Complex negotiations in multi-agent systems ⋆

Mihnea Scafeş * **Costin Bădică** **

* *University of Craiova, Software Engineering Department, Bvd.Decebal 107, Craiova, 200440, Romania (e-mail: scafes_mihnea@software.ucv.ro)*
** *University of Craiova, Software Engineering Department, Bvd.Decebal 107, Craiova, 200440, Romania (e-mail: costin.badica@software.ucv.ro)*

Abstract: This paper presents an overview of existing models of automated negotiation in multi-agent systems with a special focus on complex negotiations involving non-linear utility functions.

Keywords: Automated negotiation, Multi-agent Systems

1. INTRODUCTION

Negotiation is used in business (to agree over the price of an item), politics (to negotiate between countries over some regional resources they want to use) and various other domains. As the information systems became more and more advanced, negotiation started to be used between such systems, using computers. Over the last years researchers tried to automate the negotiation process. They used computer science for implementation and analysis of negotiation algorithms. They spent considerable effort trying to find better negotiation models that lead to better outcomes, as well as to improve the algorithms to automatically compute the negotiation outcome. Automated negotiation is often studied in the field of multi-agent systems (MAS), a research field in which researchers combine techniques derived from distributed systems, artificial intelligence and computer science.

Generally, negotiation brings together three topics (Jennings et al., 2001): negotiation protocols, negotiation subject and negotiation strategies. The protocols are the rules that negotiation participants must obey during negotiation. A protocol describes the steps of a negotiation, what messages can be sent, and what actions participants are allowed to take during each phase of the negotiation. There are various types of protocols, for different types of negotiation:

- Contract Net protocol (CNP) has been introduced by Smith (Smith, 1980) for distributed problem solving. It is largely used for task allocation problems. Using this protocol, agents negotiate about tasks. One agent (called *manager*), which is interested in performing a task, announces the other agents that the task is available. Either the manager agent might not be capable of performing the task on its own or it might try to find other agents (called *contractors*) that are able to perform the task more effectively (less cost, more precision a.o., depending on the problem domain). The agents that are interested in carrying out the task (called *contractors*) submit bids to the manager. The manager awards the task to the contractor that sent the most satisfactory bids. The contractor then starts the task and eventually informs the manager when the execution of the task is finished. The bidding, bid processing and task processing phases are strongly dependent on the problem. For related work on task allocation, see Section 2.
- Rubinstein's alternating offers protocol (Osborne and Rubinstein, 1994). Multiple agents negotiate by taking actions in turns. At each turn one agent makes a proposal to the other agents, which can accept or reject the proposal. If all agents accept it, negotiation ends, otherwise, the next agent makes a proposal at the next turn. This protocol is normally used for bargaining problems. For related work on bargaining, see Section 2.
- Monotonic Concession Protocol (MCP) (Rosenschein and Zlotkin, 1994) is a particular case of Rubinstein's alternating offers model. This protocol forces an agent to make a concession to the other agent at each round. It is guaranteed to stop, but it requires that agents know each other's preferences, which is undesirable in practice in a competitive setting. MCP involving multiple agents has been proposed in (Endriss, 2006; Bădică and Scafeş, 2009).
- Auctions are particular negotiation protocols used for multilateral negotiations. Agents bid for items and special agents called *auctioneers* evaluate bids and allocate items. There are many types of auctions, some well known auction types being the English auction, the Dutch auction, first-price sealed-bid, second-price sealed-bid, and combinatorial auctions (Vidal, 2006).

The negotiation subject or negotiation object, as it is also called, describes what is being negotiated between

⋆ Mihnea Scafeş was supported by IOSUD-AMPOSDRU contract 109/25.09.2008. Costin Bădică was supported by DIADEM project. Diadem project is funded by European Union under Information and Communication Technologies (ICT) theme of the 7th Framework Programme for R&D, ref. no: 224318.

the partners. It can be a state of the environment the participants what to reach, an action they would like to perform or an item they would like to have. If the subject is composed of multiple attributes it is called *multi-issue*, otherwise it is called *single-issue*. For example, when a car dealer and a client are negotiating about a car, they might negotiate about the price of the car, the engine and other options the client might want, which means they are negotiating about a multi-issue subject. Further more, the issues are *divisible* if they can be split between the participants (i.e. if all the participants can get a share of the issue, e.g. money) and *indivisible* if they cannot (e.g. a house).

The strategy of the agents represents how agents make decisions during negotiation and it is strongly dependent on the problem domain, the protocol, the subject and the information the agents have.

Typically, there are two types of negotiation problems: bargaining and task allocation (Vidal, 2006).

