also knows these utility values for the other agent. The strategy requires an agent to calculate its degree of willingness to risk a conflict at any round on the basis of the difference between its loss of utility in fully conceding at this round and accepting the other agent's current offer, relative to its loss of utility in not conceding at this round and causing a conflict deal. The agent then compares its own risk at this round with that of the other agent. If its own risk is greater than that of its opponent, then the agent should not concede; if its own risk is less than or equal to that of its opponent, it should concede just enough to shift the balance of risks between the two agents.

Let us now consider this strategy against the list of eight factors presented in Section 3 which may influence the design of a strategy. Factor 1: The negotiation goal of the agent is assumed to be the achievement of the best possible deal with utility above that of the conflict deal. If the only deals possible have utility less than the conflict deal, then the agent would prefer conflict. Factor 2: The domain is assumed to be representable by a single utility dimension. Factor 3: The Zeuthen strategy is specific to the Monotonic Concession Protocol, and so is informed by the utterances legal under this protocol, namely, repeating the most recent offer or making a concession. Factor 4: The interaction capabilities of an agent are assumed (implicitly) to be identical with the legal utterances. Factor 5: The only agent value which informs the design of the Zeuthen strategy is an assumption that each agent seeks to maximize its own utility in the negotiation, without regard to the utility achieved by the other agent in any deal. However, Harsanyi's result [8] may be seen as evidence of an *invisible hand* at work, so that a socially desirable outcome is achieved through purely selfish behaviour, provided both opponents use this strategy. This result may provide solace to a self-interested agent with a social conscience.

Factor 6: The Zeuthen strategy takes explicit account of the agent's opponent, making the (unrealistic) assumption that the agent knows its opponent's utilities. Factor 7: The strategy takes no account of the time, computational or other resources available to the participants. It would be an interesting question to determine the circumstances under which the strategy is sensible for resource-bounded agents. Factor 8: Finally, the Zeuthen strategy does consider the alternatives each agent has available via consideration of its utility loss arising from a conflict deal. However, the incorporation of alternatives in this summarized way would be considered insufficiently subtle by negotiation advisors such as [6, 34].

Note that the design of the Zeuthen strategy is influenced by a number of strong assumptions about the agents' capabilities, values, rationality and common knowledge. Moreover, the interaction protocol is fairly simple and only allows agents to either make proposed deals or concede. This is one reason why the analytical tools of game-theory were able to establish interesting properties of the negotiation process (e.g., Pareto-optimality of the outcome). Hence, given these assumptions about the agents involved and given common knowledge by agent designers of game theory, the strategy which designers will adopt becomes deterministic. In other words, strategy designers have no room for deviating from the Zeuthen strategy without risking a worse outcome than they could expect with the Zeuthen strategy.

7.2 Trading Agent Competition

When agent designers relax some of their assumptions, particularly regarding unbounded rationality and common knowledge, they immediately fall outside the region of predictability of classical game-theory. In such cases, approximate strategies (or heuristics) must be devised and analytical results become hard to achieve. Instead, experimentation through simulation becomes a more viable option. The Trading Agent Competition (TAC) (see [7] for an overview), of which three annual competitions have been conducted to date, involves multiple competing agents bidding in simultaneous auctions. In this subsection, we consider a strategy adopted by THALIS, the 3rd top scorer in TAC-02.

Eight agents participated in each TAC-02 game. Each agent performed the role of a travel agent attempting to provide booking for eight clients travelling from TACtown to Tampa and back during a five-day period. Each client was characterised by a random set of preferences for arrival and departure dates, hotels and entertainment tickets. Utility was gained by purchasing a complete package and was calculated based on comparison with the corresponding client's preferences. Package constituents were sold in separate simultaneous auctions, each with certain price dynamics. Airline tickets were sold in single round continuous auctions with biased random pricing that was more likely to increase. Hotel bookings were sold in ascending English auctions clearing every minute, while entertainment tickets were traded in continuous double auctions. The score of an agent was the difference between the total utility gained for its clients and the agent's expenditure.

We now consider the strategy adopted by THALIS, which was developed at the University of Essex, UK, for buying hotel bookings against the factors we presented in Section 3. THALIS started bidding for hotels only after the start of the fourth minute in the game, and bid relatively high (based on statistics it had collected thus far) [7]. Factor 1: The negotiation goal of each agent is to maximize the total score as described above. Factor 2: The domain is representable in terms of multiple interrelated attribute, with a corresponding preference function used to derive utility. Factor 3: The strategy operates in an ascending English auction protocol. Factor 4: All agents have symmetric abilities of making bids or skipping turns. In this particular strategy, THALIS exploits the ability to *wait* in order to obtain information about average hotel bidding prices. Factor 5: The only value in consideration is assumed to be the maximization of the overall score (as opposed to the hotel specific score).

