

# A logic-based framework to compute Pareto agreements in one-shot bilateral negotiation

Azzurra Ragone and Tommaso Di Noia and Eugenio Di Sciascio<sup>1</sup> and Francesco M. Donini<sup>2</sup>

**Abstract.** We propose a logic-based approach to automated one-shot multi-issue bilateral negotiation. We use logic in two ways: (1) a logic theory to represent relations among issues – *e.g.*, logical implication – in contrast with approaches that describe issues as uncorrelated with each other; (2) utilities over formulas to represent agents having preferences over different bundles of issues. In this case, the utility assigned to a bundle is not necessarily the sum of utilities assigned to single elements in the bundle itself. We illustrate the theoretical framework and the one-shot negotiation protocol, which makes use of a facilitator to compute some particular Pareto-efficient outcomes. We prove the computational adequacy of our method by studying the complexity of the problem of finding Pareto-efficient solutions in a propositional logic setting.

## 1 Introduction

Many recent research efforts have been focused on automated negotiation among agents in various contexts, including e-marketplaces, resource allocation settings, online auctions, supply chain management and, generally speaking, e-business processes. Basically, two different approaches to automated negotiation exist: *centralised* and *distributed* ones. In the first ones, agents elicit their preferences and then a facilitator, or some central entity, selects the most suitable deal based on them. In the latter ones, agents negotiate through various negotiation steps reaching the final deal by means of intermediate deals, without any external help [3]. Although the main target of an agent is reaching a satisfying agreement, knowing if it is Pareto-efficient<sup>3</sup> is a matter that cannot be left out. Furthermore it is fundamental to assess *how hard* it is to find a Pareto-efficient agreement, and if such an agreement actually exists. Recently, there has been a growing interest toward multi-issue negotiation, also motivated by the idea that richer and expressive descriptions of demand and supply can boost e-marketplaces (see *e.g.*, [21] for a reasonable set of motivations) but –to the best of our knowledge– also in recent literature, issues are still described as uncorrelated terms, without considering any underlying semantics, which could instead be exploited –among other reasons– to perform a pre-selection of candidate deals. Moreover, agents can have different preferences over different bundles of issues, and utility assigned to each bundle is not merely the sum of utilities assigned to a single resource in that bundle. When multiple issues are involved

in the negotiation process, it is usually the case that the issues are not completely independent; sometimes an issue implies another one (for instance, `SatelliteAlarm` implies `AlarmSystem` in an automotive domain), sometimes their relation is more complex (as in bundles of optionals). We propose to fill this lack of expressiveness by relating with each other the issues involved in the process via a logical theory. In particular, here we present a framework for automated multi-issue bilateral negotiation where logic is exploited both to represent existing relations between issues and in utilities over formulas. The rest of the paper is structured as follows: next section presents a brief summary of the huge literature on bilateral negotiation. Section 3 presents the reference scenario and assumptions we make. In Section 4 the **Multi-issue Bilateral Negotiation** problem (MBN) is formally defined based on a logic framework. We define utility functions over a set of logic formulas representing preferences on bundles of interrelated issues. In Section 5, we prove the computational adequacy of our method by studying the complexity of the problem of finding Pareto-efficient solutions in a propositional logic setting. The bargaining process is detailed in Section 6 and an illustrative example is given in Section 7. Conclusion and future work close the paper.

## 2 Related Work

Automated bilateral negotiation among agents is a challenging research topic, widely investigated, both in artificial intelligence (AI) and in microeconomics research communities. Several definitions have been proposed in the literature. Rubinstein [19] defined the *Bargaining Problem* as the situation in which “two individuals have before them several possible contractual agreements. Both have interests in reaching agreement but their interests are not entirely identical. What ‘will be’ the agreed contract, assuming that both parties behave rationally?” Also Nash [13] defined: “a TWO-PERSON bargaining situation involves two individuals who have the opportunity to collaborate for mutual benefit in more than one way.” The negotiation is *distributive*, also called Win-Lose negotiation [18], when a single item is involved and bargainers have opposite interests: if one gets more the other part gets less. If several items are negotiated simultaneously the negotiation is called *integrative* (Win-Win negotiation), each issue has a different utility (or score) for each player and differently from the distributive case, the players are not strict competitors: they can cooperate and if one gets more, the other player does not necessarily get less. Finally it is possible to have *many issues* to be negotiated among *many parties*; this is a more complex setting, because of the interplay among shifting coalitions: parties can join and act against other parties involved [17]. In game theory, the bargaining problem has been modeled either as *cooperative* or

<sup>1</sup> SisInfLab - Technical University of Bari, Italy, email: {a.ragone, t.dinoia, disciascio} @poliba.it

<sup>2</sup> University of Tuscia, Italy, email: donini@unitus.it

<sup>3</sup> An agreement is Pareto-efficient if there is no other agreement that will make at least one participant better off without making at least one other participant worse off. If a negotiation outcome is not Pareto-efficient, then there is another outcome that will make at least one participant happier while keeping everyone else at least as happy [7].

