

Complex Negotiations in Multi-Agent Systems



Víctor Sánchez Anguix

Departamento de Sistemas Informáticos y Computación

Universitat Politècnica de València

A thesis submitted for the degree of

Doctor of Philosophy in the subject of Computer Science

Under the supervision of:

Dr. Ana García Fornes and Dr. Vicente Julián Inglada

2013 February

PhD Thesis

Title: Complex Negotiations in Multi-Agent Systems

Author: Víctor Sánchez Anguix

Advisors: Dr. Ana García Fornes
Dr. Vicente Julián Inglada

Reviewers: Prof. Carles Sierra
Prof. Peter McBurney
Prof. Paulo Novais

Examination Board:

President Prof. Carles Sierra
Secretary Prof. Vicente Botti
Member Prof. Peter McBurney
Member Prof. Paulo Novais
Member Prof. Pavlos Moraitis

Day of the defense: 8th February 2013

Abstract

Multi-agent systems (MAS) are distributed systems where autonomous entities called agents, either human or software, pursue their own goals. The MAS paradigm has been proposed as the appropriate modeling approach for the deployment of applications like electronic commerce, multi-robot systems, security applications, and so forth. In the MAS community, the vision of open multi-agent system, where heterogeneous agents can enter and leave the system dynamically, has gained strength as a potentially interesting modeling paradigm due to its conceptual relation with technologies like world wide web, grid computing, and virtual organizations. Given the heterogeneity and agent's self-interest, conflict is a candidate phenomenon to arise in multi-agent systems.

In the last few years, the term agreement technologies has been used to address all the mechanisms that, directly or indirectly, promote the resolution of conflicts in computational systems like multi-agent systems. Among agreement technologies, automated negotiation is proposed as one key mechanism in conflict resolution due to its analogous use in human conflict resolution. Automated negotiation consists of an automated exchange of proposals carried out by software agents on behalf of their users. The final goal is the achievement of an agreement with all the involved parts.

Despite being studied by scholars in Artificial Intelligence for several years, several problems have not been addressed by the scientific community yet. The main objective of this thesis is proposing negotiation models for complex scenarios where the complexity may stem from (i) limited computational capabilities or (ii) the necessity to accommodate the preferences of multiple individuals. In the first part of the thesis we propose a bilateral negotiation model for the problem of negotiation in Ambient Intelligence (AmI), a domain with a special emphasis

on computational efficiency due to the limited capability of AmI devices. In the second part of the thesis we propose several negotiation models for agent-based negotiation teams. A negotiation team is a group of individuals that acts together as single negotiation party due to its common interests in the negotiation at hand. The complexity of negotiation teams resides in the fact that despite having common interests, intra-team conflict is also present. As far as we are concerned, the topic of negotiation teams in MAS is introduced with this thesis.

Resumen

Los sistemas multi-agente (SMA) son sistemas distribuidos donde entidades autónomas llamadas agentes, ya sean humanos o software, persiguen sus propios objetivos. El paradigma de SMA ha sido propuesto como la aproximación de modelo apropiada para aplicaciones como el comercio electrónico, los sistemas multi-robot, aplicaciones de seguridad, etc. En la comunidad de SMA, la visión de sistemas multi-agente abiertos, donde agentes heterogéneos pueden entrar y salir del sistema dinámicamente, ha cobrado fuerza como paradigma de modelado debido a su relación conceptual con tecnologías como la Web, la computación grid, y las organizaciones virtuales. Debido a la heterogeneidad de los agentes, y al hecho de dirigirse por sus propios objetivos, el conflicto es un fenómeno candidato a aparecer en los sistemas multi-agente.

En los últimos años, el término tecnologías del acuerdo ha sido usado para referirse a todos aquellos mecanismos que, directa o indirectamente, promueven la resolución de conflictos en sistemas computacionales como los sistemas multi-agente. Entre las tecnologías del acuerdo, la negociación automática ha sido propuesta como uno de los mecanismos clave en la resolución de conflictos debido a su uso análogo en la resolución de conflictos entre humanos. La negociación automática consiste en el intercambio automático de propuestas llevado a cabo por agentes software en nombre de sus usuarios. El objetivo final es conseguir un acuerdo con todas las partes involucradas.