Factor 6: It is hard to assess the extent to which the designers of the THALIS agent considered strategies of counterparts in determining their agent's strategy. It may be that the designers believe, perhaps through previous experience, that the data gathered in the first four minutes would be sufficient to give a good estimate of prices and price dynamics. Alternatively, the designers may be assuming that hotel booking prices are not likely to increase significantly over this initial period, and so there is no need to bid. Factor 7: The strategy takes account of the time resources needed to gather the statistics needed to estimate the average bidding price. Factor 8: Finally, assessment of an agent's alternatives enters only indirectly and incompletely, via the agent's utility function. There appears to be no allowance, for example, for opportunity costs, such as the cost of having an airline ticket but no hotel reservation, and so having to travel without a guaranteed room.

From this example, it is clear that as the trading mechanism becomes more complex, the factors that may influence the design of a strategy also increase in complexity. This can be seen by comparing the analyses we have undertaken of the Zeuthen and THALIS TAC-02 strategies. We believe that our proposed list constitutes a comprehensive set of classes of strategic factors which could guide strategy analysis and generation. In particular, this structure allows us to see how strategies may be extended or refined.

8. CONCLUSIONS

In this paper, we have defined a number of factors which may influence the generation of strategies for agents engaged in negotiations with one another over the allocations of potentially-scarce resources. Previous work on defining and studying negotiation strategies has done so within the context of a particular interaction protocol, or class of protocols. In order to develop a protocol-independent theory of strategy, we have identified the factors which may influence the creation of a strategy for a rational agent, and then considered some of these in more detail. As an example, we applied our generic approach to the Zeuthen strategy of the

Monotonic Concession Protocol and a particular strategy exercised by an agent participating in the Trading Agent Competition in 2002.

Every science begins with a classification. The research reported here is at a very preliminary stage, and much remains to be done before a full theory of strategy can been developed. For example, we have not attempted yet to formalize these notions. It may be the case that a formal, protocol-independent theory of strategy is impossible without a formal, generic theory of protocols. While formal theories of some types of negotiation protocols exist — e.g., of auction mechanisms [18], or of dialogue game protocols [16] — there is as yet no formal theory of interaction protocols in general.

An interesting question to raise is: how would it be possible to evaluate a theory of strategy? We believe that one answer lies in the usefulness of the theory. Such a theory would potentially allow comprehensive analysis of strategies across different protocols, and so provide a better understanding of the differences between protocols and between strategies. The theory would thus be useful to the study of protocols in a manner similar to Ariel Rubinstein's view of the usefulnes of modeling to economics: *"Models of economic theory are meant to establish "linkages" between the concepts and statements that appear in our daily thinking on economic situations"* [23, p. 191]. In addition, a theory of strategy may aid the construction of a repertoire of strategies and tactics available to agent strategy designers. As a further possible development, an associated compositional theory of tactics, structuring the process of constructing complex tactics and strategies from other, simpler tactics, may potentially lead to the automation of the strategy design process. This would extend the application of the agent paradigm from automated negotiation agents to automated negotiation agent designers.

Acknowledgments

The first author is grateful for financial support through a Ph.D. top-up scholarship from CSIRO Mathematical and Information Sciences, Australia, and a PORES scholarship from the University of Melbourne. This work commenced while the first author was a visitor to the Agent ART Group of the Department of Computer Science at the University of Liverpool, UK.

9. REFERENCES

[1] P. E. Dunne. Prevarication in dispute protocols. In G. Sartor, editor, *Proceedings of the 9th Internatonal Conference on AI and Law (ICAIL-2003)*, New York, NY, USA, 2003. ACM Press.

[2] P. Faratin. *Automated Service Negotiation Between Autonomous Computational Agents*. PhD thesis, Department of Electronic Engineering, Queen Mary and Westfield College, University of London, London, UK, 2000.

[3] S. Fatima, M. Wooldridge, and N. R. Jennings. Multi-issue negotiation under time constraints. In C. Castelfranchi and W. L. Johnson, editors, *Proc. 1st Intern. Joint Conf. on Autonomous Agents and Multiagent Systems (AAMAS-2002)*, pages 143–150. ACM Press, 2002.

[4] S. Fatima, M. J. Wooldridge, and N. R. Jennings. Optimal negotiation strategies for agents with incomplete information. In J-J. Meyer and M. Tambe, editors, *Intelligent Agent series VIII: Proc. 8th Intern. Workshop on Agent Theories, Architectures, and Languages (ATAL 2001)*, volume 2333 of *LNCS*, pages 53–68. Springer, 2001.

[5] FIPA. Communicative Act Library Specification. Technical Report XC00037H, Foundation for Intelligent Physical Agents, 10 August 2001.

[6] R. Fisher, W. Ury, and B. Patton. *Getting to Yes: Negotiating Agreement Without Giving In*. Century Business, London, UK, second edition, 1991.

[7] A. Greenwald. The 2002 trading agent competition: An overview of agent strategies. *AI Magazine*, 24(1):83–91, 2003.

[8] J. C. Harsanyi. Approaches to the bargaining problem before and after the theory of games: a critical discussion of Zeuthen's, Hicks', and Nash's theories. *Econometrica*, 24:144–157, 1956.