*non-cooperative* games [6]. The first approach, given a set of axioms and a coalition, determines how to split the surplus among the participants; so the aim is finding a solution given a set of possible outcomes. Instead, in non-cooperative games there are a well-defined set of rules and strategies; in such a game it is possible to define an *equilibrium* strategy, which ensures the rational outcomes of a game: no player could benefit by unilaterally deviating from her strategy, given the other players follow their own strategies [9]. On the other hand, AI-oriented research has been more focused on automated negotiation among agents and on designing high-level protocols for agent interaction [10]. Agents can play different roles: act on behalf of buyer or seller, but also play the role of a mediator or facilitator. Agents are considered individually rational: they will never accept a deal which involves a loss, *i.e.*, a disadvantageous deal [5]. Different approaches exist to automate negotiation. Many of these consider self-interested agents negotiate over a set of resources in order to obtain an optimal allocation of such resources [5, 4, 3]. Endriss et al. [5] propose an optimal resource allocation in two different negotiation scenarios: one, with money transfer, results in an allocation with maximal social welfare; the second is a money-free framework which results in a Pareto outcome. In [3] agents negotiate over small bundles of resources and a mechanism of resource allocation is investigated, which maximizes the social welfare by means of a sequence of deals involving at most  $k$  items each. Both papers [5, 3] extend the framework proposed in [20] concerning negotiation for (re)allocating tasks among agents. We share with [22] the definition of agreement as a model for a set of formulas from both agents. However, in [22] only multiple-rounds protocols are studied, and the approach taken leaves the burden to reach an agreement to the agents themselves, although they can follow a protocol. Results do not take preferences into account, so that it is not possible to guarantee the reached agreement is Pareto-efficient. Our approach, instead, aims at giving an *automated* support to negotiating agents to reach, in one shot, Pareto agreements. In [16] an initial propositional logic framework was presented. Here we extend and generalize the approach, study the computational complexity of the problem showing how to compute a solution of the problem through combinatorial optimization techniques, and prove the computational adequacy of our method.

### 3 Scenario and assumptions

We illustrate our logic-based negotiation framework in accordance with [8], which defines: the *Space of possible deals*, that is the set of potential agreements; the *Negotiation Protocol*, the set of rules that an agent follows in order to reach an agreement; the *Negotiation Strategy* that an agent adopts given the set of rules specified in the negotiation protocol. In our framework the *Space of possible deals* is computed with the aid of a logic language. The *Negotiation Protocol* we adopt here is a *one-shot* protocol. We refer to *Single-shot* bargaining [17], also called *ultimatum game*, where one player proposes a deal and the other player may only accept or refuse it [2]. Therefore we focus our initial investigation on a one-shot protocol, although several other negotiation protocols exist, like *e.g.*, *alternate-offers* protocol [19] or *monotonic concession* protocol [8]. In some negotiations the parties cannot reach an agreement without some external help, that is without the intervention of a third-party. The third-party can be defined as a *facilitator*, *mediator* or *arbitrator*, depending on her influence and power in the negotiation, that is depending on her role: the more the role is evaluative, more appropriate is to define the third-party intervention as arbitration rather than facilitation or

mediation. We adopt in this paper the centralized approach, so we hypothesize the presence of an electronic facilitator, who may automatically explore the negotiation space and discover Pareto-efficient agreements to be proposed to both parties. The presence of a facilitator and the one-shot protocol is an incentive for the two parties to reveal the true preferences, because they can trust in the facilitator and they have a single possibility to reach the agreement with that counterpart. Usually bargainers may not want to disclose their preferences or utility function to the other party, but they can be ready to reveal these information to a trusted – automated – facilitator helping negotiating parties to achieve efficient and equitable outcomes [18, p.311]. Therefore we propose a one-shot protocol with the intervention of a *facilitator* with a proactive behavior: it suggests to each participant the solution which is Pareto-efficient. For what concerns *strategy*, the players can choose to accept or refuse the solution proposed by facilitator; they refuse if they think possible to reach a better agreement looking for another partner, or another shot, or for a different set of bidding rules. Here we do not consider the influence of the *outside options* in the negotiation strategy [12]. Obviously, in e-marketplaces, both buyer and seller know that the opponent is probably not the only partner available for negotiating with, and that there might be more than one single e-marketplace.

