[2] Andrew C. Revkin: "Climate and the Web: Electronic Democracy on Steroids", http://dotearth.blogs.nytimes.com/2008/04/25/dclimate-and-the-web-electronic-democracy-on-steroids/ (2008).

[3] T. Bosse and C. M. Jonker: "Human vs. computer behaviour in multi-issue negotiation", In Proceedings of 1st International Workshop on Rational, Robust, and Secure Negotiations in Multi-Agent Systems (RRS-2005), pp. 11–24 (2005).

[4] P. Faratin, C. Sierra and N. R. Jennings: "Using similarity criteria to make issue trade-offs in automated negotiations", Artificial Intelligence, pp. 142:205–237 (2002).

[5] S. Fatima, M. Wooldridge and N. R. Jennings: "Optimal negotiation of multiple issues in incomplete information settings", In Proceedings of Third International Joint Conference on Autonomous Agent and Multi-Agent Systems (AAMAS-2004), pp. 1080–1087 (2004).

[6] T. Ito, H. Hattori and M. Klein: "Multi-issue negotiation protocol for agents : Exploring nonlinear utility spaces", In Proceedings of 20th International Joint Conference on Artificial Intelligence (IJCAI-2007), pp. 1347–1352 (2007).

[7] K. J. Arrow: "Social Choice and Individual Values", Yale Univ. Press (1970).

[8] K. Fujita, T. Ito, H. Hattori and M. Klein: "An approach to implementing a threshold adjusting mechanism in very complex negotiations: A preliminary result", In Proceedings of The 2nd International Conference on Knowledge, Information and Creativity Support Systems (KICSS-2007) (2007).

[9] S. J. Russell and P. Norvig: "Artificial Intelligence : A Modern Approach", Prentice Hall (2002).

[10] G. Werner, S. Rolf and S. Bernd: "An experimental analysis of ultimatum bargaining", Journal of Economic Behavior and Organization, **3**, pp. 367–388 (1982).

[11] R. J. Lin and S. T. Chou: "Bilateral multi-issue negotiations in a dynamic environment", In Proceedings of the International Workshop on Agent-mediated Electronic Commerce (AMEC2003) (2003).

[12] M. Barbuceanu and W.-K. Lo: "Multi-attribute utility theoretic negotiation for electronic commerce", In Proceedings of the International Workshop on Agent-mediated Electronic Commerce (AMEC2000), pp. 15–30 (2000).

[13] X. Luo, N. R. Jennings, N. Shadbolt, H. Leung and J. H. Lee: "A fuzzy constraint based model for bilateral, multi-issue negotiations in semi-competitive environments", Artificial Intelligence, **148**, pp. 53–102 (2003).

[14] M. Klein, P. Faratin, H. Sayama and Y. Bar-Yam: "Negotiating complex contracts", Group Decision and Negotiation, **12**, 2, pp. 58–73 (2003).

[15] G. Lai, C. Li and K. Sycara: "A general model for pareto optimal multi-attribute negotiations", In Proceedings of The 2nd International Workshop on Rational, Robust, and Secure Negotiations in Multi-Agent Systems (RRS-2006) (2006).

[16] G. Lai, K. Sycara and C. Li: "A decentralized model for multi-attribute negotiations with incomplete information and general utility functions", In Proceedings of The 2nd International Workshop on Rational, Robust, and Secure Negotiations in Multi-Agent Systems (RRS-2006) (2006).

[17] V. Robu and H. L. Poutre: "Retrieving the structure of utility graphs used in multi-item negotiation through collaborative filtering of aggregate buyer preferences", In Proceedings of The 2nd International Workshop on Rational, Robust, and Secure Negotiations in Multi-Agent Systems (RRS-2006) (2006).

[18] V. Robu, D. J. A. Somefun and J. L. Poutre: "Modeling complex multi-issue negotiations using utility graphs", In Proceedings of the 4th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2005) (2005).

[19] S. S. Fatima, M. Wooldridge and N. R. Jennings: "Approximate and online multi-issue negotiation", In Proceedings of th 6th Inernational Joint Conference on Autonomous Agents and Multi-agent Systems (AAMAS-2007), pp. 947–954 (2007).

[20] J. Shew and K. Larson: "The blind leading the blind: A third-party model for bilateral multi-issue negotiations under incomplete information", In Proceedings of The 1st International Workshop on Agent-based Complex Automated Negotiations (ACAN-2008) (2008).

[21] K. Hindriks, C. Jonker and D. Tykhonov: "Avoiding approximation errors in multi-issue negotiation with issue dependencies", In Proceedings of The 1st International Workshop on Agent-based Complex Automated Negotiations (ACAN-2008) (2008).

Auton Agent Multi-Agent Syst
DOI 10.1007/s10458-010-9159-9

Addressing stability issues in mediated complex contract negotiations for constraint-based, non-monotonic utility spaces

Miguel A. Lopez-Carmona · Ivan Marsa-Maestre · Mark Klein · Takayuki Ito

Abstract Negotiating contracts with multiple interdependent issues may yield non- monotonic, highly uncorrelated preference spaces for the participating agents. These scenarios are specially challenging because the complexity of the agents' utility functions makes traditional negotiation mechanisms not applicable. There is a number of recent research lines addressing complex negotiations in uncorrelated utility spaces. However, most of them focus on overcoming the problems imposed by the complexity of the scenario, without analyzing

M. A. Lopez-Carmona and I. Marsa-Maestre are visiting scholars at Massachusetts Institute of Technology (MIT) from Universidad de Alcala.

M. A. Lopez-Carmona (✉) · I. Marsa-Maestre
Computer Engineering Department, Edificio Politecnico, Universidad de Alcala,
Ctra. N-II, Km 33.6, 28805, Alcala de Henares, Madrid, Spain
e-mail: miguelangel.lopez@uah.es

I. Marsa-Maestre
e-mail: ivan.marsa@uah.es

M. A. Lopez-Carmona · I. Marsa-Maestre · M. Klein
Center for Collective Intelligence, MIT Sloan School of Management, Massachusetts Institute of Technology (MIT), 5 Cambridge Center, NE25-749A, Cambridge, MA 02142, USA

M. Klein
e-mail: m_klein@mit.edu

T. Ito
Department of Computer Science and Engineering, Graduate School of Engineering,
Nagoya Institute of Technology, Gokiso, Showa-ku, Nagoya 466-8555, Japan
e-mail: ito.takayuki@nitech.ac.jp

T. Ito
Todai Policy Alternatives Research Institute, University of Tokyo, 7-3-1 Hongo,
Bunkyo-ku, Tokyo 113-0033, Japan

T. Ito
Precursory Research for Embryonic Science and Technology (PRESTO),
Japan Science and Technology Agency (JST), Tokyo, Japan

Published online: 04 December 2010

the potential consequences of the strategic behavior of the negotiating agents in the models they propose. Analyzing the dynamics of the negotiation process when agents with different strategies interact is necessary to apply these models to real, competitive environments. Specially problematic are high *price of anarchy* situations, which imply that individual rationality drives the agents towards strategies which yield low individual and social welfares. In scenarios involving highly uncorrelated utility spaces, "low social welfare" usually means that the negotiations fail, and therefore high price of anarchy situations should be avoided in the negotiation mechanisms. In our previous work, we proposed an auction-based negotiation model designed for negotiations about complex contracts when highly uncorrelated, constraint-based utility spaces are involved. This paper performs a strategy analysis of this model, revealing that the approach raises stability concerns, leading to situations with a high (or even infinite) price of anarchy. In addition, a set of techniques to solve this problem are proposed, and an experimental evaluation is performed to validate the adequacy of the proposed approaches to improve the strategic stability of the negotiation process. Finally, incentive-compatibility of the model is studied.