The bargaining problem has been studied in economics and has a strong theoretical support from *game-theory*. In bargaining problems, each agent tries to maximize its own preference measure. *Utility functions* are generally used to measure preferences. Agents typically use an *alternating offers* type protocol protocol when exchanging proposals and employ various strategies when computing proposals.

In a task allocation problem, agents are able to perform tasks with an associated cost and sometimes delegate tasks to other agents. An agent might choose to delegate a task either because it might not be able to perform it or because it might cost less than performing it itself. The Contract Net Protocol (CNP) (Smith, 1980) is usually studied for task contracting.

This paper presents an overview of the state-of-the-art in negotiations with a focus on complex negotiations and their challenges. The paper is structured as follows. Section 2 consists of related work in negotiation models with relatively simple preferences (e.g. linear utility functions). This section provides a background in negotiation (task allocation and bargaining), emphasizes achieved results and problems that researchers have encountered. Section 3 provides guidelines for developing negotiation models. These guidelines result from the related work and consist of the key elements researchers should take into account when developing negotiation models. Section 4 discusses complex negotiations, which are typically encountered when the utility functions are non-linear and non-monotonic. Section 5 contains conclusions and directions for future work.

2. EXISTING NEGOTIATION MODELS

For the rest of this section we will discuss the work of various researchers in the field of automated negotiation, and emphasize main ideas of their research.

(Jennings et al., 2001) discusses agent organization (cooperation and coordination) by means of automated negotiation. The authors discuss existing methods and emphasize challenges encountered by researchers. They present the three main components of negotiation: protocols, objects and decision making models (strategies). During negotiation, agents search in their private deal spaces (the space of potential outcomes) for offers to make to the other agents. Thus, negotiation is presented as a distributed search problem. The deal spaces might change as the negotiation progresses, as a result of context changes or persuasion. Agents can critique the received offers and can make counter-offers. Various negotiation techniques bring particular elements to this framework. Game-theoretic techniques can help design negotiation protocols and strategies and provide strong theoretical methods for the analysis of negotiations. Using such techniques, a negotiation is modeled as a strategic game. But they assume complete knowledge of information, an assumption which rarely holds in real-world situations. Moreover, searching for solutions is often an intractable problem (e.g. finding equilibria). The authors claim that game-theoretic techniques are much harder to use for multi-issue negotiations. Heuristic approaches attempt to overcome the disadvantages of game-theoretic approaches. Mainly, they tend to find acceptable solutions (rather than optimal solutions) in a reasonable amount of time, that are then experimentally evaluated.

Argumentation-based approaches offer the possibility for agents to exchange more information than just proposals or counter-proposals – they can offer some details about their decisions (for example, why has an agent rejected a proposal?).

In a related paper, Kraus discusses automated negotiation for various domains (Kraus, 2001). The concept of *equilibrium* (used in game-theory) is explained in the context of strategic negotiation. The author also discusses other negotiation methods such as auctions, task allocation, coalition formation and argumentation.

2.1 Task allocation

The Contract Net Protocol (CNP) (Smith, 1980) is a well known mechanism for task allocation. An agent called *manager* makes tasks available to *contractor* agents. After a selection process, the manager awards one or more tasks to a contractor, supervises the execution of tasks and processes task results. The contractor is in charge of task execution. The original specification of the CNP does not contain any strategy model. The decision making models for bidding, bid processing and task awarding phase are not specified. Sandholm proposes a study of the CNP based on marginal cost calculations (Sandholm, 1993). Agents use local utility functions for computing costs that are taken into account during the bidding and awarding phase. This model formalizes the decision steps of CNP, but the local decision models (the utility functions for example) are not specified. This type of negotiation has been applied to a vehicle routing problem, where dispatch centers negotiate in order to route vehicles efficiently and reduce costs.

Because in the CNP the manager tries to optimally select contractors, it would be nice to have some information about the performance and efficiency about contractors prior to negotiations. (Yu-Sheng et al., 2007) uses this approach in the context of task allocation among robots. They record successes and failures of contractors and use this information to select potential contractors. The manager uses credibility and relevance in order to select the contractor after the bidding phase. They also discuss commitment and permit breach of contract in their negotiations. A contractor is allowed to breach a contract by paying a penalty to the manager. The experiments prove that the negotiation mechanism is feasible.

Researchers have tried to extend the CNP according to their needs, for various negotiation problems. A survey of extensions to the CNP, as well as a theoretical evaluation of these extensions can be found in (Bozdag, 2008).