### 4 Logic-based Multi Issue Bilateral Negotiation

We use propositional formulas to model the buyer’s demand and the seller’s supply. Relations among issues are represented by a set  $\mathcal{T}$  – for Theory – of propositional formulas. We point out that each of the following definitions can be reformulated and still keep its validity within First Order Logic as long as only closed formulas are used.

Given a set  $\mathcal{A} \doteq \{A_1, \dots, A_n\}$  of propositional atoms (Attributes), we denote with  $\mathcal{L}_{\mathcal{A}}$  the set of propositional formulas (denoted by greek letters  $\beta, \sigma, \dots$ ) built from  $\mathcal{A}$  using  $\wedge, \vee, \neg, \Rightarrow$ . A *propositional theory* is a finite subset of  $\mathcal{L}_{\mathcal{A}}$ , which we denote by  $\mathcal{T} \subset \mathcal{L}_{\mathcal{A}}$ . An interpretation  $m$  assigns values *true*, *false* to elements of  $\mathcal{A}$ , and we extend  $m$  to  $\mathcal{L}_{\mathcal{A}}$  according to truth tables for connectives. We denote with  $m \models \beta$  the fact that  $m$  assigns *true* to  $\beta \in \mathcal{L}_{\mathcal{A}}$ . A *model* of  $\mathcal{T}$  is an interpretation  $m$  assigning *true* to all formulas of  $\mathcal{T}$ , written  $m \models \mathcal{T}$ . A theory is *satisfiable* if it has a model.  $\mathcal{T}$  logically implies a formula  $\varphi$ , denoted by  $\mathcal{T} \models \varphi$  iff  $\varphi$  is true in all models of  $\mathcal{T}$ . We denote with  $\mathcal{M}_{\mathcal{T}} = \{m_1, \dots, m_n\}$ , the set of all models for  $\mathcal{T}$ , and omit the subscript when no confusion arises.

Usually, in a bilateral negotiation the issues within both the buyer’s request and the seller’s offer can be split into *strict requirements* and *preferences*. The former are constraints the buyer and the seller want to be necessarily satisfied to accept the final agreement – this is what we call *demand/supply* – while the latter are issues they may accept to negotiate on – this is what we call *preferences*.

Anticipating our illustrative example, let us consider an automotive domain and suppose the buyer is asking for a sedan with leather seats, preferably red with car CD reader. She strictly wants a sedan with leather seats and expresses her preferences (she is willing to negotiate on) on the color *red* and the car *CD reader* in order to reach an agreement with the seller.

**Definition 1 (Demand, Supply, Agreement)** *Let  $\mathcal{A}$  be a set of propositional atoms, and  $\mathcal{T} \subset \mathcal{L}_{\mathcal{A}}$ .*

- *a buyer’s demand is a propositional formula  $\beta$  (for Buyer) in  $\mathcal{L}_{\mathcal{A}}$  such that  $\mathcal{T} \cup \{\beta\}$  is satisfiable.*

- a seller's supply is a propositional formula  $\sigma$  (for Seller) in  $\mathcal{L}_A$  such that  $\mathcal{T} \cup \{\sigma\}$  is satisfiable.
- $m$  is a possible deal between  $\beta$  and  $\sigma$  iff  $m \models \mathcal{T} \cup \{\sigma, \beta\}$ , that is,  $m$  is a model for  $\mathcal{T}$ ,  $\sigma$ , and  $\beta$ . We also call  $m$  an agreement.

Intuitively,  $\sigma$  and  $\beta$  stand for the minimal requirements that each partner accepts for the negotiation. On the other hand, if seller and buyer have settled strict attributes that are in conflict with each other, that is  $\mathcal{M}_{\mathcal{T} \cup \{\sigma, \beta\}} = \emptyset$ , the negotiation ends immediately because, *sic stantibus rebus*, it is impossible to reach an agreement. If the participants are willing to avoid the *conflict deal* [8], and continue the negotiation, it will be necessary they revise their strict requirements.

In the negotiation process both the buyer and the seller set some preferences on attributes, or combination of attributes, they want to negotiate on. The utility function is usually defined based on these attributes. We start defining buyer's and seller's preferences and the associated utilities. Obviously, in a negotiation process there are two utility functions:  $u_\beta$  for the buyer, and  $u_\sigma$  for the seller.

**Definition 2 (Preferences)** Let  $\mathcal{A}$  be a set of propositional atoms, and  $\mathcal{T} \subset \mathcal{L}_A$ .

