

Multi-player Multi-issue Negotiation with Mediator using CP-nets

Thiri Haymar Kyaw¹, Sujata Ghosh² and Rineke Verbrugge³

¹University of Technology (Yatanarpon Cyber City), Pyin Oo Lwin, Myanmar

²Indian Statistical Institute, SETS Campus, MGR Knowledge City, CIT Campus, Chennai 600 113, India

³Institute of Artificial Intelligence, University of Groningen, PO Box 407, 9700 AK Groningen, The Netherlands

Keywords: Multi-agent Negotiation, Multi-issue Negotiation, Conditional Preference Networks, Incomplete Information.

Abstract: This paper presents a simple interactive negotiation approach for conflicts in everyday life with incomplete information. We focus on mediation to obtain an agreement while going through alternating offers over a finite time bargaining game. The mediator searches and proposes a jointly optimal negotiation text for all players participating in the negotiation process based on their conditional preference networks (CP-nets). The players make a decision whether to accept or reject by examining their utility CP-nets. We develop two algorithms for the mediator and the players. If the first negotiation text cannot be accepted by all players, the mediator offers the next negotiation texts by searching for jointly optimal solutions. This negotiation process continues until an agreement is achieved or a deadline is reached. This proposed approach can support multi-issue, multi-party negotiation to achieve an agreement during a finite number of rounds with near optimal outcomes.

1 INTRODUCTION

Negotiation occurs in several areas of real-world problems: personal cases such as marriage, divorce, and parenting; business cases such as pricing between seller and buyer and sharing a market between organizations; international crisis cases like the Cuba missile crisis, the North Korean crisis, and Copenhagen climate change control. Negotiation is a process for agents to communicate and compromise in order to reach beneficial agreements. In such situations, the agents have a common interest in cooperating, but have conflicting interests over exactly how to cooperate (Fatima et al., 2005).

Bargaining is a simple form of a negotiation process. It is used to establish a price to trade a fixed and defined commodity between seller and buyer. One party usually attempts to gain advantage over another to obtain the best possible agreement. Splitting a pie between two players is a simple bargaining example. In such games with many periods of offers and counteroffers, strategies are not just actions, but rather ways for choosing actions based on the actions chosen by both agents in earlier periods (Rasmusen, 2007).

In competitive bargaining, the process is viewed as a competition that is to be won or lost. Positional bargaining is a negotiation strategy that involves holding on to a fixed idea, or position, of what you want

and arguing for it and it alone, regardless of any underlying interests. The classic example of positional bargaining is the haggling that takes place between proprietor and customer over the price of an item. The customer has a maximum amount she will pay and the proprietor will only sell something for a price above a certain minimum amount. Each side starts with an extreme position, which in this case is a monetary value, and proceeds from there to negotiate and make concessions. Eventually a compromise may be reached. A position is usually determined by the interests of a negotiating party, and reflected in a contract that it puts forward to its counterpart.

Integrative bargaining (also called “interest-based bargaining” or “win-win bargaining”) is a negotiation strategy in which parties collaborate to find a “win-win” solution to accommodate their different interests. This strategy focuses on developing mutually beneficial agreements based on the interests of the disputants. Interests include the needs, desires, concerns, and fears important to each side (Raiffa et al., 2002). Integrative bargaining usually produces more satisfactory outcomes for the players involved than does positional bargaining. Positional bargaining is based on fixed, opposing positions and tends to result in a compromise or no agreement at all. Our negotiation approach focuses on integrative bargaining for achieving a satisfactory agreement for all players.

Interest-based negotiation either can get the parties to an agreement point where they can bargain or even better, to a point where they do not need to bargain at all. Interest-based negotiation typically entails two or more issues to be negotiated. It involves an agreement process that better integrates the aims and goals of all the negotiating parties through creative and collaborative problem solving.

Mediation usually consists of a negotiation process that employs a mutually agreed upon third party to settle a dispute between negotiating parties in order to find a compatible agreement to resolve disputes (Fisher, 1978). In negotiation, the parties agree to work with each other to resolve a dispute. In mediation, the parties agree to work with a facilitator or mediator to resolve a conflict. In many cases, international negotiations aim to achieve an agreement on various issues between multiple parties. “Camp David” is an interesting example of a negotiation that happened between Egypt and Israel in 1978, resulting in a more or less successful agreement with the help of a mediator, the United States.

In this paper, we propose a simple mediated approach for multi-issue negotiation with incomplete information based on adjusting the players’ preferences. In this approach, the mediator searches for a jointly optimal negotiation text for all players in their conditional preference networks (CP-nets), using a depth-first search-based algorithm. The players use utility-based CP-nets to make a decision for agreement. The purpose of the paper is to achieve near optimal joint preference for all players while each player has imperfect information about his opponents.

The rest of the paper is organized as follows. In Section 2, we briefly explain preliminaries about CP-nets, which are conditional preference networks for representing and reasoning with qualitative preferences. We also discuss utility CP-nets and mediation approaches using a single negotiation text in Section 2. In Section 3, we describe the proposed negotiation approach with algorithms and illustrations. We discuss other closely related approaches to negotiation based on CP-nets and compare our approach to them in Section 4. Finally, in Section 5, we conclude the paper and mention some future directions.

2 PRELIMINARIES

We begin with background concepts of conditional preference networks (CP-nets), their induced preference graphs and the utility-based CP-nets in this section. We will also discuss a mediation approach using single negotiation text (SNT).