An experimental scalability analysis of the CNP is done in (Juhasz and Paul, 2002). Here the authors conclude that the performance depends on the system load and that the CNP cannot be used without a deadline.

The CNP has been studied for task allocations in the context of an insurance application (Paurobally et al., 2007). Here the authors characterize the negotiation issues with several properties (preferred value, reserved value, utility and weight) and describe various implemented strategies for generating proposals and calls for proposals.

An experimental comparison of three different methods for task allocation (sequential auctions, multi-issue MCP and mediator-based simulated annealing) has been done in (Chakraborty et al., 2006), for the scope of allocating tasks for monitoring the environment using negotiations between ground control stations and orbiting space probes.

2.2 Bargaining

One of the first models of bargaining is Rubinstein's alternating offers game with an infinite horizon (no deadlines) and complete information (Rubinstein, 1982). In this game, two players are bargaining on how to split a pie. Each player proposes a partition of a pie at each turn. The other player must either accept the offer or reject it and propose another partition. The game takes the time preferences into account as fixed bargaining costs or fixed discounting factors. Rubinstein later extended this model (Rubinstein, 1985) for incomplete information. He introduces an uncertainty element in his previous model, namely one of the two agents can be of two types: *strong* or *weak*. The *weak* player is more "impatient", i.e. it loses utility more rapidly with time than the *strong* player. The scenario presented in both models of Rubinstein (two players trying to split a pie) is characterized as negotiation about a single issue.

Another interesting model of bilateral single-issue negotiation is that of (Sandholm and Vulkan, 1999). The authors prove that in incomplete information settings with time effects (deadlines and discount factors), the bargainers prefer to wait until the earlier deadline. That is, agreement is reached at the latest possible step.

Multi-issue negotiations are more interesting to study mainly because different procedures can be used to negotiate the issues. They can be negotiated all together (*bundled*, in package) or sequentially (*issue-by-issue*). It is also more difficult to design negotiation models with multiple issues using game-theoretic techniques (Jennings et al., 2001). Using the *issue-by-issue* approach, it has been shown that the order in which the issues are negotiated (the agenda) influences the negotiation outcome (Fershtman, 1990). When the decision of what issue to negotiate next is taken during negotiation, the agenda is *endogenous*, otherwise it is *exogenous*.

An agenda-based approach to negotiations has been proposed in (Fatima et al., 2004b). It is a multi-issue, issue-by-issue negotiation model with incomplete information. Agents (two agents, a buyer and a seller) use utility functions with time discounts to evaluate received proposals. Depending on the discount factors, they can be *patient* (gain utility with time) or *impatient* (lose utility with time). Three different types of tactics are used to generate counter-offers (Faratin et al., 1998): (i) *boulware* – the agent concedes very slowly during the negotiation, but near the deadline it concedes rapidly to the reservation price; (ii) *conceder* – the agent concedes rapidly at the beginning of the negotiation to the reservation value; (iii) *linear* – the agent concedes linearly. Agents have incomplete information about their opponents in the form of two probability distributions over many values, one for the opponent's deadline and the other for the opponent's reservation price. The relationships between agent deadlines and discount factors lead to six negotiation scenarios that are then analyzed. For each one of the six scenarios the authors determine the optimal strategy, i.e. the strategy that gives the maximum expected utility. The authors also determine conditions for convergence of the optimal strategies, study the properties of the equilibrium and determine conditions under which the equilibrium is Pareto optimal. This model is extended to multiple issues by extending the information model of the agents. The analysis of negotiation with multiple issues includes a comparison of two implementations: exchange of an issue takes place immediately after the issue has been negotiated – *sequential implementation*, or after all the issues have been negotiated – *simultaneous implementation*. The negotiation agenda is endogenous, the order in which the issues are negotiated is determined by the equilibrium.

A heuristic approach is presented in (Bosse et al., 2005, 2004). The authors have created a *System for Analysis of Multi-Issue Negotiation (SAMIN)* for the purpose of testing and improving negotiations among humans. Negotiation processes are formalized using a *temporal trace language (TTL)* which makes it easy to add useful properties that the system should be tested against (e.g. Pareto monotony). Their model also uses incomplete information: one party can estimate the other's issue weights. This guessing is based on the assumption that an agent concedes more on a less important issue. Time effects are not present in the model. Also, the model is not theoretically analyzed.

The authors of (Kraus and Wilkenfeld, 1993) present a negotiation model with applications to an international crisis. They take time into account, describe the equilibrium and show that, in theory, if an agreement can be reached, it will be reached in the first or the second round of negotiation.