Pese a haber sido estudiada por la Inteligencia Artificial durante años, distintos problemas todavía no han sido resueltos por la comunidad científica todavía. El principal objetivo de esta tesis es proponer modelos de negociación para escenarios complejos donde la complejidad deriva de (i) las limitaciones computacionales o (ii) la necesidad de representar las preferencias de múltiples individuos. En la

primera parte de esta tesis proponemos un modelo de negociación bilateral para el problema de las negociaciones en la Inteligencia Ambiental (AmI), un dominio con un énfasis especial en la eficiencia computacional debido a las características de los dispositivos que podemos encontrar en el escenario. En la segunda parte de esta tesis proponemos diversos modelos de negociación para equipos de negociación. Un equipo de negociación es un grupo de individuos que actúa como una única parte en el proceso de negociación debido a sus intereses comunes. La complejidad en los equipos de negociación reside en el hecho de que, pese a tener intereses comunes, el conflicto dentro del equipo también está presente. En lo que nos concierne, el tema de los equipos de negociación en SMA es introducido con esta tesis.

Resum

Els sistemes multi-agent (SMA) són sistemes distribuïts on entitats autònomes anomenades agents, ja siguin humans o programes, persegueixen els seus propis objectius. El paradigma de SMA ha sigut proposat com una aproximació apropiada per a aplicacions com el comerç electrònic, els sistemes multi-robot, aplicacions de seguretat, etc. En la comunitat de SMA, la visió de sistemes multi-agents oberts, on agents heterogenis poden entrar i eixir del sistema dinàmicament, ha pres força com a paradigma de modelatge degut a la seua relació conceptual amb tecnologies com la Web, la computació grid, i les organitzacions virtuals. Degut a la heterogeneïtat dels agents, i al fet d'estar dirigits pel seus propis objectius, el conflicte és un fenomen candidat a aparèixer en els sistemes multi-agent.

En els darrers anys, el terme tecnologies de l'acord ha sigut usat per a referir-se a tots aqueixos mecanismes que, directa o indirectament, promouen la resolució de conflictes en sistemes computacionals com són els sistemes multi-agent. Entre les tecnologies de l'acord, la negociació automàtica ha sigut proposta com a un dels mecanismes clau en la resolució de conflictes degut al seu ús anàleg en la resolució de conflictes entre humans. La negociació automàtica consisteix en l'intercanvi automàtic de propostes per part d'agents software en el nom dels seus usuaris. L'objectiu final es aconseguir un acord amb totes les parts involucrades.

Malgrat haver sigut estudiada per la Intel·ligència Artificial durant anys, diversos problemes encara no han sigut resolts per la comunitat científica. El principal objectiu d'aquesta tesi és proposar models de negociació per a escenaris complexos on la complexitat deriva de (i) les limitacions computacionals o (ii) la necessitat de representar les preferències de múltiples individus. En la primera part d'aquesta tesi proposem un model de negociació bilateral per al problema de la Intel·ligència Ambiental (AmI), un domini amb un èmfasi especial en la eficiència computacional

degut a les característiques dels dispositius que podem trobar en l'escenari. En la segona part d'aquesta tesi proposem diversos models de negociació per a equips de negociació. Un equip de negociació és un grup d'individus que actua com a una única part en el procés de negociació degut als seus interessos comuns. La complexitat en els equips de negociació resideix en el fet que, encara que tenen interessos comuns, el conflicte dins de l'equip també està present. En allò que ens concerneix, el tòpic dels equips de negociació en SMA és introduït en aquesta tesi.

To my family.

Without your dedication and support, this thesis would not be possible.

Acknowledgements

It is incredibly hard to thank all of the people that, in some way or another, have helped me during these graduate years. First of all, I would like to thank Vicente, Vicent, and Ana for their patience. I know that I can be stubborn, and that makes me “difficult” sometimes.

I would like to acknowledge Prof. Katia Sycara for her supervision during my visit at Carnegie Mellon University. I should also thank Prof. Catholijn Jonker and Dr. Reyhan Aydogan for their invaluable help with the last chapter of this thesis.

It would be unconsiderate if I forgot about my labmates. The moments of fun and joy have been uncountable. Of course, I should also thank LAB 208 since I have been their “special guest” quite often. I should specially thank those “miserable” inhabitants of LAB 208, because they have contributed to make my graduate years the best academic cycle of my life.

I would like to thank my parents, my sister, and Barbara, since they have been there for me whenever I needed them.