2.1 CP-nets and UCP-nets

Boutilier and colleagues introduced CP-nets as a graphical representation of conditional preference networks that can be used for specifying preference relations in a relatively compact, intuitive, and structured manner using conditional ceteris paribus (all other things being equal) preference statements (Boutilier et al., 2004a; Boutilier et al., 2004b). CP-nets can be used to specify different types of preference relations, such as a preference ordering over potential decision outcomes or a likelihood ordering over possible states of the world.

CP-nets are similar to Bayesian networks (Pearl, 1988). Both utilize directed graphs; however, the aim of CP-nets in using the graph is to capture statements of qualitative conditional preferential independence. A CP-net over variables $V = X_1, \dots, X_m$ is a directed graph G over X_1, \dots, X_m whose nodes are annotated with conditional preference tables $CPT(X_i)$ for each $X_i \in V$. Each conditional preference table $CPT(X_i)$ associates a total order \succ_i^u with each instantiation u of X_i ’s parents $Pa(X_i) = U$ (Boutilier et al., 2004a).

Let $V = X_1, \dots, X_m$ be a demand set of m attributes; $X_i \in V$ ($i = 1$ to m). $D(X_i)$ is the domain of X_i and is represented as $D(X_i) = x_1, \dots, x_n$. There are $D(X_1) \times D(X_2) \times \dots \times D(X_m)$ possible alternatives (outcomes), denoted by O . Elements of O are denoted by o, o', o'' etc. and represented by concatenating the values of the variables (Li et al., 2011b). For example, if $V = \{A, B, C\}$, $D(A) = \{a_1, a_2, a_3\}$, $D(B) = \{b_1, b_2\}$ and $D(C) = \{c_1, c_2, c_3\}$, then the assignment $a_2b_2c_1$ assigns a_1 to variable A , b_2 to B and c_1 to C .

The preference information captured by an acyclic CP-net N can be viewed as a set of logical assertions about a user’s preference ordering over complete assignments to variables in the network. These statements are generally not complete, that is, they do not determine a unique preference ordering. Those orderings consistent with N can be viewed as possible models of the user’s preferences, and any preference assertion that holds in all such models can be viewed as a consequence of the CP-net (Boutilier et al., 2004b).

The set of consequences $o \succ o'$ of an acyclic CP-net constitutes a partial order over outcomes: o is preferred to o' in this ordering iff $N \models o \succ o'$. This partial order can be represented by an acyclic directed graph, referred to as the induced preference graph:

- The nodes of the induced preference graph correspond to the complete assignments to the variables of the network; and
- There is an edge from node o' to node o if and only if the assignments at o' and o differ only in the value of a single variable X , and given the values

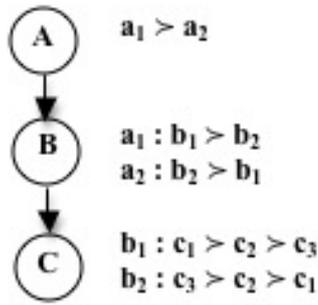


Figure 1: CP-net, for Player 1.

assigned by o' and o to $Pa(X)$, the value assigned by o to X is preferred to the value assigned by o' to X .

For example, consider the CP-net given in Figure 1, whose variables are A , B and C . The preference statements mean that a_1 is strictly preferred to a_2 while the preferences of variable B are conditioned on the variable A . If a_1 is chosen, then b_1 is preferred to b_2 and if a_2 is chosen, b_2 is preferred to b_1 . The preferences of variable C are also conditioned on the variable B . The preference graph induced by the CP-net of Figure 1 is shown in Figure 2.

The concept of Utility CP-net (UCP-net) was also introduced by Boutilier and colleagues. It extends the concept of CP-net by allowing quantification over nodes with conditional utility information. Semantically, Boutilier et al. treat the different factors $V = X_1, \dots, X_m$ as *generalized additive independent* of one another for an underlying utility function u (Boutilier et al., 2001); this means intuitively that the expected value of u is not affected by correlations between the variables, and implies that u can be decomposed as a sum of factors over each set of variables X_i (Bacchus and Grove, 1995). For example, the CP-net in Figure 1 can be extended with utility information by including a factor (f) for each variable in the network, specifically $f_1(A)$, $f_2(A, B)$ and $f_3(B, C)$ as shown in Figure 3. We calculate the total utility of preference strings as follows: $u(A, B, C) = f_1(A) + f_2(A, B) + f_3(B, C)$. Each of these factors is quantified by the *CPT* (conditional preference tables) in the network. For example, the utility of $a_1b_1c_1$ is as follows: $u(a_1, b_1, c_1) = f_1(a_1) + f_2(a_1, b_1) + f_3(b_1, c_1) = 5 + 0.6 + 0.6 = 6.2$ according to Figure 3.