In (Fatima et al., 2007), the authors study bilateral multi-issue negotiation with indivisible issues, incomplete information and time constraints in the form of deadlines and discount factors. For indivisible issues, finding the equilibrium is an NP-hard problem (similar to the 0-1 knapsack problem). Thus, the authors find approximate strategies and show that they form approximate equilibria. The study is extended to online negotiation, when issues become available later in the negotiation and the agents are unsure about their preferences. This situation also leads to an approximate equilibrium.

In (Fatima et al., 2006a,b), the same authors study the negotiation procedure. They compare the *package deal* (bundled issues, negotiated together), *simultaneous* (negotiated at the same time, but independent of each other) and *sequential* (independently and one after another) procedures in bilateral, multi-issue negotiations with complete and incomplete information. They found that for the package deal procedure with complete information, the negotiation outcome can be computed in linear time by solving a *fractional knapsack problem.* The difference between the two papers is that they focus on different types of uncertainties, but the key result of both is that the optimal procedure is the package deal.

The same researchers also study optimal agendas. In (Fatima et al., 2004a) they study the optimal agenda for bilateral, 2-issue negotiations with incomplete information. Four negotiation scenarios are defined, depending on the zone of agreement for the two issues. Then they identify the optimal agenda for each one of the two procedures. In (Fatima et al., 2003), the authors decompose the negotiation issues into a number of stages. At each stage, the issues to be negotiated are determined exogenously and the order is determined endogenously. Therefore, the authors study an agenda that is both exogenous and endogenous. They determine negotiation scenarios that give agents higher utilities when the issues are negotiated in stages, but the agents are not able to identify the scenarios, since they do not have complete information. A mediator is used to help the agents identify those scenarios.

3. DEVELOPING NEGOTIATION MODELS

There are lots of things to be taken into consideration when designing negotiation models. Most often the protocols, strategies and negotiation objects are domain dependent. That means that the negotiation designer first analyzes the problem to be resolved (the environment, the available resources) and then chooses the most appropriate negotiation protocol, strategies and subject.

When there are only two agents that will interact, a model for bilateral negotiation will be used, like (Fatima et al., 2004b; Sandholm and Vulkan, 1999). When more that two agents will interact, multilateral negotiation models will be used. Multilateral negotiation models can be divided into two categories: (i) one-to-many negotiations in which one of the agents has a special role (e.g. *coordinator*, *beneficiary* in service contracting, task allocation; *auctioneer* in auctions) and (ii) many-to-many negotiations, where all the agents have the same privileges and negotiate for the partitions of an object.

The negotiation subject is one of the most important things to take into consideration, as it is the element that drives the negotiation (e.g. a pie that will be split (Rubinstein, 1982), a house that will be sold) and gives the outcome structure. If there is only one atomic item/issue [1] that will be negotiated, the subject is said to be *single-issue.* If the negotiation subject can be split into multiple atomic items/issues, the subject is said to be *multi-issue.* Single-issue negotiations are easier to study using game-theoretic techniques (Jennings et al., 2001) because the preferences over a single issue can be represented easier. Humans usually prefer one issue to another, and one value of an issue to another. In economics, these preferences are represented using *indifference curves* (Varian, 2005), but in this form they cannot be compared. Therefore, the concept of utility has been introduced.

Utility functions are largely used to represent and compare preferences in multi-agent negotiation. In negotiation, a utility function $u(x) : \mathcal{X} \to \mathcal{R}$ maps a set of outcomes to real numbers. It shows how much an agent prefers an outcome using a real number. Usually this function is bounded to a positive interval of the real numbers, e.g. $[0, 1]$. A utility of 0 denotes no valuation. Usually 0 utility is typically assigned also to the conflict deal (i.e. no deal), this having the effect that a rational agent does not want to reach a conflict but it tries to reach an agreement and gain an outcome. In multi-agent negotiations, agents use utility to represent their preferences and evaluate outcomes and try to maximize their utilities (von Neumann and Morgenstern, 1947) (Von Neumann-Morgenstern utility maximizers). If the agents value the issues independently (issues do not depend on one another), a weighted-sum utility function can be used to evaluate outcomes.

If the agent considers dependencies between issues, the utility function can take more complex forms, depending on how these interdependencies are modeled (Hemaissia et al., 2007; Ito et al., 2007; Klein et al., 2003).