Keywords Automated multi-issue negotiation · Complex utility spaces · Strategy analysis

1 Introduction

Automated negotiation provides an important mechanism to reach agreements among distributed decision makers [4,42,43,71]. It has been extensively studied from the perspective of e-commerce [23,25,49,76], though it can be seen from a more general perspective as a paradigm to solve coordination and cooperation problems in complex systems [38,32], providing a mechanism for autonomous agents to reach agreements on, e.g., task allocation, resource sharing, or surplus division [15,37].

A variety of negotiation models have been proposed according to the many different parameters which may characterize a negotiation scenario [6,43]. We briefly review the key concepts about multi-attribute negotiation and the most relevant works in the field in Sect. 2.1. In the last years, there has been an increasing interest in complex negotiations [39]. Complexity of a negotiation scenario may depend on several factors, like the cardinality of the solution space, the number of negotiating agents, the number of issues under negotiation, the degree of interdependency between the issues, and structural properties of the preference landscape of the different agents, like ruggedness, modality or correlation length [80]. Specially challenging are those scenarios involving high cardinality solution spaces, since they tend to make exhaustive search in the solution space highly inefficient, and those involving highly rugged or highly uncorrelated utility spaces, since traditional negotiation approaches (mostly intended for linear or quasi-concave utility functions) cannot be applied to these scenarios. We briefly discuss utility space complexity and the techniques used to measure it in Sect. 2.2.

We can find some successful research works in the literature addressing negotiation in nonlinear utility spaces. [39] presented, as far as we are aware, the first negotiation protocol specific for complex preference spaces, based on using simulated annealing to progressively enhance an agreement between two agents. In [26], a different approach is taken, reducing the complexity of the agent's preference space by using approximations of the agents' utility functions where issue interdependency has been removed. [17] do not study the inherent complexity of agent preference spaces, but the complexity introduced in a negotiation when agent preferences change over time. We comprehensively review these and other related works in Sect. 2.3.

In our previous work [54], we proposed a mediated, auction-based protocol for nonlinear utility spaces generated using weighted constraints, such as the ones we may encounter when negotiating complex contracts with multiple, interdependent clauses [30]. We also proposed a set of decision mechanisms to generate bids at the negotiating agents and to identify feasible deals at the mediator once the bids from the negotiating agents have been received [54]. We briefly summarize the approach in Sect. 3. Experiments showed that these approaches achieve high effectiveness (measured as high optimality rates and low failure rates for the negotiations) in moderately rugged utility spaces.

In [55], we extended this work to address highly-rugged utility spaces. We proposed the use of a technique to balance utility and deal probability in the negotiation process, which we called *quality factor*. This quality factor is used to bias bid generation and deal identification taking into account the agents' attitudes (e.g. risk attitude, selfishness, willingness to cooperate). From the mechanisms we proposed to take into account quality factor in the negotiations, the most successful ones are detailed in Sect. 3.4. The experiments showed that this balance between utility and deal probability greatly improves the effectiveness of the negotiation in highly-rugged utility spaces.

However, the proposed approach draws several concerns. Though the quality factor is supposed to be able to model agents' attitudes, our previous experiments limited these attitudes to a somewhat "cooperative" environment, where all agents have the same, neutral attitude. In a real, competitive environment, we expect to have agents with different attitudes interacting. This raises the problem of agent strategic behavior, which is introduced in Sect. 4. What happens when risk averse agents interact with risk willing agents? Is there a an individually optimal strategy? If so, does this individually optimal strategy lead to satisfying solutions, or is the approach prone to situations where individual rationality lead to solutions of low social value? Furthermore, since the complexity of the utility spaces of the agents may also vary, it seems logical to think that agent strategies should vary accordingly. In this paper, we intend to address these questions in the following ways:

- We perform a strategy analysis of the auction-based protocol for constraint-based utility spaces. This analysis allows us to determine the individually optimal strategy and the socially optimal strategy for different utility space complexity levels. From the results of the analysis we conclude that the auction-based protocol, as described in [55], has stability problems, leading to situations resulting in high expected price of anarchy (Sect. 4).
- We propose a set of mechanisms intended to improve protocol stability. These approaches are based on decoupling the agent's strategies from the deal identification process, by applying different techniques on the mediator after the agents have sent their bids (Sect. 5).
- We separately study a specific stability concern, incentive compatibility, related to the possibility of agents manipulating the protocol by means of insincere revelation of information (Sect. 6).

For each contribution, an experimental evaluation has been performed to validate our hypothesis and evaluate its effect. The experimental settings are described in Sects. 4.2, 5.2, 6.2 and 6.3, along with the discussion of the results obtained. Finally, the last section summarizes our conclusions and sheds light on some future research.

2 Complex negotiation scenarios

In the last years, there has been an increasing interest in complex negotiation scenarios, where agents negotiate about multiple, interdependent issues [39]. These scenarios are specially challenging, since issue interdependency yields nonlinear utility spaces, which make classic negotiation approaches not applicable [30]. In this section we first briefly review existing research on multi-attribute negotiation and outline the key components of any negotiation model. Then we discuss the most relevant works so far on the field of agent-based complex automated negotiations. Finally, some of the issues raised by complex negotiation scenarios, which are directly relevant to our research, are described.

2.1 Multi-attribute negotiation

Multi-attribute negotiation may be seen as an interaction between two or more agents with the goal of reaching an agreement about a range of issues which usually involves solving a conflict of interests between the agents. This kind of interaction has been widely studied in different research areas, such as game theory [71], distributed artificial intelligence [13] and economics [67]. Using a notation similar to that used in [71] and [88], we can formally define a *multi-attribute negotiation domain* as a tuple

$$\langle X, D, Ag, U \rangle$$

where

- $X = \{x_i | i = 1, \dots, n\}$ is a finite set of variables, called attributes or issues;
- $D = \{d_i | i = 1, \dots, n\}$ is a finite set of domains, such that each domain d_i represents the feasible values of the variable x_i;
- $Ag = \{1, \dots, m\}$ is the set of negotiating agents, also assumed finite;
- $U = \{U^j | j = 1, \dots, m\}$, where $U^j : D \rightarrow \mathbb{R}$ represents the preference structure or utility for agent j.