Finally, since the list of people grows exponentially with each paragraph, I would like to acknowledge all of them:

$$\forall x, \quad \textit{Supported}(x, \textit{Victor}) \longrightarrow \textit{Thanks}(\textit{Victor}, x).$$

Contents

List of Figures	ix
List of Tables	xi
1 Motivation	1
1.1 Objectives	6
1.2 Contributions	8
1.3 Document Structure	9
2 State of the Art	11
2.1 Introduction to Agreement Technologies	11
2.2 Automated Negotiation	15
2.2.1 Automated Negotiation and Ambient Intelligence	17
2.2.2 Concession Strategies	18
2.2.3 Similarity Mechanisms in Negotiation	20
2.2.4 Bayesian Learning in Negotiation	22
2.2.5 Genetic Algorithms in Negotiation	25
2.2.6 Offline Learning in Negotiation	29
2.2.7 Complex Interdependent Utility Functions	31
2.3 Negotiation Teams	35

CONTENTS

2.4	Conclusions & Discussion	38
2.4.1	Ambient Intelligence	38
2.4.2	Agent-Based Negotiation Teams	43
3	Bilateral Negotiation for Limited Devices	47
3.1	Introduction	47
3.2	Motivating Scenario: Product Fairs	48
3.3	Negotiation Model	51
3.3.1	Negotiation Protocol	52
3.3.2	Negotiation Strategy	54
3.3.2.1	Pre-negotiation: Self-sampling	56
3.3.2.2	Negotiation: Concession strategy	59
3.3.2.3	Negotiation: Acceptance criteria	61
3.3.2.4	Negotiation: Evolutionary sampling	62
3.3.2.5	Negotiation: Select which offers to send	65
3.3.2.6	Negotiation Trace	66
3.4	Experiments	73
3.4.1	Negotiation Setting	74
3.4.2	Results	76
3.4.2.1	Experiment 1: Performance Study on the Number of Issues	77
3.4.2.2	Experiment 2: Performance Study on p_{pevo}	79
3.4.2.3	Experiment 3: Performance Study on k	80
3.4.2.4	Experiment 4: Performance Study on $ P $ and Mem- ory Performance	81
3.4.2.5	Experiment 5: Time Performance	83
3.5	Conclusions	84

4	A General Workflow of Tasks for Negotiation Teams	91
4.1	Introduction	91
4.2	Identify Negotiation	93
4.3	Team Formation	96
4.4	Opponent Selection	100
4.5	Understand the Negotiation Domain	100
4.6	Agree Negotiation Issues	103
4.7	Plan Negotiation Protocol	104
4.8	Agree Negotiation Protocol	105
4.9	Decide Intra-team Strategy	107
4.10	Select Individual Strategy	109
4.11	Negotiation & Adaptation	110
4.12	Conclusions	112
5	Intra-Team Strategies for Negotiation Teams in Predictable Do-	119
	main	119
5.1	Introduction	119
5.2	Negotiation Setting	122
5.3	Motivating Example: Group Booking	126
5.4	Representative (RE)	129
5.4.1	Offer proposal	130
5.4.2	Evaluation of Opponent's Offer	131
5.4.3	Discussion	131
5.5	Similarity Simple Voting (SSV)	132
5.5.1	Offer proposal	133
5.5.2	Evaluation of Opponent's Offer	135
5.5.3	Discussion	136
5.6	Similarity Borda Voting (SBV)	139

CONTENTS

5.6.1	Offer proposal	140
5.6.2	Offer acceptance	140
5.6.3	Discussion	141
5.7	Full Unanimity Mediated (FUM)	142
5.7.1	Pre-negotiation: information sharing	143
5.7.2	Negotiation: Offer proposal	146
5.7.3	Negotiation: observing opponent's concessions and building an issue agenda	149
5.7.4	Negotiation: Offer acceptance	150
5.7.5	Discussion	151
5.8	Empirical Analysis of Full Unanimity Mediated Intra-Team Strategy	156
5.8.1	Studying the Impact of Intra-Team Agenda	156
5.8.2	Studying the Impact of ϵ_{a_i}	160
5.8.3	Strategy Deviation	162
5.9	Studying the Impact of the Negotiation Environment on Intra-Team Strategies' Performance	166
5.9.1	Negotiation Environment Conditions & Team Performance	166
5.9.1.1	Environment Condition: Opponent Deadline Length	167
5.9.1.2	Environment Condition: Team Deadline Length .	167
5.9.1.3	Environment Condition: Team Similarity	167
5.9.1.4	Environment Condition: Opponent Concession Speed	169
5.9.1.5	Environment Condition: Number of Team Members	169
5.9.1.6	Team Performance: Number of Negotiation Rounds	170
5.9.1.7	Team Performance: Minimum Utility of Team Mem- bers	170
5.9.1.8	Team Performance: Average Utility of Team Mem- bers	171