F. Rossi and colleagues presented an extension of the CP-net, called mCP-nets (Rossi et al., 2004), to model the qualitative and conditional preferences of multiple agents. They allowed the individual agents to vote to obtain mCP-nets by combining several partial CP-nets. K.R. Apt and colleagues proposed an approach for analyzing strategic games that can be used to study CP-nets (Apt et al., 2005). They in-

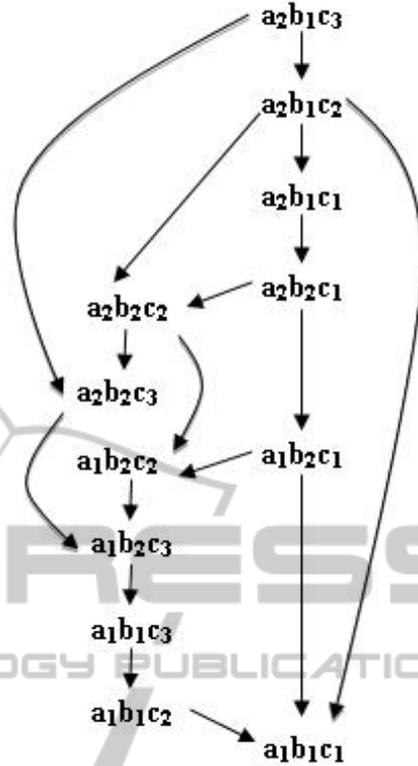


Figure 2: Induced Preference Graph of the CP-net of Figure 1, for Player 1.

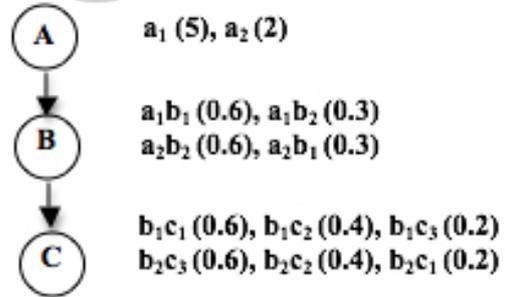


Figure 3: UCP-net corresponding to the CP-net of Figure 1, for Player 1.

troduced a generalization of strategic games in which each player has to his disposal a strict preference relation on his set of strategies, parameterized by a joint strategy of his opponents. They showed that optimal outcomes in CP-nets are Nash equilibria of strategic games with parameterized preferences. Z. Liu and colleagues also focused on the relationship between CP-nets and strategic games (Liu et al., 2012). They proposed a solution to resolve the optimal outcomes of CP-nets by transforming a CP-net to a game tree and using a tree algorithm to find Nash equilibria.

2.2 Mediation using a Single Negotiation Text (SNT)

The concept of a single negotiation text (SNT) was suggested as a mediation device by Roger Fisher (Fisher, 1978). SNT is often employed in international negotiations, especially with multi-party negotiations (Raiffa et al., 2002), (Korhonen et al., 1995) and (Ehtamo et al., 2001). For example, the SNT approach was applied by the United States in mediating the Egyptian–Israeli conflict, which is known as the Camp David Negotiations (Raiffa 1982). During an SNT negotiation, a mediator first devises and proposes a deal (SNT-1) for the two protagonists to consider. The first proposal is not intended as the final agreement. It is meant to serve as an initial, single negotiating text: a version to be privately criticized by both sides and then modified in an iterative manner.

The SNT is utilized as a method of focusing the parties' attention on the same composite text (Raiffa et al., 2002). The important aspect of the process is that it appears to be fair to both sides, and not divisive. Based upon the criticisms by the parties, the mediator prepares another proposal, which is not perfect, but which improves both parties' positions. Again, both parties provide suggestions on improving the proposal, and this new proposal is again criticized by the parties. This process continues until all the issues are settled and the final agreement is achieved or it is clear that no agreement is achievable. P. Korhonen and colleagues discussed the importance of the starting point of the single negotiation text (Korhonen et al., 1995). They argue that, if the path taken in subsequent steps does not compensate for a biased starting point, the bias will have considerable impact on the final outcome of the negotiations.

3 THE PROPOSED APPROACH TO NEGOTIATION

In real-world negotiations, negotiators need to achieve an agreement on multiple issues with multiple players. Sometimes, a mediator is included to facilitate the negotiation process. Some negotiations fail because the parties have too many conflicts and they cannot work with each other. Therefore, a mediator may be used if the parties prefer a third party who is neutral and does not represent any party's interests. Also in situations where the parties cannot meet to negotiate directly, a mediator may be needed.

Our approach is based on a natural way to negotiate in the real world. The proposed framework con-

sists of two types of individuals: the mediator and the players. All players and the mediator specify the issues that they need to negotiate before the negotiation process starts. Each player keeps his own private information and he does not know his opponents' private information. Each player reports his partial CP-nets to the mediator. They do not directly come to know their opponents' preferences at any stage. In addition, each player defines his own utility values for each attribute and calculates his total utility for all combinations of variables, as we mentioned in Section 2.1. Each player creates his own UCP-net that is used for proposing a maximum preference and for deciding to accept or reject the proposal by the mediator. The mediator seeks to propose a single negotiation text that gains optimal joint outcomes for all players by comparing all players' proposed preferences based on depth-first search (Li et al., 2011b; Li et al., 2011a; Aydogan et al., 2011).

3.1 Case Study: Negotiation with Mediation

As an example case, let us consider three players and one mediator for three issues. In this framework, the players and mediator can be run on different computers. When starting the negotiation process, all players report their overall preference information about negotiation issues to the mediator. Then, the mediator creates induced preference graphs for all players based on their CP-nets. Let N be the set of players: $N = \{1, 2, 3\}$. We consider the three variables A, B, C as three negotiation issues. The domains of the variables are $D(A) = a_1, a_2$; $D(B) = b_1, b_2$; and $D(C) = c_1, c_2, c_3$.

This can be seen as a simple real-world negotiation between different preferences of family members. Suppose that a family including father, mother and 20 years old son decided to buy a new house. They have different preferences for the three issues:

Type of House (A): house with small garden (a_1);
condo apartment (a_2);

Place Near (B): market (b_1); park (b_2);

Price Range (C): high (c_1); medium (c_2); low (c_3).