The interaction protocols define the conditions under which agents exchange messages (offers, counter-offers). For bargaining problems the most encountered protocol is the alternating offers (Osborne and Rubinstein, 1994) protocol and variations. Using this protocol, agents start by offering deals that give maximum utility and then continue making concessions to the other agent using different tactics (for an example of tactics, including time-dependent, see (Faratin et al., 1998)) until an agreement or the conflict deal is reached. A typical flow of this protocol can be seen in Figure 1, where two agents a and b reach an agreement after their utilities dropped from maxima to the scoring of the agreement outcome. After receiving an offer, an agent evaluates the offer using its utility function and decides whether to accept the offer or make a

[1] Please note that atomic here means that the negotiator does not have any interest to divide the issue, i.e. he does not have different preferences for smaller parts of the issue

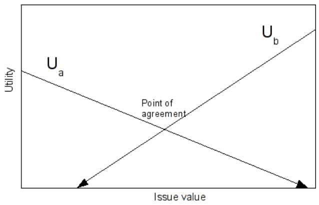

Fig. 1. Bargaining with linear functions

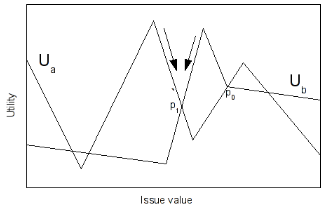

Fig. 2. Bargaining with non-monotonic functions

counter-offer. The process of evaluating the offers and making concessions are influenced by the preferences the agents have about the negotiation subject, the environment (e.g. time) and the other agents. For example, time can influence an agent's utility function and how it makes proposals. The agent can evaluate more the negotiation subject as time passes (gains utility with time), or the negotiation subject might loose importance as time passes (looses utility with time). These time constraints are usually modeled using *deadlines* and/or *discount factors* (Rubinstein, 1982, 1985; Fershtman, 1990; Fatima et al., 2004b; Sandholm and Vulkan, 1999; Faratin et al., 1998; Hemaissia et al., 2007). Bargaining problems (taking the model from game-theory) have been divided into: negotiation with complete information and negotiation with incomplete information (von Neumann and Morgenstern, 1947). Information completeness means agents completely know their and the other agents' preferences. Information incompleteness refers to the lack of knowledge about some information, either about themselves or about the other agents and it is the most encountered situation in real world. However, agents can still try to model these parameters, often by using probability distributions (see for example the model of (Fatima et al., 2004b)). An interesting fact has been pointed out in the literature (e.g. (Sandholm and Vulkan, 1999)): an agent reveals information about itself when making offers or counter-offers. Negotiation with incomplete preferences is harder as compared to negotiation with complete information. With complete information means that at least theoretically, if the point of agreement exists then it can be computed (even if sometimes this computation is too expensive). Note than knowing if a solution exists is a different thing from knowing the actual solution.

Just having a mechanism is not sufficient. It must be stable, i.e. the agents must not have the incentive to deviate from their strategies. Game-theory provides the concept of equilibria (Osborne and Rubinstein, 1994). Various types of equilibria are available, depending on the situation. Sometimes equilibria might be hard to compute, even though it can be proven that they exist. If there are equilibria, the optimal one (i.e. the one that gives agents maximum utility) is desirable. But in some cases the existence of equilibria is not easy to prove and the negotiation model is evaluated through experiments.

4. ADVANCED NEGOTIATIONS

Bargainers typically start by proposing the deal that gives them maximum utility and then make concessions following a certain strategy. They seek to maximize their utility functions. Figure 1 shows a simplified process of this kind. This hill-climbing approach works perfectly as long as the utility functions are monotonic. As monotonic functions have only one local optimum (which is also the global optimum), the hill-climbing method stops at the global optimum. But this situation changes when the utility functions are not monotonic, as the hill-climbing method is not able to get past local optima.

For example, the non-monotonic utility functions of two agents, a and b, are represented in Figure 2 [2]. It can be observed that the utility function of agent a has 3 local maxima (one of which is the global maximum). The utility function of agent b has two local maxima.

When they are brought together in a bargaining process (Figure 2), the agents use the hill-climbing method and reach an agreement point, p_1, but the outcome is not optimal. Instead, they should have reached the agreement in p_0 in order to get the maximum outcome.

The utility functions become non-monotonic in complex situations (for example, when agents negotiate about interdependent issues). In order to use hill-climbing, one must design methods for detection of local optima and take actions accordingly. One possibility to do that is to explore the entire deal space and mark the points that give local optima. But if the number of issues is very large, the deal space becomes too complex to be easily explored. Researchers have tried to overcome this problem. All the known proposed solutions are heuristic. Many of them make use of the simulated annealing optimization algorithm (Russell and Norvig, 2003). This algorithm, might accept, with a certain probability, solutions that might not be better than the current solution. It has been shown in practice that simulated annealing can get past local optima and reach global optimum in many situations, unlike hill-climbing, which gets stuck in local optima.