[9] R. Kowalczyk and V. Bui. On constraint-based reasoning in e-negotiation agents. In F. Dignum and U. Cortés, editors, *Agent-Mediated Electronic Commerce III*, volume 2003 of *LNCS*, pages 31–46. Springer, 2001.

[10] S. Kraus. *Strategic Negotiation in Multi-Agent Environments*. MIT Press, Cambridge MA, USA, 2001.

[11] S. Kraus, K. Sycara, and A. Evenchik. Reaching agreements through argumentation: A logical model and implementation. *Artificial Intelligence*, 104(1–2):1–69, 1998.

[12] S. Kumar, M. J. Huber, and P. R. Cohen. Representing and executing protocols as joint actions. In C. Castelfranchi and W. L. Johnson, editors, *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2002), Bologna, Italy*, pages 543–550, New York City, NY, USA, 2002. ACM Press.

[13] R. J. Lewicki, D. M. Saunders, and J. W. Minton, editors. *Essentials of Negotiation*. McGraw-Hill, Boston, MA, 2003.

[14] G. L. Lilien, P. Kotler, and K. S. Moorthy. *Marketing Models*. Prentice-Hall, Englewood Cliffs, NJ, USA, 1992.

[15] J. Lukacs. *Five Days in London: May 1940*. Yale University Press, New Haven, CT, USA, 1999.

[16] P. McBurney and S. Parsons. Games that agents play: A formal framework for dialogues between autonomous agents. *Journal of Logic, Language and Information*, 11(3):315–334, 2002.

[17] P. McBurney, R. M. van Eijk, S. Parsons, and L. Amgoud. A dialogue-game protocol for agent purchase negotiations. *Journal of Autonomous Agents and Multi-Agent Systems*, 7(1–2), 2003. *In press.*

[18] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, Cambridge, MA, USA, 1994.

[19] S. Parsons, C. Sierra, and N. Jennings. Agents that reason and negotiate by arguing. *Journal of Logic and Computation*, 8(3):261–292, 1998.

[20] I. Rahwan, E. Sonenberg, and F. Dignum. Towards interest-based negotiation. In T. Sandholm and M. Yokoo, editors, *Proc. 2nd Intern. Joint Conf. on*

Autonomous Agents and Multiagent Systems (AAMAS 2003), New York, NY, USA, 2003. ACM Press. *To appear.*

[21] J. H. Roberts and J. M. Lattin. Development and testing of a model of consideration set composition. *Journal of Marketing Research*, 28:429–440, 1991.

[22] J. S. Rosenschein and G. Zlotkin. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*. MIT Press, Cambridge MA, USA, 1994.

[23] A. Rubinstein. *Modeling Bounded Rationality*. MIT Press, Cambridge, MA, USA, 1997.

[24] F. Sadri, F. Toni, and P. Torroni. Logic agents, dialogues and negotiation: an abductive approach. In M. Schroeder and K. Stathis, editors, *Proc. Symposium on Information Agents for E-Commerce, Artificial Intelligence and the Simulation of Behaviour Conference (AISB-2001)*, York, UK, 2001. AISB.

[25] T. Sandholm. eMediator: A next generation electronic commerce server. *Computational Intelligence*, 18(4):656–676, 2002.

[26] T. Sandholm and V. Lesser. Leveled commitment contracting: A backtracking instrument for multiagent systems. *AI Magazine*, 23(3):89–100, 2002.

[27] M. Schut and M. Wooldridge. The control of reasoning in resource-bounded agents. *Knowledge Engineering Review*, 16(3):215–240, 2001.

[28] C. Sierra, N. R. Jennings, P. Noriega, and S. Parsons. A framework for argumentation-based negotiation. In M. Singh, A. Rao, and M. J. Wooldridge, editors, *Intelligent Agent IV: 4th International Workshop on Agent Theories, Architectures and Languages (ATAL 1997)*, volume 1365 of *LNCS*, pages 177–192. Springer, 1998.

[29] B. von Stengel. Computing equilibria for two-person games. In R. Aumann and S. Hart, editors, *Handbook of Game Theory*, volume 3, chapter 45, pages 1723–1759. North-Holland, Amsterdam, 2002.

[30] K. Sycara. The PERSUADER. In D. Shapiro, editor, *The Encyclopedia of Artificial Intelligence*. John Wiley and Sons, January 1992.

[31] F. Tohmé. Negotiation and defeasible reasons for choice. In *Proc. Stanford Spring Symp. on Qualitative Preferences in Delbiberation and Practical Reasoning*, pages 95–102, 1997.

[32] H. R. Varian. Mechanism design for computerized agents. In *Proc. USENIX Workshop on Electronic Commerce*, New York, NY, USA, July 1995.

[33] D. N. Walton and E. C. W. Krabbe. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. SUNY Press, Albany, NY, USA, 1995.

[34] M. A. Young. *Rational Games: A Philosophy of Business Negotiation from Practical Reason*. Quorum Books, Westport, CT, USA, 2001.

[35] F. L. B. Zeuthen. *Problems of Monopoly and Economic Warfare*. Routledge and Sons, London, UK, 1930.