Multi-attribute negotiation is seen as an important challenge for the multi-agent system research community [43], and there is a great variety of negotiation models and protocols intended to address different parts of this challenge. These models may be classified according to different criteria [6], such as their structure, the dynamics of the negotiation process, or the different constraints (e.g. deadlines, information availability...). According to the theoretical foundations of the negotiation models, we can find approaches based on game theory, heuristic approaches and argumentation-based approaches. Game theory approaches aim to find optimal solutions analytically, analyzing equilibrium conditions [59]. These models are mathematically sound and elegant, but their pratical use in some negotiation scenarios is somewhat restricted due to the assumptions usually made: unlimited computation and memory resources, perfect rationality and complete information. In *heuristic* approaches, however, these assumptions are relaxed, and participants attempt to find an "approximately-optimal" under bound rationality using heuristic search and evaluation methods [12–14,20,31,39,44,70]. In *argumentation* based negotiation, agents are given the ability to reason their positions, including a meta-information level which allows them to use promises, rewards, threats and other incentives [66].

Regardless of the theoretical approach involved, different authors agree that there are three key components in a negotiation model [16,33,41]:

- An *interaction protocol*, which defines the rules of encounter among the negotiating agents, including what kind of offer exchange is allowed and what kind of deals may be reached and how they are established.
- The *preference* sets of the different agents, which allow them to assess the different solutions in terms of gain or utility and to compare them.

– A set of *decision mechanisms* and *strategies*, which govern agents' decision making, allowing them to determine which shall be their next action for a given negotiation state.

2.1.1 Interaction protocols for negotiation

The most-widespread interaction protocol for negotiation is based on the exchange of offers and counter offers, which are expressed as an assignation of values to the different attributes. This kind of negotiation protocols are known as positional bargaining. In argumentation based negotiation, however, this exchange of offers also includes meta-information, in order to allow reasoning about the positions of the different agents. A particular protocol family for multi-lateral negotiations are *auction-based protocols*, where negotiating agents send their offers (also called *bids*) to a mediator, which then decides the winning deal [77]. Auction-based protocols allow to efficiently deal with one-to-many and many-to-many negotiations. Another important division regarding interaction protocols is between *one-shot* protocols and *iterative* protocols. In one-shot protocols, there is a single interaction step between the agents [59]. In iterative protocols, on the other hand, agents have the opportunity to refine their positions in successive protocol iterations [62].

2.1.2 Preference sets, utility functions and the use of constraints

From the decision theory perspective, preferences express the absolute or relative satisfaction for an individual about a particular choice among different options [36]. [7] classify agent preference structures in four broad families: binary, ordinal, cardinal and fuzzy preference structures. Among these families, cardinal preference structures are probably the most widely used in complex negotiations. In particular, it is usual to define agent preferences by means of utility functions.

Formally, for a given multi-attribute domain $\langle X, D, Ag, U\rangle$, the *utility function* for each agent $j \in Ag$ is defined as

$$U^j : D \to \mathbb{R},$$

assigning to each possible combination of values in X or *deal* $s = \{s_i | i = 1, \ldots, n; s_i \in d_i\}$ a real number, which represents the utility that deal s yields for agent j.

The most basic form to represent a utility function is to make an enumeration of the points in the solution space which yield a non-zero utility value. In this way, an agent's utility function may be represented as a set of pairs $\langle s, u(s)\rangle \,|u(s) \neq 0$, where $u(s)$ is the utility of the solution s for the agent. It is easy to see that, though this representation for utility functions is fully expressive, its cardinality may grow greatly with the number of issues or with the cardinality of each issue's domain. Because of that, more succinct representations for utility functions are used in most cases. Examples of such representations which are widely used in the negotiation literature are linear-additive utility functions [14] or k-additive utility functions [22].

Another widely used way to represent preferences and utility functions is the use of constraints over the values of the attributes. There is a vast variety of multi-attribute negotiation models and approaches making use of constraints in different forms, from hard constraints to soft, probabilistic or fuzzy constraints [31,47,52]. There are several reasons which favor the use of constraints in negotiation models. First, they allow for efficient methods for preference elicitation. Moreover, constraints allow to express dependencies between the possible values of the different attributes. Finally, the use of constraints for offer expression allow to limit the

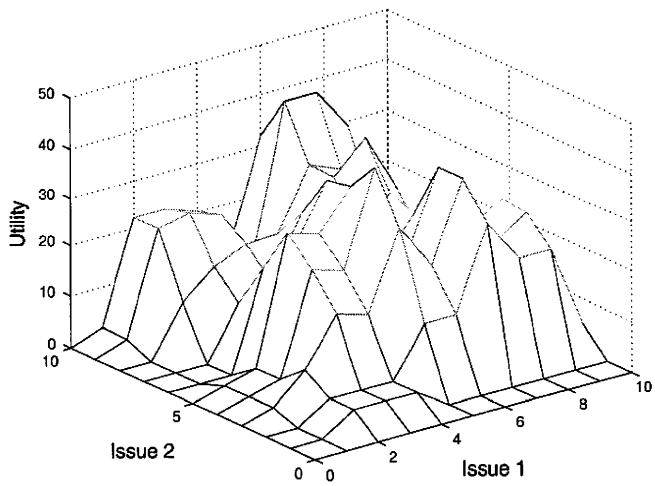

Fig. 1 Example of a nonlinear utility space defined by means of weighted constraints

region of the solution space which has to be explored in a given negotiation step. Reducing the region of the utility space under exploration according to the constraints exchanged by agents is a widely used technique in automated negotiation [48], since it makes the search for agreements a more efficient process than when using positional bargaining, specially in complex negotiation scenarios.

A particular case of constraint-based utility representation which has been used to model complex utility spaces for negotiation are *weighted constraints*. There is a utility value for each constraint, and the total utility is defined as the sum of the utilities of all satisfied constraints.

More formally, the utility space of the agents may be defined as a set of constraints $C = \{c_k | k = 1, \ldots, l\}$. Each constraint c_k has an associated utility value $u(c_k)$. If we note as $s \in x(c_k)$ the fact that a given contract $s = \{s_i | i = 1, \ldots, n\}$ is in the set of contracts that satisfy constraint c_k, an agent's utility for contract s may be defined as

$$u(s) = \sum_{c_k \in C | s \in x(c_k)} u(c_k),$$

that is, the sum of the utility values of all constraints satisfied by s. This kind of utility functions produces nonlinear utility spaces, with high points where many constraints are satisfied, and lower regions where few or no constraints are satisfied. Figure 1 shows an example of the kind of utility spaces which may be modeled using weighted constraints.