For example, mother (Player 1) prefers to buy a house with a small garden, situated near a market. If the house is situated near a market, she prefers a high price to a medium or a low price. If it is situated near a park, she prefers a low price to a high price. Father (Player 2) prefers a place near a park to a place near a market. If the place is near a park, he prefers a condominium apartment and if it is near a market, he prefers a house with a small garden. He prefers a high

Table 1: Utility Table for Players 1, 2 and 3.

Strings	Player 1	Player 2	Player 3
$a_1b_1c_1$	6.2	3	7.2
$a_1b_1c_2$	6	2.7	7.4
$a_1b_1c_3$	5.8	2.9	7.6
$a_1b_2c_1$	5.5	5.8	8.6
$a_1b_2c_2$	5.7	5.5	8.4
$a_1b_2c_3$	5.9	5.7	8.2
$a_2b_1c_1$	2.9	2.7	4.6
$a_2b_1c_2$	2.7	2.4	4.4
$a_2b_1c_3$	2.5	2.6	4.2
$a_2b_2c_1$	2.8	6.1	4.2
$a_2b_2c_2$	3	5.8	4.4
$a_2b_2c_3$	3.2	6	4.6
Average	4.35	4.27	6.15

price or a low price rather than a medium price. Their son (Player 3) prefers a house with a small garden to a condo apartment. He also prefers a place situated near a park to a place near a market. If it is a house near a market or an apartment near a park, he prefers a high price to a medium or a low price. Otherwise, he prefers a low price to a medium or a high price.

Let us suppose that the real estate agent acts as a mediator and that the family members do not want to share their preferences with one another. After proposing three negotiation texts by the mediator, all family members agree to buy a new house near a park with low price. We will illustrate the details of the negotiation process in Section 3.2.

3.2 Negotiation Process for Case Study

Assume that the CP-net and the induced preference graph given in Figure 1 and Figure 2 have been proposed by Player 1. The CP-nets and their induced graphs of Player 2 are shown in Figures 4, 5, and those for Player 3 in Figure 7, 8. All players prepare UCP-nets with their private utility values as illustrated in Figure 3, 6 and 9. Each player also calculates the total utility of strings in their UCP-nets, as shown in Table 1. They pick up one string with maximum utility outcomes in their UCP-nets and propose it to the mediator. In this example, Players 1, 2 and 3 propose a_1, b_1, c_1 , a_2, b_2, c_1 and a_1, b_2, c_1 respectively. These are the bottommost strings of the induced graphs (see Figure 2, 5 and 7).

We have developed two algorithms: Algorithm 2 (see page 10) that helps the mediator decide which single negotiation texts to propose, given the players' CP-nets and their answers to previous proposals; and Algorithm 1 (see 10), that helps each of the other players to decide whether to accept a proposal by the mediator or not. We now proceed to show how the

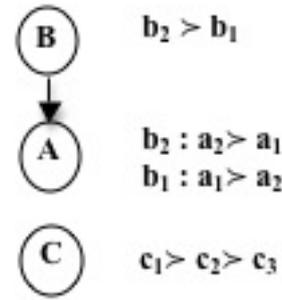


Figure 4: CP-net for Player 2.

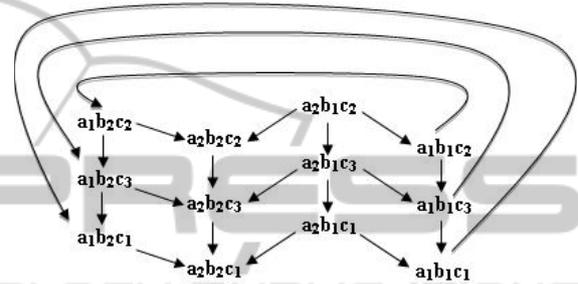


Figure 5: Induced Preference Graph of the CP-Net of Figure 4, for Player 2.

algorithms work for the case study.

The mediator generates a single negotiation text that we call “the proposal”, by searching jointly optimal gains of all players according to Algorithm 2 (see page 10). After receiving the maximum preferred strings of all players, the mediator searches acceptable probability to the other players' strings (line 11, Algorithm 2, see page 10). Starting point is the bottommost string of the induced preference graph and we assume that the string whose edge directly points to the bottommost string gets the acceptable probability 0.9. We define probability of a string on the graph by going backward from the bottommost or maximum preferred string. We count the intermediate edges from the bottommost string to the particular string by reducing the probability by 0.1 for one edge. This searching process continues until the acceptable probability 0.5 is reached. Otherwise, the probability is assigned to zero.

In this case study, the acceptable probability from Player 1's preferred string a_1, b_1, c_1 to Player 2's preferred string a_2, b_2, c_1 is 0.8 and to Player 3's preferred string a_1, b_2, c_1 it is 0.9 on Player 1's induced graph (see Figure 2). For Player 2, probability from his string a_2, b_2, c_1 to Player 1's proposed string a_1, b_1, c_1 is 0.8 and probability to Player 3's proposed string a_1, b_2, c_1 is 0.9 (see Figure 5). For Player 3, probability from his string a_1, b_2, c_1 to Player 1's proposed string a_1, b_1, c_1 is 0.9 and probability to Player 2's proposed string a_2, b_2, c_1 is 0.9 (see Figure 7). Ac-

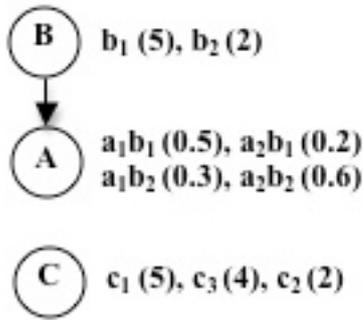


Figure 6: UCP-net for Player 2.