- the buyer's negotiation preferences  $\mathcal{B} \doteq \{\beta_1, \dots, \beta_k\}$  are a set of propositional formulas in  $\mathcal{L}_A$ , each of them representing the subject of a buyer's preference, and a utility function  $u_\beta : \mathcal{B} \rightarrow \mathbb{R}^+$  assigning a utility to each formula, such that  $\sum_i u_\beta(\beta_i) = 1$ .
- analogously, the seller's negotiation preferences  $\mathcal{S} \doteq \{\sigma_1, \dots, \sigma_h\}$  are a set of propositional formulas in  $\mathcal{L}_A$ , and a utility function  $u_\sigma : \mathcal{S} \rightarrow \mathbb{R}^+$  such that  $\sum_j u_\sigma(\sigma_j) = 1$ .

Buyer's/seller's preferences are used to evaluate how good is a possible agreement and to select the best one. As usual, both agents' utilities are normalized to 1 to eliminate outliers, and make them comparable. Since we assumed that utilities are additive, the *global utility* is just a sum of the utilities of preferences satisfied in the agreement. To lower the number of symbols, we use  $u_\beta$  and  $u_\sigma$  also to denote the functions computing the global utility of agents over an agreement.

**Definition 3 (Global Utilities)** Let  $\mathcal{B}$  and  $\mathcal{S}$  be respectively the buyer's and seller's preferences, and  $\mathcal{M}$  be their agreements set. The global utility of an agreement  $m \in \mathcal{M}$  for a buyer and a seller, respectively, are defined as:

$$u_\beta(m) \doteq \Sigma\{u_\beta(\beta_i) \mid m \models \beta_i\}, \quad u_\sigma(m) \doteq \Sigma\{u_\sigma(\sigma_i) \mid m \models \sigma_i\}$$

where  $\Sigma\{\dots\}$  stands for the sum of all elements in the set.

Based on the previous definitions it is possible now to define a Multi-issue Bilateral Negotiation problem. The only other elements we still need to introduce are the *disagreement thresholds*, also called disagreement payoffs, or reservation values,  $t_\beta, t_\sigma$ , that are the minimum utility that each agent requires to pursue a deal. They may incorporate an agent's attitude toward concluding the transaction, but also overhead costs involved in the transaction itself, e.g., fixed taxes.

**Definition 4 (MBN)** Given a demand  $\beta$  and a set of buyer's preferences  $\mathcal{B}$  with utility function  $u_\beta$ , a supply  $\sigma$  and a set of seller's preferences  $\mathcal{S}$  with utility function  $u_\sigma$ , a propositional theory  $\mathcal{T}$  and two disagreement thresholds  $t_\beta, t_\sigma$ , a **Multi-issue Bilateral Negotiation problem (MBN)** is finding a model  $m$  (agreement) such that  $m \models \mathcal{T} \cup \{\sigma, \beta\}$  and both  $u_\beta(m) \geq t_\beta$  and  $u_\sigma(m) \geq t_\sigma$ .

Observe that not every agreement  $m$  is a solution of an MBN, if either  $u_\sigma(m) < t_\sigma$  or  $u_\beta(m) < t_\beta$ . Such an agreement represents

a deal which, although satisfying strict requirements, it is not worth the transaction effort.

Obviously among all possible agreements that we can compute given a theory  $\mathcal{T}$ , we are interested in agreements that are Pareto-efficient, so we define:

**Definition 5 (Pareto)** A Pareto agreement for an MBN is an agreement  $m$ , such that for no agreement  $m'$  both  $u_\beta(m') \geq u_\beta(m)$  and  $u_\sigma(m') \geq u_\sigma(m)$ , with at least one strict inequality.

We can also define the search of particular agreements for an MBN, maximizing the welfare  $u_\beta + u_\sigma$ , or maximizing the product  $u_\beta \cdot u_\sigma$

**Definition 6** Given an MBN, we define the following problems:

- MAX-SUM-MBN is the problem of finding an agreement  $m$  for which  $u_\sigma(m) + u_\beta(m)$  is maximal
- MAX-PROD-MBN is the problem of finding an agreement  $m$  for which  $u_\sigma(m) \cdot u_\beta(m)$  is maximal

Clearly, every solution for MAX-SUM-MBN and MAX-PROD-MBN is also a Pareto agreement, but not vice versa.

## 5 Computational issues

In this section we start by proving membership in NPO [1] of MAX-SUM-MBN and MAX-PROD-MBN. Then, we outline how to compute solutions for MAX-SUM-MBN and MAX-PROD-MBN through general combinatorial optimization techniques. Finally, we prove NPO-hardness of both problems, that is, tailored algorithms yielding solutions approximated within a polynomial bound from the optimum are unlikely to exist.

**Theorem 1** Given an instance of MBN, deciding whether it has a solution can be checked in nondeterministic polynomial-time, even if a lower bound  $k > 0$  is given on either the sum or the product of  $u_\sigma, u_\beta$ .