2.1.3 Agent strategies, mechanism stability and incentive-compatibility

In an automated negotiation, a strategy guides the decision making process of an agent throughout the different stages of the negotiation protocol [41]. The main challenge in an automated negotiation scenario as far as decision mechanisms are concerned is to design *rational* agents, able to choose an adequate negotiation strategy. In negotiations among selfish agents, negotiation mechanisms must be designed in a way that makes them stable, understanding stability as the impossibility (or at least difficulty) of the strategic manipulation of the mechanisms. This means that the mechanisms should motivate the agents to act in an adequate way, since if a rational, selfish agent may benefit from taking a strategy which is different to the one expected by the protocol, it will do so. This problem is closely related to the notion of ***equilibrium*** defined in game theory. In an equilibrium, each player of the game has adopted a strategy that they have no rational incentive to change (because it is the best alternative, given the circumstances). There are different equilibrium conditions which can be defined, like dominant strategies [40,83], Nash equilibrium [60] or Bayes-Nash equilibrium [24].

Achieving stability in a negotiation mechanism does not guarantee to reach solutions maximizing social welfare. Therefore, stability must not be used as a single criterion to evaluate decision mechanisms, and social welfare should also be considered. An specially illustrative example is the *prisoner's dilemma* [65], which describes an scenario where Nash equilibrium yields low utility values for the agents involved. A more generic concept which is becoming widely used to characterize situations where individual rationality leads agents to results which yield low social welfares is the notion of *price of anarchy*. The price of anarchy was first introduced in [63] in the context of selfish routing, as a measure of loss of social efficiency due to selfish behavior. In the context of a problem of social welfare maximization, price of anarchy can be defined as follows:

Definition 1 *Price of anarchy.* [63] The *price of anarchy (PoA)* in a given game is defined as the ratio between the social welfare of the best possible outcome of the game and the social welfare of the worst Nash equilibrium in the game:

$$PoA = \frac{\max_{s \in S} sw(s)}{\min_{s \in S_{\text{Nash}}} sw(s)},$$

where S is the set of all possible outcomes of the game, $S_{\text{Nash}} \subseteq S$ is the set of all possible outcomes induced by a Nash equilibrium in the game, and $sw(s)$ is the social welfare of a given outcome s.

Defined in this way, price of anarchy gives an indication of the potential loss in a given game when individually rational agents are confronted. In situations where PoA is high, additional mechanisms which incentivize social behavior are desirable, in order to modify the equilibrium conditions of the game and reduce this value of PoA, thus improving the stability of the protocol. Stability, however, may also come at a price. Even when worst-case equilibria can be avoided, equilibrium conditions may lead to solutions which are distant to the social optimum (generally due to the fact that stability enhancing measures favor "fair" solutions against Pareto-optimal ones). To measure this, price of stability is introduced in an analogous manner:

Definition 2 *Price of stability.* [2] The *price of stability (PoS)* in a given game is defined as the ratio between the social welfare of the best possible outcome of the game and the social welfare of the best Nash equilibrium in the game:

$$PoS = \frac{\max_{s \in S} sw(s)}{\max_{s \in S_{\text{Nash}}} sw(s)},$$

where S is the set of all possible outcomes of the game, $S_{\text{Nash}} \subseteq S$ is the set of all possible outcomes induced by a Nash equilibrium in the game, and $sw(s)$ is the social welfare of a given outcome s.

Taking this into account, when mechanisms are introduced to reduce price of anarchy in a game, their impact over price of stability should also be evaluated.

Another threat to mechanism stability is strategic revelation of information. In incomplete information scenarios [34], since the agents' beliefs about the preferences of a given agent may influence the decision mechanisms they use, an agent may use as a strategy to lie about its own preferences in order to manipulate the decision mechanisms of the rest of the agents to its own benefit. This raises an additional concern to mechanism design [83].

It would be desirable to design protocols which are not prone to be manipulated through insincere revelation of information. *Incentive-compatibility* is defined as the property of a negotiation mechanism which makes telling the truth the best strategy for any agent, assuming the rest of the agents also tell the truth.Though incentive-compatibility is usually independently studied, it is closely related to the notions of strategic equilibrium seen above. In particular, incentive-compatibility may be seen as the property of a negotiation model where, regarding the possibility of telling or not telling the truth, having all agents telling the truth is a Nash equilibrium. A more restrictive property is (*strategy-proofness*), which imposes truthful revelation of information to be a dominant strategy. This means that for any agent the best choice is to tell the truth regardless of the other agents' attitudes towards sincerity [19].

An example of an incentive-compatible protocol is the *Vickrey auction*. The Vickrey auctions are second-price, sealed, one shot auctions. In this kind of auction, that an agent i bids above its real utility value $u_i(s)$ is a bad strategy, since there is a chance that the second highest bid is also above that utility value, which would imply that the agent would have to pay for the product more than its value. Furthermore, as Vickrey auction is second price, bidding below the utility level $u_i(s)$ is also a bad strategy, since it reduces the chance to bid without any advantage, as the price the agent will have to pay for the product is not given by its bid, but by the second highest bid. Another incentive compatible mechanism is the Clarke tax method [11], where a tax is imposed to each agent once the negotiation has ended, and this tax makes each agent "pay" for the impact that its participation had over other agents' utilities, showing that, in this way, if an agent's false valuation changes the negotiation result, the utility obtained by that agent (after taxes are applied) is never higher that the utility it would have gained using truthful valuations [83].

2.2 Negotiation, optimization and complexity

Though there has been an increasing interest in complex negotiations in the last years, little efforts have been made to study complexity itself within negotiation (apart from computational complexity, which has been thoroughly studied in many scenarios). Therefore, if we want to be able to assess complexity in negotiations, we need to resort to other knowledge areas. One area where many authors have dealt with complexity characterization and measurement is optimization. In fact, negotiation scenarios and optimization problems are often closely related, since there are many similarities in the ways both problem families are defined and addressed. For example, negotiating agents are usually utility optimizers, and negotiation mechanisms are often evaluated in terms of their ability to reach Pareto-optimal solutions. In negotiation, Pareto-optimal solutions are those where payoff cannot be improved for any of the agents without decreasing the payoff for another agent. This concept of Pareto-efficiency is also sought in multi-objective optimization, trying to find solutions where no further gains can be achieved in one of the objectives without losing in another [74]. Multi-objective optimization has been widely used for negotiation support [84], and negotiation mechanisms have also been used to solve multiobjective optimization problems, usually by distributing the different objectives among negotiating agents [75]. Therefore, some of the concepts studied in multiobjective optimization may be used in negotiation, and vice versa.

In the context of a multi-attribute negotiation, complexity of a given scenario may depend at least on the number of issues, the level of interdependency between the preferences on the issues, the domain of the issues, the possibility of change over time of the negotiation context, the method used to describe preferences and the structural properties of the agent's utility spaces. In general, a large number of issues with a high interdependency and a large domain contribute to more complex preference spaces. If the negotiation context changes over time, complexity also increases. The method to describe preferences also has an influence in the complexity of the negotiation scenario. This is specially true when optimization techniques are used to find high utility regions within an agent's utility spaces, or to find deals among different agents. A constraint-based preference space, for instance, may present discontinuities which make gradient based optimizers not applicable, while differentiable utility functions contribute to a faster local optimization. Therefore, to study complexity in negotiation scenarios, we may find useful to characterize structural complexity of the agents' utility spaces, and to this end we may benefit from existing research on function characterization for optimization.