The problem of complex negotiations is described in (Klein et al., 2003). The authors develop mediated bilateral negotiations about interdependent boolean issues. The preference model is simple, but it defines a very large space that cannot be easily explored. They study the outcome of negotiations between two types of agents: hill-climbing (accepting only contracts better than the last accepted contracts) and annealing (can accept worse contracts with a certain probability). When pairing hill-climbers only, the outcome is poor (they get stuck in local optima), while when pairing annealers, the outcome is good (as they can get past local optima). However, when pairing one hill-climber with one annealer, the hill-climber does very well because the annealer, mostly at the beginning, accepts even non-beneficial contracts. By improving

[2] The functions are depending on one issue only for ease of representation, as multiple issues require multiple dimensions

their model, they come up with a final solution in which they put the annealing part inside the mediator and extend the negotiation protocol to allow agents to vote the contracts proposed by the mediator. The solution works very well and the work is very valuable.

(Ito et al., 2007) studies a model of multi-issue negotiation with nonlinear and non-monotonic utilities and interdependent issues. The interdependencies between issues are represented with constraints. As it is very hard to completely explore the deal space in this situation, the agents first take random contract samples from the deal space, then adjust the samples to find local optima using simulated annealing and then make bids and submit them to a mediator, which computes bid intersection and determines the outcome. The model limits the number of bids the agent can make to make computations finish in a reasonable amount of time. The performance drops exponentially with the number of issues. A comparison with a hill-climbing method is provided, showing that the proposed protocol performs better, especially in complex deal spaces with local optima, where the hill-climbing method blocks. The model is somehow improved in (Hattori et al., 2007) by using a three-stage protocol which reduces complexity, but the disadvantages of the previous model, namely the presence of a mediator, a bid limit per agent and comparison with hill-climbing only, still remain. The presence of a mediator is disadvantageous because agents have to give their private preferences to a third-party agent. The model has not been analyzed using game-theoretic tools. In a more recent work (Fujita et al., 2008) the model is improved in such a way that it becomes scalable by using a protocol based on several representative agents.

Because the previously described negotiation models using constraints do not perform well when the constraints are narrow, they have been improved in a series of works (Marsa-Maestre et al., 2009a,b, 2010). Instead of taking contract samples from the deal space, the authors take constraint samples and instead of relying solely on the utility of a contract, they make use of a quality factor of a contract or a constraint, which takes into account the width of the contract. They also manage to get socially optimal outcomes by changing the computation in the mediator (Marsa-Maestre et al., 2010).

Another negotiation model with simulated annealing is (Kardan and Janzadeh, 2008). Agents negotiate until a deadline. They propose contracts by mutating the last accepted contract and accept using a simulated annealing method. This work does not involve a mediator.

There are, however, other methods that do not use simulated annealing and focus on the approximation of the utilities, or try to reduce the complexity of the problem.

Results of finding the optimal procedure in case of nonlinear utility functions are shown in (Fatima et al., 2009). Computing equilibrium for the package deal procedure if utility functions are nonlinear is hard. The authors compare the package deal procedure (when issues are bundled together) for a linear approximation of the utility functions with the simultaneous procedure (when issues are negotiated independently of each other, in parallel), resulting that an equilibrium can be computed in polynomial time for both procedures and that the package deal procedure leads to Pareto optimality.

Another model that works with interdependent issues, (Hemaissia et al., 2007), considers more complex interdependencies with the help of the Choquet integral. They study a cooperative protocol and show that the protocol has subgame perfect equilibria.

By modeling agent preferences using utility graphs and trying to decompose them, the complexity of the problem can be exponentially reduced (although it remains exponential) (Robu et al., 2005). The outcomes are close to maximum efficiency.

5. CONCLUSIONS AND FUTURE WORK

Negotiation with non-linear utility functions is a complex problem. Such situations are common in the real world. So far, exploration of the deal space using simulated annealing and with the help of a mediator has been the typical method used for solving this problem. More research should be done to find better approaches to this problem, better methods for exploration of complex deal spaces and for reaching agreements in a reasonable amount of time. Designing such methods might require a complete design and analysis of agent preferences. The approaches might be extended with methods for representing and learning the other agents' preferences (mainly the interdependencies between issues) and finding optimal strategies. If the solutions cannot be theoretically verified various experiments will be performed to measure their performance, efficiency and other characteristics.