5.9.2	Results	171
5.9.2.1	Number of Negotiation Rounds	171
5.9.2.2	Same Type of Deadlines	173
5.9.2.3	Different Types of Deadlines	177
5.9.2.4	Team size effect on intra-team strategies	183
5.10	Conclusions	185
6	Negotiation Teams in Unpredictable Domains	195
6.1	Introduction	195
6.2	Negotiation Setting	196
6.3	Motivating Example: Advanced Group Booking	198
6.4	Intra-Team Protocol	200
6.4.1	Pre-negotiation Phase	200
6.4.2	Negotiation Phase	204
6.4.2.1	Evaluation of Opponent's Offer	205
6.4.2.2	Offer Proposal	205
6.5	Team Members' Strategies	207
6.5.1	Basic Strategy for Team Members	209
6.5.1.1	Evaluation of Opponent's Offer	209
6.5.1.2	Offer Proposal	210
6.5.2	Bayesian-based Strategy for Team Members	211
6.5.2.1	Bayesian Learning	211
6.5.2.2	Offer Proposal	212
6.6	Unanimously Acceptable Proof	214
6.7	Experiments	216
6.7.1	First Experiment: Measuring Model Performance	217

CONTENTS

6.7.2	Second Experiment: Analyzing the impact of parameters for the proposal of unpredictable partial offers	223
6.7.3	Third Experiment: Analyzing the Impact of the Reservation Utility	227
6.8	Conclusions	232
7	Conclusions	233
7.1	Contributions	233
7.2	Future Work	236
7.3	Scientific Publications	240
7.3.1	Publications in SCI Journals	240
7.3.2	Publications in Conferences	241
7.4	Scientific Research Stays	242
	References	245

List of Figures

2.1	An example of linear (left) and complex (right) utility functions . . .	32
3.1	An example of the furniture fair application	51
3.2	An example of two agents negotiating in the k-alternating protocol.	54
3.3	An example of a crossover operation.	64
3.4	Distance to the Nash Product, distance to the closest Pareto Point, and number of negotiation in Experiment 1.	87
3.5	Distance to the Nash Product, distance to the closer Pareto Point, and number of negotiation rounds in Experiment 2.	88
3.6	Distance to the Nash Product, distance to the closest Pareto Point, and number of negotiation rounds in Experiment 3.	89
3.7	Distance to the Nash Product, distance to the closest Pareto Point, and number of negotiation rounds in Experiment 4.	90
4.1	A general workflow of tasks for agent-based negotiation teams. . .	94
5.1	Average utility reported to the opponent by team's proposals per round in short (upper graphic) and long (lower graphic) deadline scenarios.	158

LIST OF FIGURES

5.2	Results for very similar team members when the team has a long deadline, the opponent has a short deadline and the opponent uses conceder or very conceder tactics.	189
5.3	Results for very similar team members when the team has a long deadline, the opponent has a short deadline and the opponent uses boulware or very boulware tactics.	190
5.4	Results for very dissimilar team members when the team has a long deadline, the opponent has a short deadline and the opponent uses conceder or very conceder tactics.	191
5.5	Results for very dissimilar team members when the team has a long deadline, the opponent has a short deadline and the opponent uses boulware or very boulware tactics.	192
5.6	Effect of the size of the team on team performance when both parties have the same type of deadline and team members are very similar.	193
5.7	Effect of the size of the team on team performance when both parties have the same type of deadline and team members are very dissimilar.	194
6.1	Representation of the negotiation space of an agent	202
6.2	Overview of the intra-team protocol carried out during the pre-negotiation	204

List of Tables

2.1	Categorization of computational negotiation models.	41
3.1	Approximate amount of memory needed by the <i>NES</i> and <i>ES</i> model.	83
3.2	Average time performance of the <i>ES</i> model for 30 negotiations. . .	84
5.1	Average impact of ϵ_{a_i} on team performance. Min: Minimum utility of team members, Ave: Average utility of team members	161
5.2	Confidence intervals (upper table) for the average utility of team members and number of failed negotiations (lower table) depending on the number of <i>deviated</i> agents and d_i	165
5.3	Average number of rounds when both parties have a long deadline.	173
5.4	Average of the minimum utility of team members (Min.) and the average of the average utility of team members (Ave.).	174
5.5	Average for the minimum utility of team members (Min.) and the average utility of team members (Ave.) when the team has a short deadline and the opponent has a long deadline.	181
6.1	Average joint Utility (T. Joint) and average opponent Utility (Op) for the first set of experiments.	220