In this context, and more specifically in the field of optimization using evolutionary algorithms, structural complexity analysis plays a crucial role, since algorithm search capabilities are greatly impacted by some structural properties of the optimized function, which is usually known as *fitness landscape* in evolutionary computation.

An interesting detail about fitness landscapes is that they include the definition of a *neighborhood operator* ϕ, which expresses the probability that the search function (usually, a genetic algorithm) passes from one point in the landscape to another [27]. This operator is directly related to the search mechanism used and its parameters (e.g. simulated annealing temperature or mutation probability for genetic algorithms), which implies an important consequence: the complexity of a utility space may be different depending on the considered search algorithm and its parameters. This operator also defines the concept of *neighbor solutions* in the space, which in turn influences the definition of local optima (maxima and minima), and therefore the structural properties of a fitness landscape which are interesting regarding search complexity within the space, such as modality [28], ruggedness, smoothness and neutrality [80].

Once the properties which has an influence on the complexity of a fitness landscape or a solution space have been studied, techniques which allow to measure the complexity of a given space are needed. Most of the approaches we can find in the literature are based on the correlation between different samples of the fitness function f, like *fitness distance correlation* metrics [79] or stochastic models representing the correlation structure of the space [27]. A metric which is easy to compute in most scenarios and allows to make quantitative evaluations about the complexity of a fitness or utility landscape is *correlation length* or *correlation distance*. Correlation distance is defined as the minimum distance ψ which makes correlation fall below a given threshold (usually 0.5), which gives an idea of the distance we can move throughout the solution space while keeping a certain correlation between samples [53].

2.3 Related research on automated negotiation in complex utility spaces

Klein et al. [39] present, as far as we are aware, the first negotiation protocols specific for complex preference spaces. They propose a simulated annealing-based approach, a refined version based on a parity-maintaining annealing mediator, and an unmediated version of the negotiation protocol. Of great interest in this work are the positive results about the use of simulated annealing as a way to regulate agent decision making, along with the use of agent expressiveness to allow the mediator to improve its proposals. However, this expressiveness is somewhat limited, with only four possible valuations which allow the mediator to decide which contract to use as a parent for mutation, but not in which direction to mutate it. On the other hand, the performed experiments only consider the bilateral negotiation scenario, though authors claim that the multiparty generalization is simple. Finally, the family of negotiation protocols they propose are specific for binary issues and binary dependencies. Higher-order dependencies and continuous-valued issues, common in many real-world contexts, are known to generate more challenging utility landscapes which are not considered in their work.

Luo et al. [51] propose a fuzzy constraint based framework for multi-attribute negotiations. In this framework a buyer agent defines a set of fuzzy constraints to describe its preferences. The proposals of the buyer agent are a set of hard constraints which are extracted from the set of fuzzy constraints. The seller agent responds with an offer or with a relaxation request. The buyer then decides whether to accept or reject an offer, or to relax some constraints by priority from the lowest to highest. In Lopez-Carmona and Velasco [49], Lopez-Carmona et al. [50] an improvement to Luo's model is presented. They devise an expressive negotiation protocol where proposals include a valuation of the different constraints, and seller's responses may contain explicit relaxation requests. It means that a seller agent may suggest the specific relaxation of one or more constraints. The relaxation suggested by a seller agent is based on utility and viability criteria, which improves the negotiation process. Though these constraint-based works model discontinuous preference spaces, the operators used to compute utility and the utility spaces defined yield monotonic preference spaces, which are far from the complex preference spaces covered in our work.

Another interesting approach to solve the computational cost and complexity of negotiating interdependent issues is to simplify the negotiation space. Hindriks et al. [26] propose a weighted approximation technique to simplify the utility space. They show that for smooth utility functions the application of this technique results in an outcome that closely matches the outcome based on the original interdependent utility structure. The method is evaluated for a number of randomly generated utility spaces with interdependent issues. Experiments show that this approach can achieve reasonably good outcomes for utility spaces with simple dependencies. However, an approximation error that deviates negotiation outcomes from the optimal solutions cannot be avoided, and this error may become larger when the approximated utility functions become more complex. Authors acknowledge as a necessary future work to study which kind of functions can be approximated accurately enough using this mechanism. Another limitation of this approach is that it is necessary to estimate a region of utility space where the actual outcome is expected to be (i.e. it is assumed that the region is known a priori by the agents).

In Robu et al. [69] utility graphs are used to model issue interdependencies for binary-valued issues. Utility graphs are inspired by graph theory and probabilistic influence networks to derive efficient heuristics for non-mediated bilateral negotiations about multiple issues. The idea is to decompose highly non-linear utility functions in sub-utilities of clusters of inter-related items. They show how utility graphs can be used to model an opponent's preferences. In this approach agents need prior information about the maximal structure of the utility space to be explored. Authors argue that this prior information could be obtained through a history of past negotiations or the input of domain experts. However, our approach has the advantage that outcomes can be reached without any prior information and that it is not restricted to binary-valued issues.

There are several proposals which employ genetic algorithms to learn opponent's preferences according to the history of the counter-offers based upon stochastic approximation. In Choi et al. [9] a system based on genetic-algorithms for electronic business is proposed. In this work the utility functions are restricted to take a product combination form (i.e. utility of an outcome is the product of the utility values of the different issues). The objective function used is based on the comparison of the changes of consecutive offers. Small changes of an issue suggest that this issue is more important. For each new population, the protocol enforces that the generated candidates cannot be better than the previous offer. Unlike other negotiation models based on genetic algorithms, this proposal adapts to the environment by dynamically modifying its mutation rate. Lau et al. [45] have also reported a negotiation mechanism for non-mediated automated negotiations based on genetic algorithms. The fitness function relies on three aspects: an agent's own preference, the distance of a candidate offer to the previous opponent's offer, and time pressure. In this work agents' preferences are quantified by a linear aggregation of the issue valuations. However, non-monotonic and discontinuous preference spaces are not explored. In Chou et al. [10] a genetic algorithm is proposed which is based on a joint elitism operation and a joint fitness operation. In the joint elitism operation an agent stores the latest offers received from the opponent. The joint fitness operation combines agent's own utility function and euclidean distance to the opponent's offer. In this work two different negotiation scenarios are considered. In the first one utility is defined as the weighted sum of the different issue values (i.e. issues are independent). The second scenario defines a utility function where there is a master issue and a set of slave issues. Utility is calculated as the weighted sum of the different issue values, but the weights of the slave and master issues change according to the value of the master issue.

In Yager [87] a mediated negotiation framework for multi-agent negotiation is presented. This framework involves a mediation step in which the individual preference functions are aggregated to obtain a group preference function. The main interest is focused on the implementation of the mediation rule where they allow a linguistic description of the rule using fuzzy logic. A notable feature of their approach is the inclusion of a mechanism rewarding the agents for being open to alternatives other than simply their most preferred. The negotiation space and utility values are assumed to be arbitrary (i.e. preferences can be non-monotonic). However, the set of possible solutions is defined a priori and is fixed. Moreover, the preference function needs to be provided to the mediation step in the negotiation process, and pareto-optimality is not considered. Instead, the stopping rule is considered, which determines when the rounds of mediation stop.