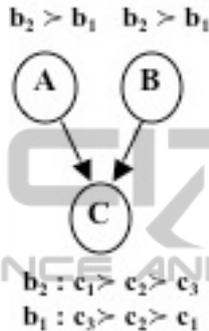


Figure 7: CP-net for Player 3.

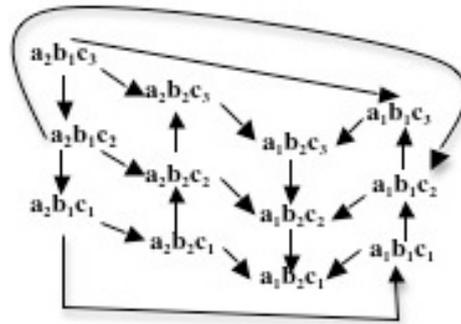


Figure 8: Induced Preference Graph of the CP-Net of Figure 7, for Player 3.

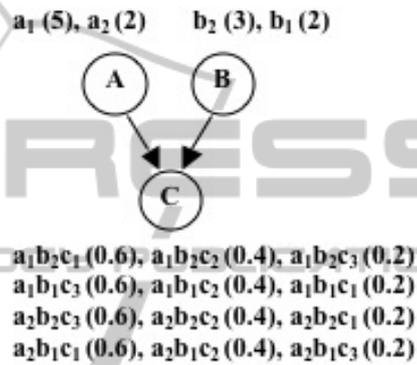


Figure 9: UCP-net for Player 3.

According to this example, a_1, b_2, c_1 is jointly optimal for all players because it obtains reachable probability 0.9 from their maximum preference strings. Therefore, the mediator proposed a_1, b_2, c_1 as a first jointly optimal negotiation text.

Moreover, the mediator marks all acceptable probabilities less than threshold (line 14, Algorithm 2, see page 10). This threshold can be changed when the mediator cannot find any jointly optimal proposal within the threshold. If the players reject the proposal and the mediator has an alternative jointly optimal proposal in his previous marked list, the mediator can use the alternative as the next proposal.

If there is no jointly optimal proposal among the players' proposals, the mediator tries to search for an alternative jointly optimal proposal (line 21, Algorithm 2, see page 10) that has the same acceptable probability of the players' previous proposals. The mediator searches all strings that have one backward edge from the maximum preferred string (probability 0.9) for all players. He then searches a common string of all players. If there is no common string, the mediator continues to search all possible strings for probability 0.8. This process continues until an average maximum probability is found.

After receiving a new negotiation text from the mediator, all players check their utility outcomes (see Table 1) of the text to make a decision "accept"

or "reject". Player 1 rejects the proposal a_1, b_2, c_1 because the difference between her maximum utility and the utility of the current text is greater than the threshold ($maxU - a_1, b_2, c_1 = 6.2 - 5.5 = 0.7$) according to Algorithm 1 (page 10). We assume that the starting threshold is 0.4 in Algorithm 1 (page 10). The threshold can be changed according to the type of utility values. Player 3 strongly accepts the current proposal because he gets his maximum utility. Player 2 also accepts because the difference between his maximum utility and the current text is less than the threshold ($maxU - a_1, b_2, c_1 = 6.1 - 5.8 = 0.3$).

The negotiation process continues until the agreement is achieved by the mediator acting according to Algorithm 2 and the Players according to Algorithm 1. If there is no agreement until the final round before the deadline, the mediator can announce this to all players and the players can evaluate all negotiation texts and consider if they are willing to reduce their maximum utility.

Finally, in this case, the mediator proposes $a_1 b_2 c_3$ as a jointly optimal negotiation text. Player 1 accepts the proposal because the difference between her maximum utility and the current text is less than the threshold; ($maxU - a_1, b_2, c_3 = 6.2 - 5.9 = 0.3$). Player 2 and 3 also accept the proposal because the difference between their maximum utilities and the

current text is equal to threshold; ($\max U - a_1, b_2, c_3 = 6.1 - 5.7 = 0.4$ and $\max U - a_1, b_2, c_3 = 8.6 - 8.2 = 0.4$, respectively). All players achieve a jointly optimal outcome, which is greater than their average outcomes (see Table 1), although they do not achieve their maximum outcomes.

3.3 Cyclic CP-nets

In addition, our approach can achieve a negotiation outcome even when players' CP-nets are cyclic as shown in Figure 10 (a). Let us show a simple example of negotiation between two players. Player 1's preferences and induced graph are shown in Figure 10 (b) and 10 (c). Player 1's private utilities are: a_1b_1 (4), a_2b_2 (4), a_2b_1 (2) and a_1b_2 (2). Player 2's preferences and induced graph are shown in Figure 10 (d) and 10 (e). Player 2's private utilities are: a_1b_1 (3), a_2b_2 (3), a_2b_1 (3) and a_1b_2 (3). Actually, player 2's preference utilities are all the same and its induced graph is not satisfiable (Boutilier et al., 2004a).