LIST OF TABLES

6.2	Impact of w_A and w_{op} on the team joint utility and the opponent utility in different scenarios	226
6.3	Average percentage of unpredictable partial offers pruned in the pre-negotiation.	229
6.4	Average joint Utility (T. Joint) and average opponent Utility (Op) for teams composed by Bayesian team members with different reservation utilities.	230

1

Motivation

Multi-agent systems (MAS) are distributed systems where autonomous entities called agents, either human or software, pursue their own goals in reactive, proactive and social ways (1). This paradigm has been proposed as an adequate modeling approach for the deployment of applications like electronic commerce (2), multi-robot systems (3), security applications (4), and so forth. Inside the MAS community, the vision of open multi-agent system, where heterogeneous agents can enter and leave the system dynamically, has gained strength as a potentially interesting modeling paradigm due to its conceptual relation with technologies like world wide web, grid computing, and virtual organizations (5, 6). Given the heterogeneity and agent's self-interest, conflict is a candidate phenomenon to arise in multi-agent systems. In the last few years, the term agreement technologies (7, 8) has been used to address all the mechanisms that, directly or indirectly, promote the resolution of conflicts in computational systems like multi-agent systems.

Among agreement technologies, negotiation is proposed as one key mechanism in conflict resolution due to its analogous use in the resolution of human conflicts. Pruitt (9) defines negotiation as a process in which a joint decision is made by two

1. MOTIVATION

or more parties by verbalizing contradictory demands and then moving towards and agreement. Classically, negotiation has been studied in the social sciences and game theory. On the one hand, the social sciences mainly study how humans behave and act in real negotiation processes (10). On the other hand, game theory researchers focus on looking for optimal agreements under the assumptions of unbounded computational resources and complete/partial information regarding opponent's preferences and strategies. Some of the most important theoretical results come from game theory (e.g., the work of Nash (11), Rubinstein's work (12), and Binmore's work (13)). Although game theory provides interesting theoretical results, most of game theory's assumptions do not hold in computer systems since there are limitations on the information regarding players, and computational resources are limited.

Artificial Intelligence (AI) scholars have focused on solving the problem of negotiation in computer systems. Thus, works in AI usually assume that information regarding the opponent is usually imperfect or non-existent, and computational resources are bounded and limited. The goal in AI has been reaching near optimal solutions at reasonable computational costs.

The first works in negotiation from the perspective of AI are related to the area of Negotiation Support Systems (NSS) (14, 15, 16, 17): decision support systems that assist humans in real negotiations by providing communication infrastructures, predictions, strategy suggestions, and analysis tools for the available information. Nowadays, given the increasing implantation of large scale open multi-agent systems, the number of available partners with whom one may negotiate/interact has increased exponentially. Since human negotiation across the Internet could be extremely time consuming, automated negotiation has arisen as a solution for large scale systems. As its name indicates, automated negotiation

consists in autonomous software agents reaching agreements on behalf of their users.

Despite the fact that automated negotiation has been studied since the 90's decade (18, 19, 20, 21), there is still a wide range of problems whose solution has not been treated in the literature.

On the one hand, nowadays the number of computational devices present in our everyday life has grown considerably. The use of technology helps us to achieve a better quality of life, to make our life easier and more comfortable. However, due to the increasing number of devices, it is necessary that the technology itself adapts to the needs of the user, instead of the human being the one that adapts to technology. In that sense, Ambient Intelligence (AmI) tries to cover that necessity: it looks to offer personalized services and provide users with easier and more efficient ways to communicate and interact with other people and systems (22, 23). In Ambient Intelligence domains, users enter and leave the system in a very dynamic way. Applications are usually embedded in devices with very limited capabilities like smartphones, mobile phones, PDAs, and so forth. Given the heterogeneity of Ambient Intelligence domains, conflict may be present among users' goals. In that case, coordination and negotiation mechanisms are needed in order to solve conflicting situations. Putting a special emphasis on computational efficiency of negotiations carried out in Ambient Intelligence is of extreme importance. With computational efficiency, we refer to factors such as the number of offers sampled and the number of messages exchanged. We argue that most negotiation models, even though they care about efficiency, they have not focused on the particularities of Ambient Intelligence domains. The design of new computational models for negotiation in Ambient Intelligence domains may lead to the implantation of ubiquitous electronic commerce applications.