Fatima et al. [18] analyze bilateral multi-issue negotiation involving nonlinear utility functions. They consider the case where issues are divisible and there are time constraints in the form of deadlines and discounts. They show that it is possible to reach Pareto-optimal agreements by negotiating all the issues together, and that finding an equilibrium is not computationally easy if the agents' utility functions are nonlinear. In order to overcome this complexity they investigate two solutions: approximating nonlinear utilities with linear ones; and using a simultaneous procedure where the issues are discussed in parallel but independently of each other. This study shows that the equilibrium can be computed in polynomial time. An important part of this work is the complexity analysis and estimated approximation error analysis performed over the proposed approximated equilibrium strategies. Heuristic approaches have generally the drawback of the lack of a solid mathematical structure which guarantees their viability, which raises the need of an exhaustive experimental evaluation. An adequate complexity analysis and establishing a bound over the approximation error contribute to give heuristic approaches part of the technical soundness they usually lack. Among the limitations of the proposal, we can point out that this work is focused on symmetric agents where the preferences are distributed identically, and the utility functions are separable in nonlinear polynomials of a single variable. This somewhat limits the complexity of the preference space.

Finally, combinatorial auctions [21,29,72,73,82,86] can enable large-scale collective decision making in nonlinear domains, but only of a very limited type (i.e. negotiations consisting solely of resource allocation decisions). Multi-attribute auctions, wherein buyers advertise their utility functions, and sellers compete to offer the highest-utility bid [5,78,64] are also aimed at a fundamentally limited problem (a purchase negotiation with a single buyer) and require full revelation of preference information.

In summary, in the existing research nearly all the models which assume issue interdependency rely on monotonic utility spaces, binary valued issues, low-order dependencies, or a fixed set of defined a priori solutions. Simplification of the negotiation space has also been reported as a valid approach for simple utility functions, but it cannot be used with higher-order issue dependencies, which generate highly uncorrelated utility spaces. Therefore, new approaches are needed if automated negotiation is to be applied to settings involving non-monotonic, highly uncorrelated preference spaces.

3 An auction based approach for negotiations in highly uncorrelated, constraint based utility spaces

In this work we analyze agents' strategic behavior and mechanism stability for a mediated, auction-based negotiation approach we designed for highly uncorrelated, constraint based utility spaces [30,54]. To make such strategic analysis easier to understand, in this section we motivate and review the most relevant aspects of our negotiation model.

3.1 Negotiation domain and agent preference model

We explore the problem of negotiating complex contracts, which was first introduced by Klein et al. [39]. Contracts are defined as a set of issues or *clauses*, each of which may have a value. The aforementioned authors limited encounters to bilateral negotiations (i.e. two negotiating agents), and clauses were limited to binary values, meaning than the clause was or not present in a given contract. Even with such restrictions, the domain of the solution space may become very large. For instance, a negotiation scenario with 50 possible clauses would yield a search space of about 10^{15} possible contracts. This, along with the assumption of non-linearity in the agents' preference spaces, imposed serious difficulties for the negotiation. First, agents needed to use nonlinear optimization mechanisms to try to find desirable contracts within their own preference spaces. Once desirable contracts for each agent were identified, building agreements had its own difficulties, since the scenario was assumed competitive, and thus agents were not inclined to fully disclose their preferences.

Though there are negotiation scenarios about complex contracts which may be modeled with such a solution space, in many cases more than two agents are involved in a negotiation. Also, most contracts may have non-binary clauses. In a rental agreement, for instance, clauses may state the rent, the security deposit or the length of the lease. A labor agreement may include different insurance options. Such issues may have a larger domain, which can greatly increase the solution and preference space complexity.

Taking this into account, in this work we focus in the general case of multilateral negotiations of complex contracts, where the issues or clauses included in the contracts have discrete domains. We also assume that agents' preferences about the different issues are not independent, which means that the utility that a given clause in the contract yields for an agent may depend on the presence of other clauses. Interdependence between attributes in agent preferences can be described by using different categories of functions, like K-additive

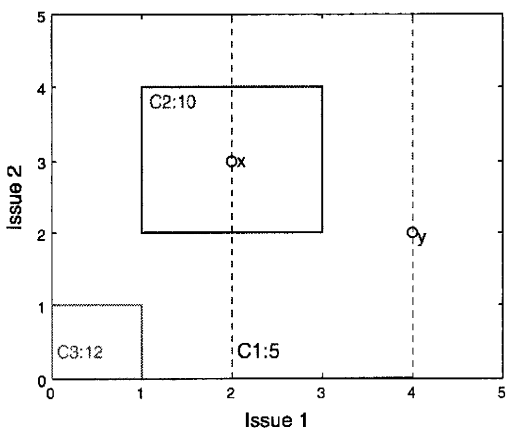

Fig. 2 Example of a utility space with two issues and three constraints

utility functions [8,22], bidding languages [61], or weighted constraints [31]. In this work, we model this dependency in agent preferences by means of weighted constraints, which are a natural way to model user preferences and to express dependencies between issues [51]. These constraints may represent ranges of values over different issues, meaning that when *all* the clauses affected by the constraint have values satisfying it, that yields a given utility value for the agent. The set of an agent's constraints and their associated utility values builds the preference space of the agent.

From a geometrical point of view, each constraint represents a region with one or more dimensions, and has an associated utility value. The number of dimensions of the space is given by the number of issues n under negotiation, and the number of dimensions of each constraint must be lesser than or equal to n. The utility yielded by a given potential solution (contract) in the utility space for an agent is the sum of the utility values of all the constraints that are satisfied by that contract. Figure 2 shows an example for two issues and three constraints: a unary constraint $C1$ and two binary constraint $C2$ and $C3$. The utility values associated to the constraints are also shown in the figure. In this example, contract x would yield a utility value for the agent $u(x) = 15$, since it satisfies both $C1$ and $C2$ (that is, constraints $C1$ and $C2$ overlap, creating a region of higher utility). Contract y, on the other hand, would yield a utility value $u(y) = 5$, because it only satisfies $C1$. It can also be noted that unary constraint $C1$ can be seen as a binary constraint where the width of the constraint for issue 2 is all the domain of the issue, so we can generalize and say that all constraints have n dimensions.

More formally, we can define the negotiation domain and an agent's preference model by means of a set of definitions:

Definition 3 *Issues under negotiation.* The *issues under negotiation* are defined as a finite set of variables $X = \{x_i | i = 1, \ldots, n\}$.

Definition 4 *Solution space.* The negotiation *solution space* is defined by the values that the different values may take. To simplify, we assume that issues take values from the domain of integers $[0, x_D^{\max}]$:

$$D = [0, x_D^{\max}]^n$$

Definition 5 *Contract or potential solution.* A *contract* or *potential solution* to the negotiation problem is a vector $s = \{s_i | i = 1, \ldots, n\}$ such that $s \in D$ defined by the issues' values.