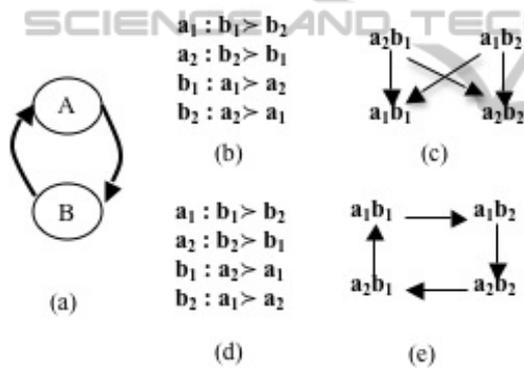


Figure 10: Examples of cyclic CP-nets and their induced preference graphs.

In our negotiation process, player 1 proposes a_1b_1 as her maximum preferred string and player 2 proposes a_2b_2 . When the mediator computes the acceptable probability, a_2b_2 has the same probability as a_1b_1 for Player 1 and a_1b_1 has 0.8, an acceptable probability for Player 2. If there is no backward edge from the maximum preferred string to a particular string, the two strings may have the same probability. For instance, if two strings, a_1b_1 and a_2b_2 (see Figure 10 (c)) have a backward edge from the same string (a_2b_1), then, a_1b_1 and a_2b_2 have the same probability. This reasoning can easily be applied to search on the preference graphs given from the cyclic CP-nets. In our example, the mediator proposes a_2b_2 which is Player 2's proposed string. It also has the same acceptable probability as Player 1's proposed string a_1b_1 . Both players accept the mediator's proposal because it meets their maximum preference utility.

3.4 Negotiating International Conflict

Our approach can be applied to international conflict resolution as well. Camp David is a well-known negotiation process that happened between Egypt and Israel in 1978. The negotiation process lasted for 13 days and the United States acted as a mediator. U.S. mediators had already known deeply about the preferred solutions of Egypt and Israeli and they decided to use a single negotiation text (SNT). The U.S. started by offering its first SNT-1 but was not trying to push this first proposal. It was meant to serve as an initial SNT; a text to be criticized by both sides, then modified, and remodified in an iterative manner. These modifications would be made by the U.S. based on the recommended changes by both sides. After playing six rounds, a satisfactory agreement, the Camp David accord, was reached (Raiffa et al., 2002). We can simulate this negotiation using the proposed approach. In Camp David, there are two players, Egypt and Israel, and four basic negotiation issues (Telhami, 1993; Oakman, 2002):

1. A peace treaty and normalization of relations between Egypt and Israel (A);
2. Demilitarization and removal of Israeli settlements from Sinai (B);
3. Linkage between these issues and the future of the West Bank and Gaza (C);
4. A statement on principles, including Israeli withdrawal from all occupied territories and the right of Palestinians to self-determination (D).

The Camp David negotiation process concerns international issues and foreign affairs and we omit the details of Egyptian and Israeli preferences. Let us consider that $D(A) = (a_1, a_2, a_3)$, $D(B) = (b_1, b_2)$, $D(C) = (c_1, c_2, c_3)$, and $D(D) = (d_1, d_2)$. There is a total combination of 36 possible agreements. The mediator, representing U.S., prepares CP-nets and induced preference graphs for both players. We can define the utility values of all variables for both players. Then, negotiation process continues according to Algorithms 1 and 2. We found that indeed, a final agreement is achieved within a finite number of rounds. For details, see <http://www.ai.rug.nl/SocialCognition/experiments/>.

4 DISCUSSION ON RELATED WORK

This paper presents an approach for negotiation over

multiple players on multiple issues with the support of a mediator. To achieve a jointly optimal agreement, the mediator offers a single negotiation text based on all players' preference graphs given from their CP-nets. Every player can decide to accept or reject the offer by checking the negotiation text's utility on his or her private UCP-nets. We proposed two algorithms for the mediator and the players. For successful negotiation between players, they often need to give up their maximum expected preferences because otherwise the negotiation process may not achieve a satisfactory agreement within a finite number of rounds. The proposed approach is appropriate for players who are willing to accept a jointly optimal choice.

M. Li and colleagues presented a protocol for negotiation in combinatorial domains (Li et al., 2011b), which can lead rational agents to reach optimal agreements under an incomplete information setting. They proposed POANCD (Protocol to reach Optimal Agreement in Negotiation over Combinatorial Domains), which has two phases. The first phase of POANCD consists of distributed formation of a negotiation tree by the participating agents, based on CP-nets of agents. After the first phase, the agents make a few initial agreements. In the second phase, the agents act cooperatively to achieve best possible agreement by exploring possible mutually beneficial alternatives. Li and colleagues also proved their approach to dominance testing in CP-nets (Li et al., 2011a). Their approach did not have a solution on cyclic CP-nets yet in (Li et al., 2011b). Recently, however, Li and colleagues proposed an approach called MajCP (Majority-rule-based collective decision-making with CP-nets) that can work with cyclic CP-nets as well (Li et al., 2011c).

The purpose of the proposed approach in this paper is to deal with similar situations, players do not only have preference orderings in their CP-nets, but they also have their private utility values (possibly different utilities for the same preference order). Our approach provides negotiation in an interactive way, with mediator and players as a game where the mediator proposes a single negotiation text and players can decide about agreement themselves.