REFERENCES

Bosse, T., Jonker, C.M., Meij, L., Robu, V., and Treur, J. (2005). A system for analysis of multiissue negotiation. In *Software Agent-Based Applications, Platforms and Development Kits*, 253–280. Birkhaeuser Publishing Company.

Bosse, T., Jonker, C.M., and Treur, J. (2004). Experiments in human multi-issue negotiation: Analysis and support. *Autonomous Agents and Multiagent Systems, International Joint Conference on*, 2, 671–678.

Bozdag, E. (2008). A survey of extensions to the contract net protocol. Technical report, CiteSeerX - Scientific Literature Digital Library and Search Engine [http://citeseerx.ist.psu.edu/oai2] (United States).

Bădică, C. and Scafeş, M. (2009). One-to-many monotonic concession negotiation protocol. In *Proc.of the* 4^{th} *Balkan Conference in Informatics, BCI'2009*, 8–13. IEEE Computer Society.

Chakraborty, D., Saha, S., Sen, S., and Clement, B. (2006). Negotiating monitoring task allocation for orbiters. In S. Chaudhuri, S.R. Das, H.S. Paul, and S. Tirthapura (eds.), *ICDCN*, volume 4308 of *Lecture Notes in Computer Science*, 282–287. Springer.

Endriss, U. (2006). Monotonic concession protocols for multilateral negotiation. In *AAMAS '06: Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*, 392–399. ACM, New York, NY, USA.

Faratin, P., Sierra, C., and Jennings, N.R. (1998). Negotiation decision functions for autonomous agents. *Int. Journal of Robotics and Autonomous Systems*, 24(3 - 4), 159–182.

Fatima, S., Wooldridge, M., and Jennings, N.R. (2003). Optimal agendas for multi-issue negotiation. In *2nd International Joint Conference on Autonomous Agents and Multi-Agent Systems*, 129–136.

Fatima, S.S., Wooldridge, M., and Jennings, N. (2009). An analysis of feasible solutions for multi-issue negotiation involving non-linear utility functions. In *Proc. 8th Int. Conf on Autonomous Agents and Multi-Agent Systems*, 1041–1048.

Fatima, S., Wooldridge, M., and Jennings, N.R. (2004a). Optimal negotiation of multiple issues in incomplete information settings. In *AAMAS '04: Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems*, 1080–1087. IEEE Computer Society, Washington, DC, USA.

Fatima, S.S., Wooldridge, M., and Jennings, N.R. (2004b). An agenda-based framework for multi-issue negotiation. *Artif. Intell.*, 152(1), 1–45.

Fatima, S., Wooldridge, M., and Jennings, N.R. (2006a). Multi-issue negotiation with deadlines. *Journal of Artificial Intelligence Research*, 27, 381–417.

Fatima, S., Wooldridge, M., and Jennings, N.R. (2006b). On efficient procedures for multi-issue negotiation. In *8th Int. Workshop on Agent-Mediated Electronic Commerce*, 71–84.

Fatima, S., Wooldridge, M., and Jennings, N.R. (2007). Approximate and online multi-issue negotiation. In *6th International Joint Conference on Autonomous Agents and Multi-agent Systems.*, 947–954.

Fershtman, C. (1990). The importance of the agenda in bargaining. *Games and Economic Behavior*, 2(3), 224–238.

Fujita, K., Ito, T., and Klein, M. (2008). A preliminary result on a representative-based multi-round protocol for multi-issue negotiations (short paper). In Padgham, Parkes, Müller, and Parsons (eds.), *Proc. of 7th Int. Conf on Autonomous Agents and Multi-Agent Systems (AAMAS 2008)*, 1573–1576.

Hattori, H., Klein, M., and Ito, T. (2007). A multi-phase protocol for negotiation with interdependent issues. In *IAT '07: Proceedings of the 2007 IEEE/WIC/ACM International Conference on Intelligent Agent Technology*, 153–159. IEEE Computer Society, Washington, DC, USA.

Hemaissia, M., El Fallah Seghrouchni, A., Labreuche, C., and Mattioli, J. (2007). A multilateral multi-issue negotiation protocol. In *AAMAS '07: Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*, 1–8. ACM, New York, NY, USA.

Ito, T., Hattori, H., and Klein, M. (2007). Multi-issue negotiation protocol for agents: Exploring nonlinear utility spaces. In M.M. Veloso (ed.), *IJCAI*, 1347–1352.

Jennings, N.R., Faratin, P., Lomuscio, A.R., Parsons, S., Sierra, C., and Wooldridge, M. (2001). Automated negotiation: Prospects, methods and challenges. *International Journal of Group Decision and Negotiation*, 10(2), 199–215.