Definition 6 *Constraint.* A *constraint* is a set of intervals which define the region where a contract must be contained to satisfy the constraint. Formally, a constraint c is defined as

$$c = \{I_i^c | i = 1, \ldots, n\},$$

where $I_i^c = [x_i^{\min}, x_i^{\max}]$, with $x_i^{\min}, x_i^{\max} \in [0, x_D^{\max}]$ defines the minimum and maximum values for each issue to satisfy the constraint. Constraints defined in this way describe hyper-rectangular regions in the n-dimensional space.

Definition 7 *Constraint satisfaction.* A contract s *satisfies* a constraint c if and only if $x_i^s \in I_i^c \forall i$. For notation simplicity, we denote this as $s \in x(c)$, meaning that s is in the set of contracts that satisfy c.

Definition 8 *Preference space.* An agent's *preference space* may be defined as a tuple

$$\langle C, \Omega \rangle,$$

where $C = \{c_k | k = 1, \ldots, l\}$ is a set of constraints over the values of the issues x_i for the agent and $\Omega = \{\omega(c_k) | k = 1, \ldots, l; \omega(c_k) \in \mathbb{N}^+\}$ is a set of weights or utility values, such that $\omega(c_k)$ is the associated *utility value* for constraint c_k. For simplicity, we will assume that constraint weights take values from the set of positive integers.

Definition 9 *Utility function.* An agent's *utility function* for a contract s is defined as

$$u(s) = \sum_{c_k \in C | s \in x(c_k)} u(c_k),$$

that is, the sum of the utility values of all constraints satisfied by s.

This kind of utility functions produces nonlinear utility spaces, with high points where many constraints are satisfied, and lower regions where few or no constraints are satisfied [31]. As we have seen in Sect. 2.2, the degree of complexity of the utility spaces produced depends on the number of issues, the domain of the issues and the structural properties of the utility spaces. For the purpose of this work, we make the following assumptions:

- We assume that the number of issues and the domains of the issues are such that they make exhaustive search within the utility space of the agents intractable.
- We assume that the utility spaces of the agent are highly uncorrelated, and so no *a priori* assumptions may be made about where high utility contracts may be located. Therefore, agents may need to resort to local nonlinear optimization techniques to identify such high-utility contracts.
- We assume knowledge about other agent's preferences not to be common (i.e. agents do not know their opponent preference structures, neither they can compute opponent's utility for a given contract).
- We assume that the negotiation setting is competitive, and that agents may be unwilling to reveal too much information about their preferences to the other negotiating agents.

The negotiation protocol and mechanisms proposed, which are described in the next sections, are specifically designed to address this negotiation setting. However, through the study performed in the latter sections of this paper, some of the assumptions are relaxed to evaluate the influence of agent strategies and variations in the correlation lengths of the utility spaces over the negotiation outcomes.

3.2 Interaction protocol

As we stated at the beginning of this section, our model relies on a mediated, auction-based protocol to support agent interaction. The reason for the choice of such a protocol is two-fold. On one hand, the auction-based approach allows to efficiently cope with many of the challenges imposed by multilateral interactions [39]. On the other hand, the use of a mediator allows to decouple individual agent goals (maximizing their own payoff) from social negotiation goals (usually, reaching an agreement which maximizes social welfare). This makes easier mechanism and strategy definition, since agents can be assumed selfish and competitive, while the mediator can be entitled with the more-cooperative task of pursuing social welfare.

Since the main focus of this work is on agent strategic behavior, we have chosen a simple, one shot, auction based interaction protocol for the negotiation, which mainly consists of two steps:

1. *Bidding:* Each agent j generates a set of n_b^j bids $B^j = \{b_i^j | i = 1, \ldots, n_b^j\}$, where each bid b_i^j represents a region within the solution space which only contains contracts that agent j would be willing to accept as solutions. Each agent sends its bid set B^j to the mediator, along with the utility associated to each bid.
2. *Deal identification:* The mediator tries to find overlaps between the bids of the different agents. The regions of the contract space corresponding to the intersections of at least one bid of each agent are tagged as potential solutions. A final deal is chosen from the set of potential solutions, according to social welfare criteria.

The protocol, as described, is fairly straightforward, and the decision mechanisms which agents employ for bidding and deal identification are the ones which mostly determine the effect of agent strategic behavior. There are many different mechanisms which can be used in this context. In the following we briefly describe the ones we have found to yield better results in terms of negotiation efficiency and failure rate. All these mechanisms rely on the concept of *quality factor*, which we introduce in the following section.

3.3 Constraint/bid quality factor

The use of weighted constraints generates a "bumpy" utility space, with many peaks and valleys. However, the degree of "bumpiness" is highly dependent on the way the constraint set is generated, and specially on the average width of the constraints. Figure 3 shows an example of the resulting two-dimensional utility space for 50 binary constraints, where the domain of the issues is chosen to be [0,9], and constraints are generated by choosing the width of each constraint in each issue randomly within the [3,7] interval. This generates rather "wide" constraints. On the other hand, Fig. 4 shows an utility space obtained using "narrow" constraints, choosing their widths from the [1,2] interval. Comparing both figures we can see that, though both utility spaces are nonlinear, the space generated using narrow constraints is more complex, with narrower peaks and valleys. As the number of issues under consideration increases, the differences between having wide or narrow constraints become more relevant. For instance, the average correlation length for utility spaces generated using [3,7] constraints for six issues is $\psi = 5.9$, while average correlation length for utility spaces generated using [1,2] constraints is $\psi = 2.8$. Though most utility-maximizing negotiation approaches work in scenarios like the example shown in Fig. 3, their performance (in terms of optimality and failure rate) decreases drastically in highly nonlinear scenarios defined using

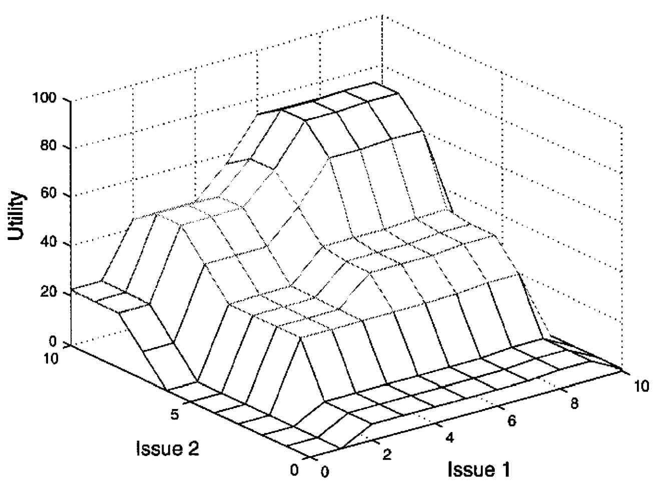

Fig. 3 Example of a nonlinear utility space generated by using "wide" constraints