*Proof.* Simply guess an interpretation  $m$ , and check in polynomial time that it is a model of  $\mathcal{T} \cup \{\sigma, \beta\}$  and that  $u_\sigma(m) \geq t_\sigma$ ,  $u_\beta(m) \geq t_\beta$ . If a lower bound  $k$  is given, also checking that either  $u_\sigma(m) + u_\beta(m) \geq k$  or  $u_\sigma(m) \cdot u_\beta(m) \geq k$  can be done in polynomial time.  $\square$

Hence both MAX-SUM-MBN and MAX-PROD-MBN belong to NPO. We now outline how an actual solution to both problems can be found by Integer Linear Programming (ILP).

Regarding MAX-SUM-MBN, proceed as follows. First of all, let  $\{B_1, \dots, B_k, S_1, \dots, S_h\}$  be  $k + h$  new propositional atoms, and let  $\mathcal{T}' = \mathcal{T} \cup \{B_i \equiv \beta_i \mid i = 1, \dots, k\} \cup \{S_j \equiv \sigma_j \mid j = 1, \dots, h\}$  – that is, every preference in  $\mathcal{B} \cup \mathcal{S}$  is equivalent to a new atom in  $\mathcal{T}'$ . Then, by using standard transformations in clausal form [11] obtain a set of clauses  $\mathcal{T}''$  which is satisfiable iff  $\mathcal{T}' \cup \{\sigma, \beta\}$  does. Then, use a well-known encoding of clauses into linear disequations (e.g., [14, p.314]) so that every solution of the disequations identifies a model of  $\mathcal{T}''$ . Let  $\{b_1, \dots, b_k\}$  the (0,1)-variables one-one with  $\{B_1, \dots, B_k\}$  and similarly  $\{s_1, \dots, s_h\}$ . Then, maximize the linear function

$$\sum_{i=1}^k b_i u_\beta(\beta_i) + \sum_{j=1}^h s_j u_\sigma(\sigma_j)$$

Regarding MAX-PROD-MBN, let  $\{B_{ij} \mid i = 1, \dots, k \ j = 1, \dots, h\}$  be a set of  $k \cdot h$  new propositional atoms. Let  $\mathcal{T}' = \mathcal{T} \cup \{B_{ij} \equiv$

$\beta_i \wedge \sigma_j \mid i = 1, \dots, k, j = 1, \dots, h$  – that is, this time every pair of conjoined preferences is equivalent to a new atom. Then, transform  $\mathcal{T}'$  into a set of clauses  $\mathcal{T}''$ , code clauses as disequations and maximize

$$\sum_{i=1}^k \sum_{j=1}^h b_{ij} u_{\beta}(\beta_i) u_{\sigma}(\sigma_j)$$

Hence, also MAX-PROD-MBN can be solved by ILP as a problem whose size is at most quadratic in the size of the original MAX-PROD-MBN.

When an optimization problem is NP-complete, one can ask whether there exists an algorithm approximating the optimum. For instance, it is known that the following problem MAX-WEIGHTED-SAT is not approximable [1] within any polynomial.

**Definition 7** MAX-WEIGHTED-SAT is the following NPO-complete problem: given set of atoms  $\mathcal{A}$ , a propositional formula  $\varphi \in \mathcal{L}_{\mathcal{A}}$  and a weight function  $w : \mathcal{A} \rightarrow \mathbb{N}$ , find a truth assignment satisfying  $\varphi$  such that the sum of the weights of true variables is maximum.

If MAX-SUM-MBN or MAX-PROD-MBN admitted some approximation algorithm, using general ILP for solving them might be an overshoot. However, the theorem below proves that this is not the case, by giving an L-reduction [1] of MAX-WEIGHTED-SAT to both problems.

**Theorem 2** MAX-SUM-MBN and MAX-PROD-MBN are NPO-complete problems, even if  $\mathcal{T}$  is in 3CNF and both  $\mathcal{B}$  and  $\mathcal{S}$  are sets of positive literals.

*Proof.* Let  $W = \langle \varphi, w \rangle$  be an instance of MAX-WEIGHTED-SAT, with  $\varphi$  in 3CNF. Define an instance  $M_+$  of MAX-SUM-MBN as follows. Let  $\mathcal{T} = \varphi$ ,  $\mathcal{B} = \{A_1, \dots, A_k\}$  with any  $k = 1, \dots, n-1$ , and let  $\mathcal{S} = \{A_{k+1}, \dots, A_n\}$ . Moreover, let  $u_{\beta}(A_i) = w(A_i)$  for  $i = 1, \dots, k$  and  $u_{\sigma}(A_j) = w(A_j)$  for  $j = k+1, \dots, n$ . Finally, let  $t_{\beta} = t_{\sigma} = 0$ . Clearly, every solution for  $W$  is also a solution for  $M_+$ , and for every model  $m$ , the value of the objective function  $u_{\beta}(m) + u_{\sigma}(m)$  is the same as the one for  $W$ . Hence, the above is an L-reduction with  $\alpha = \beta = 1$ .