Juhasz, Z. and Paul, P. (2002). Scalability analysis of the contract net protocol. In *CCGRID '02: Proceedings of the 2nd IEEE/ACM International Symposium on Cluster Computing and the Grid*, 346. IEEE Computer Society, Washington, DC, USA.

Kardan, A. and Janzadeh, H. (2008). A multi-issue negotiation mechanism with interdependent negotiation issues. In *ICDS '08: Proceedings of the Second International Conference on Digital Society*, 55–59. IEEE Computer Society, Washington, DC, USA.

Klein, M., Faratin, P., Sayama, H., and Bar-Yam, Y. (2003). Negotiating complex contracts. *Group Decision and Negotiation*, 12, 111–125.

Kraus, S. (2001). Automated negotiation and decision making in multiagent environments. *Mutli-agents systems and applications*, 150–172.

Kraus, S. and Wilkenfeld, J. (1993). A strategic negotiations model with applications to an international crisis. *IEEE Transactions on Systems, Man, and Cybernetics*, 23.

Marsa-Maestre, I., Lopez-Carmona, M.A., Velasco, J.R., and de la Hoz, E. (2009a). Effective bidding and deal identification for negotiations in highly nonlinear scenarios. In Decker, Sichman, Sierra, and Castelfranchi (eds.), *Proc. of 8th Int. Conf on Autonomous Agents and Multi-Agent Systems (AAMAS 2009)*, 1057–1064.

Marsa-Maestre, I., Lopez-Carmona, M.A., Velasco, J.R., and de la Hoz, E. (2010). Avoiding the prisoner's dilemma in auction-based negotiations for highly rugged utility spaces. In van Der Hoek, Kaminka, Lespérance, Luck, and Sen (eds.), *Proc. of 9th Int. Conf on Autonomous Agents and Multi-Agent Systems (AAMAS 2010)*, 425–432.

Marsa-Maestre, I., Lopez-Carmona, M.A., Velasco, J.R., Ito, T., Klein, M., and Fujita, K. (2009b). Balancing utility and deal probability for auction-bases negotiations in highly nonlinear utility spaces. In *Proc. of the 21st Int. Joint Conf. on Artificial Intelligence (IJCAI-09)*, 214–219.

Osborne, M. and Rubinstein, A. (1994). *A Course in Game Theory.* MIT Press, Cambridge, MA.

Paurobally, S., Tamma, V., and Wooldrdige, M. (2007). A framework for web service negotiation. *ACM Trans. Auton. Adapt. Syst.*, 2(4), 14.

Robu, V., Somefun, D.J.A., and La Poutré, J.A. (2005). Modeling complex multi-issue negotiations using utility graphs. In *AAMAS '05: Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*, 280–287. ACM, New York, NY, USA.

Rosenschein, J.S. and Zlotkin, G. (1994). *Rules of Encounter: Designing Conventions for Automated Negotiation Among Computers.* MIT Press, Cambridge, Massachusetts.

Rubinstein, A. (1982). Perfect equilibrium in a bargaining model. *Econometrica*, 50(1), 97–109.

Rubinstein, A. (1985). A bargaining model with incomplete information about time preferences. *Econometrica*, 53(5), 1151–72.

Russell, S.J. and Norvig, P. (2003). *Artificial Intelligence - A Modern Approach.* Prentice Hall.

Sandholm, T.W. (1993). An implementation of the contract net protocol based on marginal cost calculations. In *Proceedings of the 12th Interna-*

tional Workshop on Distributed Artificial Intelligence, 295–308. Hidden Valley, Pennsylvania.

Sandholm, T. and Vulkan, N. (1999). Bargaining with deadlines. In *AAAI '99/IAAI '99: Proceedings of the sixteenth national conference on Artificial intelligence and the eleventh Innovative applications of artificial intelligence conference innovative applications of artificial intelligence*, 44–51. American Association for Artificial Intelligence, Menlo Park, CA, USA.

Smith, R.G. (1980). The contract net protocol: High-level communication and control in a distributed problem solver. *IEEE Trans. Comput.*, 29(12), 1104–1113.

Varian, H.R. (2005). *Intermediate Microeconomics: A Modern Approach*. W. W. Norton & Company, 7th edition.

Vidal, J.M. (2006). *Fundamentals of Multiagent Systems: Using NetLogo Models*. Unpublished.

von Neumann, J. and Morgenstern, O. (1947). *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ.

Yu-Sheng, H., Mo-Gen, X., and Pu-Cheng, Z. (2007). Extended contract net protocol for multi-robot dynamic task allocation. *Information Technology Journal*, 6(5), 733–738.