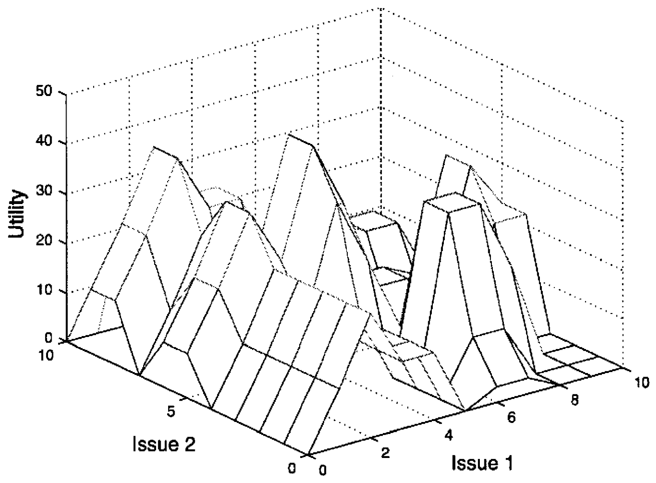

Fig. 4 Example of a highly uncorrelated utility space generated by using "narrow" constraints

narrow constraints, and therefore an alternative approach is needed to deal with these highly uncorrelated utility spaces [55].

If we compare the utility spaces shown in Figs. 3 and 4, we can see that the main difference between them (apart from the absolute utility values, but they have no effect in optimality) is the width of the peaks. Highly-nonlinear scenarios will yield narrower peaks. Utility maximizing agents tend to choose those peaks (or high-utility regions) as bids, and the result is that narrower bids will be sent to the mediator. The width of the bids (or more generally, the *volume* of the bids), will directly impact the probability that the bid overlaps a bid of another agent, and thus its *viability*, that is, the probability of the bid resulting in a deal. Intuitively, in such complex scenarios, an agent with no knowledge of the other agents' preferences should deviate from the "plain utility maximization strategy" and try to adequately balance the utility of their bids (to maximize its own profit) and the volume of those bids (to maximize the probability of a successful negotiation).

We formally represent this through the following definitions:

Definition 10 *Volume of a region.* The *volume* of a given region r within the solution space D (be it a constraint or a bid) is defined as the cardinality of the set of contracts contained within the region.

$$v_r = |r|, \quad \text{with } r \subset D$$

Definition 11 *Quality factor.* The *quality factor* of a given region r within the solution space D (be it a constraint or a bid) is defined as

$$Q_r = u_r^\alpha \cdot v_r^{1-\alpha},$$

where u_r and v_r are, respectively, the utility and volume of the bid or constraint r, and $\alpha \in [0, 1]$ is a parameter which models the attitude of the agent. A social, cooperative or risk averse agent ($\alpha < 0.5$) will tend to qualify as better bids those that are wider, and thus are more likely to result in a deal. A risk willing, highly competitive or selfish agent ($\alpha > 0.5$) will, in contrast, give more importance to bid utility.

3.4 Bid generation mechanisms

3.4.1 Contracts sampling and simulated annealing

We can see the problem of finding the adequate set of bids for an agent as a local optimization problem, since for rational agents bids should be high-utility regions or, more generally, regions of high quality factor. Therefore, nonlinear optimization mechanisms may be used by the agents to find those regions suitable to be sent to the mediator as bids. Here we describe a bidding mechanism based on simulated annealing, which consists of three steps:

1. *Sampling:* Each agent takes a fixed number of random samples from the contract space, using a uniform distribution.
2. *Adjusting:* Each agent applies simulated annealing to each sample to try to find a local optimum in its neighborhood. The function which is tried to maximize by the simulated annealing optimizer is the quality factor Q. Since the quality factor Q is a feature of a region, not a contract, the adjusted contracts must be mapped to the high utility regions where they are contained before they are accepted or rejected by the simulated annealing engine. This can be easily done by checking all constraints in the agent preference model and computing the intersection of the constraints which are satisfied by the candidate contract. The volume of this intersection can then be used to compute the quality factor Q of the region.This results in a set of high-quality contracts.
3. *Bidding:* Each agent generates a bid for each high-quality, adjusted contract. The bids are generated as the intersection of all constraints which are satisfied by the contract. Bids defined in this way represent hyper-rectangle regions in the n-dimensional solution space. Each agent sends its bids to the mediator, along with the utility associated to each bid.

The bid generation mechanism may be seen formally in Algorithm 1. Also, some details about the mechanism are highlighted. The algorithm is run for a fixed number of iterations n_b, which imposes the maximum number of generated bids (1). The function $adjust\_annealing$ $(x, Q(., \alpha), n_{SA}, T_{SA})$ uses simulated annealing to return a region of optimal quality factor using as starting point a sampled contract x (2). There are some parameters in this function which may be adjusted to influence the behavior of the simulated annealing algorithm, like the initial temperature and the number of iterations. As studied in [30], best results in term of optimality and efficiency are achieved using $n_{SA} = 30$; $T_{SA} = 30$. Moreover, the algorithm discards any contract which, once adjusted, yields less utility than the agent's reservation value u_R, which guarantees that all bids would be accepted by the agent as final solutions (3). Finally, duplicate bids or bids contained in other bids are also discarded (4). Some of these ideas are also used in the next bidding mechanism described.

Algorithm 1: Bid generation using simulated annealing over quality factor

```
  Input:
    D: solution space domain
    n_b: maximum number of bids
    u_R: reservation utility for the agent
    C: constraint set defining agent's utility space
    u: agent's utility function
    α: agent's attitude parameter
    Q: function which computes the quality factor of a region
    n_SA: iteration bound for the simulated annealing algorithm
    T_SA: initial temperature for the simulated annealing algorithm
  Output:
    B: bid set
  B = ∅;
  k = 0;
1 while k < n_b do
      k = k + 1;
      x = random_contract();
2     b = adjust_annealing(x, Q(., α), n_SA, T_SA);
3     if u(b) ≥ u_R then
          B = B ⋃ b;
  end
4 remove_duplicates(B)
```

3.4.2 Maximum weight independent set and the max-product algorithm

There have been a number of recent successful efforts in literature for using graphs to model negotiation scenarios in multi-link negotiations [89] or combinatorial auctions [21]. One of the advantages of such approaches is that they allow to use well-known graph methods for solving the negotiation problem. In our case, graphs provide an alternative perspective for the bidding process, looking at the constraint-based agent utility space as a weighted undirected graph. Consider again the simple utility space example shown in Fig. 2. Think about each constraint as a node in the graph, with an associated weight which is the utility value associated to the constraint. Now we will connect all nodes whose corresponding constraints are *incompatibles*, that is, they have no intersection. The resulting graph is shown in Fig. 5.

To find the highest utility bid in such a graph can be seen as finding the set of unconnected nodes which maximizes the sum of the nodes' weights. Since only incompatible nodes are connected, the corresponding constraints will have non-null intersection. In the example, this would be achieved by taking the set $\{C1, C2\}$. The problem of finding a maximum weight set of unconnected nodes is a well-known problem called maximum weight independent set