K. Purrington and E. H. Durfee also proposed an algorithm to find social choices of two agents by exploiting the CP-net structure (Purrington and Durfee, 2009). Their algorithm searches agents' outcome graphs from the top down, using the satisfaction interval associated with each tier, and it proceeds for each agent. A set of candidate outcomes is maintained for each agent. It contains all the outcomes that an agent is willing to accept. For each agent, the algorithm examines the highest tier of outcomes that are not cur-

rently in its candidate set, and for the agent(s) with the highest minimum for this next tier, adds those outcomes to the candidate set(s). If the intersection of the agents' candidate sets is non-empty, one of the outcomes in the intersection has maximin optimality. Their algorithm considers only the level of the preference graphs. Our framework provides not only the mediator's joint choice, but also the players' decisions based on their private UCP-nets.

R. Aydogan and P. Yolum developed a negotiation approach using heuristics for CP-nets with partial preferences (Aydogan et al., 2011). They observed three negotiation strategies: depth-based, ranking-based and utility-based. They showed an example of negotiation between a producer and a consumer agent over a service. Negotiation takes place in a turn-taking fashion, where the consumer agent starts the negotiation with a particular service request. A service request can be considered as a vector of issues (discrete or continuous), which represents the service. If the producer agent does not prefer to supply this service, then the producer generates an alternative service. The consumer agent can accept this alternative service or continue negotiation to pursue a better one. This process continues until reaching a consensus or a deadline. Aydogan and Yolum focus only on one player's preferences and do not mention the other player's preferences. They do not deal with negotiation for multiple (more than two) players. The purpose of our approach is to deal with multi-player, multi-issue negotiation via a mediator.

M. Chalamish and S. Kraus presented AutoMed, an automated mediator for bilateral negotiation under time constraints, which uses a qualitative model. AutoMed produces the negotiators' preferences using Weighted CP-networks (WCP-nets). Each disputant specifies her preferences by creating her WCP-net using a graphical interface. Next, AutoMed sorts all possible agreements according to the WCP-nets and removes all non-optimal sets. During the negotiation process, AutoMed searches for an optimal offer by finding all agreements preferred to the offer made by the opponent in each list (Chalamish and Kraus, 2012). In our approach, the mediator does not use weighted or utility CP-nets but only CP-nets based on partial preferences of the players, because the players do not want to show their private utility values.

5 CONCLUSIONS

In this paper, we present a simple negotiation approach that is useful in a practical way for negotiations. The mediator offers jointly optimal negotiation

texts based on CP-nets over a finite number of rounds and the rest of the players are willing to adjust their private interests at each round using UCP-nets. In this interactive framework, the mediator and the players can play on different machines by sending messages. The framework can deal with negotiation for multiple issues, and with multiple players who have different preferences even when their preference graphs are cyclic. The proposed approach provides a satisfactory agreement for all players with their optimal outcomes, which are not less than average utility. As a future direction, we plan to test the performance of this approach and hope to construct a more efficient search algorithm on preference graphs. For the resulting negotiation algorithms, we intend to prove formal properties related to correctness and complexity.

ACKNOWLEDGEMENTS

We gratefully acknowledge NWO research grant 600.065.120.08N- 201, Vici grant NWO 227-80-001, and a Lotus II grant from the EU.

REFERENCES

- Apt, K. R., Rossi, F., and Venable, K. B. (2005). CP-nets and Nash equilibria. In *Proceedings of the 3rd International Conference on Computational Intelligence, Robotics and Autonomous Systems (CIRAS)*, pages 13–16.
- Aydogan, R., Baarslag, T., Hindriks, K., Jonker, C., and Yolum, P. (2011). Heuristic-based approaches for CP-nets in negotiation. In *Proceedings of The 4th International Workshop on Agent-based Complex Automated Negotiations (ACAN)*.
- Bacchus, F. and Grove, A. (1995). Graphical models for preference and utility. In *Proceedings of the Eleventh conference on Uncertainty in Artificial Intelligence, UAI'95*, pages 3–10.
- Boutilier, C., Bacchus, F., and Brafman, R. I. (2001). UCP-networks: A directed graphical representation of conditional utilities. In *The 17th Conference on Uncertainty in Artificial Intelligence (UAI-01)*, pages 56–64.
- Boutilier, C., Brafman, R. I., Hoos, H. H., and Poole, D. (2004a). CP-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *Journal of Artificial Intelligence Research*, 21:135–191.
- Boutilier, C., Brafman, R. I., Hoos, H. H., and Poole, D. (2004b). Preference-based constrained optimization with CP-nets. *Computational Intelligence*, 20:137–157.
- Chalamish, M. and Kraus, S. (2012). Automed: An automated mediator for multi-issue bilateral negotiations. *Autonomous Agents and Multi-Agent Systems*, 24(3):536–564.
- Ehtamo, H., Kettunen, E., and Hamalainen, R. P. (2001). Searching for joint gains in multi-party negotiations. *European Journal of Operational Research*, 130(1):54–69.
- Fatima, S. S., Wooldridge, M., and Jennings, N. R. (2005). Bargaining with incomplete information. *Annals of Mathematics and AI*, 44(3):207–232.
- Fisher, R. (1978). *International Mediation: A Working Guide*. International Peace Academy, New York.
- Korhonen, P., Oretskin, N., Teich, J., and Wallenius, J. (1995). The impact of a biased starting position in a single negotiation text type mediation. *Group Decision and Negotiation*, 4:357–374.
- Li, M., Vo, Q. B., and Kowalczyk, R. (2011a). Efficient heuristic approach to dominance testing in CP-nets. In *Proceedings of 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 353–360.
- Li, M., Vo, Q. B., and Kowalczyk, R. (2011b). An efficient protocol for negotiation over combinatorial domains with incomplete information. In *Proceedings of the 27th Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 436–444, Barcelona, Spain.
- Li, M., Vo, Q. B., and Kowalczyk, R. (2011c). Majority-rule-based preference aggregation on multi-attribute domains with CP-nets. In Tumer, Yolum, S. and Stone, editors, *Proceedings of 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 659–666.
- Liu, Z., Li, X., Guo, S., and Liu, J. (2012). A game tree algorithm for optimal outcome of CP-nets. *International Journal of Digital Content Technology and its Applications (JDCTA)*, 6.
- Oakman, J. (2002). The Camp David Accords: A case study on international negotiation. Technical report, Princeton University, Woodrow Wilson School of Public and International Affairs.
- Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann Publishers Inc, San Francisco, CA, USA.
- Purrington, K. and Durfee, E. H. (2009). Agreeing on social outcomes using individual CP-nets. *Multiagent and Grid Systems*, 5(4):409–425.
- Raiffa, H., Richardson, J., and Metcalfe, D. (2002). *Negotiation Analysis: The Science and Art of Collaborative Decision Making*. Belknap Press of Harvard Univ. Press, Cambridge (MA).
- Rasmusen, E. (2007). *Games and Information*. Blackwell Publishing, four edition.
- Rossi, F., Venable, K. B., and Walsh, T. (2004). mCP nets: Representing and reasoning with preferences of multiple agents. In *Proceedings of the 19th National Conference on Artificial Intelligence (AAAI)*, pages 729–734.
- Telhami, S. (1993). Evaluating bargaining performance: The case of Camp David. *Political Science Quarterly*, 107:629–653.