Similarly, define an instance  $M_{\times}$  of MAX-PROD-MBN as follows. Let  $\mathcal{T} = \varphi$ ,  $\mathcal{B} = \mathcal{A}$ ,  $\mathcal{S} = \{A_0\}$  where  $A_0$  is a new atom not in  $\mathcal{A}$ . Moreover, let  $u_{\beta}(A_i) = w(A_i)$  for  $i = 1, \dots, n$  and  $u_{\sigma}(A_0) = 1$ . Finally, let  $t_{\beta} = t_{\sigma} = 0$ . Also in this case, every solution for  $W$  satisfies also  $\mathcal{T}$ , and for every model  $m$  the objective function  $u_{\beta}(m) \cdot u_{\sigma}(m)$  of  $M_{\times}$  has the same value as  $W$ 's one. Hence, also the above reduction is an L-reduction with  $\alpha = \beta = 1$ .  $\square$

The proof of the above theorem highlights the fact that a source of complexity for MAX-SUM-MBN and MAX-PROD-MBN comes from the inherent conflicts between the preferences of a single agent.

## 6 The bargaining process

We now present the bargaining process, which covers the following phases:

**Preliminary Phase.** The buyer defines  $\beta$  and  $\mathcal{B}$ , as well as the threshold  $t_{\beta}$ , and the seller defines  $\sigma$ ,  $\mathcal{S}$  and  $t_{\sigma}$ . In this paper we are not interested in how to compute these values; we assume they are determined in advance by means of either direct assignment methods (Ordering, Simple Assessing or Ratio Comparison) or pairwise comparison methods (like AHP and Geometric Mean) [15]. Both agents inform the facilitator about these specifications and the theory  $\mathcal{T}$  they refer to.

**Initial phase.** The facilitator computes  $\mathcal{M}_{\mathcal{T} \cup \{\sigma, \beta\}}$ . If  $\mathcal{M}_{\mathcal{T} \cup \{\sigma, \beta\}} = \emptyset$  the facilitator requests the buyer and the seller to refine –respectively– their  $\beta$  and  $\sigma$  and recomputes  $\mathcal{M}_{\mathcal{T} \cup \{\sigma, \beta\}}$ . Once/if buyer and seller solve any conflicts about reservation requirements, the negotiation can start.

**Negotiation-Core phase.** The facilitator proposes one of the two agreements mutually beneficial for both the buyer and the seller: an agreement which maximizes the social welfare and one which maximizes the product of utilities. If either one of the participants rejects it, then the facilitator proposes the second one. Since both of the agreements are Pareto-efficient, the order they are proposed to the participants is not relevant. From this point on, it is a *take-it-or-leave-it* offer, because the participants can either accept or reject it [7]. Notice that since both the proposed deals are Pareto-efficient, although the participants can move from the proposed deals, they might end in a less (no-Pareto) efficient agreement. Once the facilitator computes the space of all possible deals  $\mathcal{M}_{\mathcal{T} \cup \{\sigma, \beta\}}$ , it evaluates the utility of each model  $m \in \mathcal{M}_{\mathcal{T} \cup \{\sigma, \beta\}}$ , both for the buyer –  $u_{\beta}(m)$  – and for the supplier –  $u_{\sigma}(m)$ .

Among all possible models, it chooses  $\bar{m}$  maximizing the product  $u_{\beta}(\bar{m}) * u_{\sigma}(\bar{m})$  and  $\underline{m}$  maximizing the sum. The facilitator proposes such solutions to the opponents. If one of the opponents, or both of them, refuse the deals proposed by the facilitator, they reach a *conflict deal* and then they will likely look for other partners in the e-marketplace.

## 7 An illustrative example

In the following we illustrate our logic based approach to automate negotiation process with the aid of a simple example, which refers to an e-marketplace related to the automotive domain. Buyer and seller express their demand/supply and preferences. According to the previous definitions, a simple instance of an MBN can be formalized as follows:

$$\begin{aligned} \beta &= \text{StationWagon} \wedge \text{AirConditioning} \wedge \text{DriverAirbag} \wedge \\ &\quad \text{SatelliteAlarm} \\ \beta_1 &= \neg \text{ExternalColorRed} \wedge \text{ExternalColorBlack} \wedge \\ &\quad \text{InteriorColorBlack} \\ \beta_2 &= \text{Gasoline} \vee \text{Diesel} \\ u_{\beta}(\beta_1) &= 0.6 \\ u_{\beta}(\beta_2) &= 0.4 \\ t_{\beta} &= 0.3 \\ \\ \sigma &= \text{StationWagon} \wedge \text{Diesel} \\ \sigma_1 &= (\text{ExternalColorBlack} \vee \text{ExternalColorBlue}) \wedge \\ &\quad (\text{InteriorColorGrey} \vee \text{InteriorColorBlack}) \\ \sigma_2 &= \text{DriverAirbag} \wedge \neg \text{PassengerAirbag} \\ \sigma_3 &= \text{AlarmSystem} \wedge \neg \text{SatelliteAlarm} \\ u_{\sigma}(\sigma_1) &= 0.2 \\ u_{\sigma}(\sigma_2) &= 0.3 \\ u_{\sigma}(\sigma_3) &= 0.5 \\ t_{\sigma} &= 0.2 \end{aligned}$$

In this example the thresholds are set considering some transaction costs like *e.g.*, operating costs for the seller and cost of taxes for the buyer.

$$\mathcal{T} = \left\{ \begin{array}{l} \text{ExternalColorBlue} \Rightarrow \neg \text{ExternalColorBlack} \\ \text{ExternalColorBlue} \Rightarrow \neg \text{ExternalColorRed} \\ \text{ExternalColorBlack} \Rightarrow \neg \text{ExternalColorRed} \\ \text{InteriorColorBlack} \Rightarrow \neg \text{InteriorColorGrey} \\ \text{Gasoline} \Rightarrow \neg \text{Diesel} \\ \text{SatelliteAlarm} \Rightarrow \text{AlarmSystem} \\ \text{PassengerAirbag} \Rightarrow \text{FrontAirbag} \\ \text{DriverAirbag} \Rightarrow \text{FrontAirbag} \end{array} \right.$$

ExBlue	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
ExBlack	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
ExRed	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	
InBlack	1	1	0	0	0	0	1	1	0	0	0	0	0	0	0	0	0	0	0	1	1	0	0	0	0	0	0	0	0	
InGrey	0	0	1	1	0	0	0	0	0	1	1	0	0	0	0	1	1	0	0	0	0	0	0	1	1	0	0	0	0	
Gas	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
Dies	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
SatAlarm	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
Alarm	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
PassAir	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0
FrontAir	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
DrivAir	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
StatWag	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
AirCond	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
$\beta_1$	0	0	0	0	0	0	0	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
$\beta_2$	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
$u_\beta(m)$	0.4	0.4	0.4	0.4	0.4	0.4	0.4	1	1	0.4	0.4	0.4	0.4	0.4	0.4	0.4	0.4	0.4	0.4	0.4	0.4	0.4	0.4	0.4	0.4	0.4	0.4	0.4	0.4	
$\sigma_1$	1	1	1	1	0	0	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
$\sigma_2$	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	
$\sigma_3$	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
$u_\sigma(m)$	0.2	0.5	0.2	0.5	0	0.3	0.2	0.5	0.2	0.5	0	0.3	0	0.3	0	0.3	0	0.3	0	0.3	0	0.3	0	0.3	0	0.3	0	0.3		
SUM	0.6	0.9	0.6	0.9	0.4	0.7	1.2	1.5	0.6	0.9	0.4	0.7	0.4	0.7	0.4	0.7	0.4	0.7	0.4	0.7	0.4	0.7	0.4	0.7	0.4	0.7	0.4	0.7		
PROD	0.08	0.2	0.08	0.2	0	0.12	0.2	0.5	0.08	0.2	0	0.12	0	0.12	0	0.12	0	0.12	0	0.12	0	0.12	0	0.12	0	0.12	0	0.12		

**Table 1.**  $\mathcal{M}_{\mathcal{T}\cup\{\sigma,\beta\}}$  and the utility associated to each model. In bold style the values chosen by the facilitator maximizing the sum and the product.

Table 1 shows all the models in the set  $\mathcal{M}_{\mathcal{T}\cup\{\sigma,\beta\}}$ . As it is shown in Table 1, after computing  $u_\beta(m)$  and  $u_\sigma(m)$  (rows 18 and 23) the sum and the product of the two utilities are calculated. The facilitator proposes to the participants (buyer and seller) the deal that maximizes the product and the sum of utilities (in this case the same). In the final agreement we have the following propositional atoms set true:

$m = \overline{m}$  = **StationWagon**, **AirConditioning**, **DriverAirbag**, **SatelliteAlarm**, **Diesel**, AlarmSystem, FrontAirbag, ExternalColorBlack, InteriorColorBlack

where bold style atoms highlight strict attributes, while the others are preferences. The utilities related to  $\overline{m}$  are:

$$u_\beta(\overline{m}) = u_\beta(\beta_1) + u_\beta(\beta_2) = 0.6 + 0.4 = 1$$

$$u_\sigma(\overline{m}) = u_\sigma(\sigma_1) + u_\sigma(\sigma_2) = 0.2 + 0.3 = 0.5$$

while product and sum are:

$$u_\beta(\overline{m}) * u_\sigma(\overline{m}) = 1 * 0.5 = 0.5$$

$$u_\beta(\overline{m}) + u_\sigma(\overline{m}) = 1 + 0.5 = 1.5$$

## 8 Conclusion and Future Work

In this paper we proposed a logic based approach to automate multi issue bilateral negotiation (MBN), defining a protocol which leads to a Pareto-efficient deal. The relations among issues are represented through a logic theory  $\mathcal{T}$  and preferences are elicited by assigning utility to formulas, and then considering bundles of interrelated issues. Computational issues were also investigated: we proved that MAX-SUM-MBN and MAX-PROD-MBN are NPO-complete problems. In the near future, we plan to extend the approach using Description Logics, to cope with greater expressiveness in demand/supply descriptions. We are also investigating other negotiation protocols –also different from one-shot ones– still suitable for e-commerce applications. A prototype is also being implemented to validate the approach through large scale experiments.

## REFERENCES

- [1] G. Ausiello, P. Crescenzi, V. Kann, G. Gambosi, and A. M. Spaccamela. *Complexity and Approximation: Combinatorial Optimization Problems and Their Approximability Properties*. Springer, 2003.
- [2] K. Binmore. *Fun and Games. A Text on Game Theory*. D.C. Heath and Company, 1992.
- [3] Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet. Negotiating over small bundles of resources. In *Proc. of AAMAS '05*, pages 296–302, 2005.
- [4] P. E. Dunne, M. Wooldridge, and M. Laurence. The complexity of contract negotiation. *Artif. Intell.*, 164(1-2):23–46, 2005.
- [5] U. Endriss, N. Maudet, F. Sadri, and F. Toni. On optimal outcomes of negotiations over resources. In *Proc. of AAMAS '03*, pages 177–184, 2003.
- [6] E. H. Gerding, D. D. B. van Bragt, and J. A. L. Poutre. Scientific approaches and techniques for negotiation: a game theoretic and artificial intelligence perspective. Technical report, SEN-R0005, CWI, 2000.
- [7] N. Jennings, P. Faratin, A. Lomuscio, S. Parsons, M. Wooldridge, and C. Sierra. Automated negotiation: prospects, methods and challenges. *Int. J. of Group Decision and Negotiation*, 10(2):199–215, 2001.
- [8] J.S. Rosenschein and G. Zlotkin. *Rules of Encounter*. MIT Press, 1994.
- [9] C. Li, J. Giampapa, and K. Sycara. A review of research literature on bilateral negotiations. Technical report, Carnegie Mellon University, 2003.
- [10] A. R. Lomuscio, M. Wooldridge, and N. R. Jennings. A classification scheme for negotiation in electronic commerce. *Int Journal of Group Decision and Negotiation*, 12 (1):31–56, 2003.
- [11] D. W. Loveland. *Automated theorem proving: A logical basis*. North-Holland, 1978.
- [12] A. Muthoo. On the strategic role of outside options in bilateral bargaining. *Operations Research*, 43(2):292–297, 1995.
- [13] J. F. Nash. The bargaining problem. *Econometrica*, 18 (2):155–162, 1950.
- [14] C. H. Papadimitriou and K. Steiglitz. *Combinatorial optimization: algorithms and complexity*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1982.
- [15] J. Pomerol and S. Barba-Romero. *Multicriterion Decision Making in Management*. Kluwer Series in Operation Research. Kluwer Academic, 2000.
- [16] A. Ragone, T. Di Noia, E. Di Sciascio, and F. M. Donini. Propositional-logic approach to one-shot multi issue bilateral negotiation. *ACM SIGecom Exchanges*, 5(5):11–21, 2006.
- [17] H. Raiffa. *The Art and Science of Negotiation*. Harvard University Press, 1982.
- [18] H. Raiffa, J. Richardson, and D. Metcalfe. *Negotiation Analysis - The Science and Art of Collaborative Decision Making*. The Belknap Press of Harvard University Press, 2002.
- [19] A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50:97–109, 1982.
- [20] T. Sandholm. Contract types for satisficing task allocation: I theoretical results. In *Proceedings of the AAAI Spring Symposium*, 1998.
- [21] D. Trastour, C. Bartolini, and C. Priest. Semantic Web Support for the Business-to-Business E-Commerce Lifecycle. In *Proc. WWW '02*, pages 89–98, 2002.
- [22] M. Wooldridge and S. Parsons. Languages for negotiation. In *Proc of ECAI '04*, pages 393–400, 2000.