APPENDIX

Algorithm 1: Negotiation Decision by Player.

```

1: Input :
2:   UCP-nets with total utility and ordering
3:    $Agreement \leftarrow 0$ ;  $maxU \leftarrow$  maximum utility
4:    $threshold \leftarrow 0.3$ ; Proposals  $\leftarrow \emptyset$ 
5:   while  $Agreement \neq 1$  and  $t \leq finalRound$  do
6:     Search  $S$  ( $U(S) = maxU$ ) in UCP-net //search
       maximum preferred proposal
7:     Send ( $S$ ) to Mediator
8:     currentProposal  $\leftarrow$  Receive(Proposal by Media-
       tor)
9:     Proposals  $\leftarrow$  Proposals  $\cup$  currentProposal
10:    if ( $maxU - U(currentProposal) \leq threshold$ )
       then
11:      accept Proposal
12:       $Agreement \leftarrow 1$ 
13:    else
14:      reject Proposal
15:      Update  $maxU$  //search and update the utility
       less than current maxU
16:      Send ( $S$ )
17:    end if
18:  end while
19:  if Receive(finalRound) then
20:    while  $maxU > avgU$  do
21:      Update  $maxU$  //search and update the utility
       less than current maxU
22:      Evaluate proposals
23:      if  $\exists proposal : (maxU - U(proposal) \leq$ 
        $threshold)$  then
24:        accept proposal;
25:         $Agreement \leftarrow 1$ 
26:      end if
27:    end while
28:  end if

```

Algorithm 2: Negotiation by Mediator.

```

1: Input:
2:    $Player N: N = (1, 2, \dots, n)$ 
3:    $CPN_1, CPN_2, \dots, CPN_n$  //Players' CP-nets
4:    $S_1, S_2, \dots, S_n$  //Players' proposals
5:    $maxP \leftarrow 0.9$ ;  $avgP \leftarrow 0.5$  //maximum and av-
       erage acceptable probability
6:    $Agreement \leftarrow 0$ ;  $JointOptimal \leftarrow \emptyset$ ;  $t \leftarrow 0$ 
7:    $threshold \leftarrow 0.3$ 
8:   while  $Agreement \neq 1$  and  $t \leq finalRound$  do
9:     for  $i = 1$  to  $n$  do
10:      for  $j = 1$  to  $n$  do
11:        Search acceptableProbability ( $S_i, S_j : i \neq j$ )
12:      end for
13:    end for
14:    Mark all  $S_i, S_j$ : acceptableProbability  $<$ 
        $threshold$ ;
15:    while  $maxP > avgP$  do
16:      for  $i = 1$  to  $n$  do
17:        for  $j = 1$  to  $n$  do
18:          if  $\exists S_l : (acceptableProbability(S_i, S_j) =$ 
              $maxP ; i \neq j)$  then
19:             $JointOptimal \leftarrow S_l$ 
20:          else
21:            Search alternativeOptimal( $S_i$ ) //Other
             proposals with same maxP
22:            if  $\exists S_l : (acceptableProbability(S_l, S_j) =$ 
              $maxP ; l \neq j)$  then
23:               $JointOptimal \leftarrow S_l$ 
24:            else
25:               $maxP \leftarrow maxP - 0.1$ 
26:            end if
27:          end if
28:        end for
29:      end for
30:    end while
31:    Propose  $JointOptimal$ 
32:    if  $\forall Player k \in N$  accept Proposal then
33:       $Agreement \leftarrow 1$ 
34:    else if Player  $k (k \in N)$  rejects  $S_i$  then
35:      Ask new proposal to Player  $k$ 
36:      Update  $S_i (i = k)$ 
37:       $maxP \leftarrow 0.9$ 
38:    end if
39:  end while
40:  if  $Agreement = 0$  and  $t = finalRound$  then
41:    Announce finalRound and Ask for evaluating all
       proposals to Players
42:  end if

```