1. MOTIVATION

On the other hand, most negotiation models have focused on scenarios that are relatively simple compared to the scenarios that may be found in human negotiations and complex electronic applications. For instance, on the one hand, a vast majority of the negotiation models proposed in the literature are circumscribed to bilateral processes with two single individuals. Nevertheless, negotiation processes which only involve two single individuals hardly represent most negotiations carried out in the real world, which entangle much more complex processes. One of these negotiations that involve more than a single individual is negotiations where negotiation teams participate. Negotiation teams (24, 25, 26) are groups of two or more persons that join together as a single negotiation party because they share a common interest which is related to the negotiation process. Negotiations where parties are teams represent a great number of negotiations carried out in the real world. For instance, when a company wants to sell a product line to another company. It is habitual for both companies to send two negotiation teams, one per company, composed by persons from different organizational departments. Another scenario, involves a group of travelers that has decided to go on a travel together and decides to negotiate a deal with a travel agent. The group of traveler forms a negotiation team. As far as we are concerned, this thesis represents the first step in automated negotiation towards providing computational models for agent-based negotiation teams. We believe that the inclusion of agent-based negotiation teams in multi-agent systems may make possible the design of new social applications like electronic marketplaces for groups. In these applications, we believe that achieving unanimity among team members is a very important issue to be taken into account. By unanimity among team members we refer to the fact that the final agreement should be acceptable to all of the team members. This property avoids unexpected outcomes and creating discomfort among users in the long run. Therefore, we consider it to be of extreme importance.

Hence, in this thesis we pursue computational solutions for both complex scenarios: negotiation in Ambient Intelligence domains and agent-based negotiation teams. This thesis has been developed under the umbrella of several research projects in multi-agent systems. Automated negotiation, more specifically, the research carried out in this thesis plays an important role in those research projects. This thesis is developed under the framework of the following research projects funded by the Spanish Government:

- “Agreement Technologies” Consolider-INGENIO 2010 under grant CSD2007-00022 (Main Researcher: Carles Sierra, from 2007 to 2012). Agreement technologies is a term coined in the last few years to refer to those technologies that allow computational entities to automatically solve conflicts. Being used by humans so frequently, negotiation is one of the key technologies in agreement technologies. The work carried out in this thesis aims to advance the state-of-the-art in mechanisms that are able to solve conflict among computational entities.
- “Magentix2: A Multi-agent Platform for Open Multi-agent Systems” under grant TIN2008-04446 (Main Researcher: Ana Garcia-Fornes, from 2008 to 2011). Magentix2 is a multi-agent platform that aims to provide support for open systems. Open multi-agent systems are computational systems where heterogeneous agents can enter and leave the system dynamically. In such systems, conflict may arise due to the divergence of goals and interests shown by heterogeneous agents. The work of this thesis aims to provide negotiation mechanisms for Magentix2 agents.
- “Multi-agent Plan Interaction” under grant TIN2011-27652-C03-01-AR (Main Researcher: Eva Onaindia, from 2012). In this project, we aim to analyze processes where groups of agents aim to cooperate while having divergent

1. MOTIVATION

interests. Negotiation teams are groups of agents who share a common goal at the negotiation process. However, each team member may have different personal goals that are possibly in conflict. Agent-based negotiation teams are a subset of the scenarios studied by the Multi-agent Plan Interaction project.

Additionally, this work has been also supported by “Advances on Agreement Technologies for Computational Entities” PROMETEO/2008/051 (Main Researcher: Vicente Botti) funded by the Valencian Government. Moreover, the work of this thesis could have not been possible without a 4-year FPU research grant AP2008-00600 granted by the Spanish Government.

1.1 Objectives

As stated before, the main aim of this thesis is providing computational models for negotiation in complex scenarios. More specifically, we focus on negotiation in Ambient Intelligence domains and computational models for agent-based negotiation teams. For that purpose, we decided to propose the following sub-goals:

1. **State-of-the-art in Automated Negotiation:** It is necessary to survey, classify, and review the existing literature on automated negotiation and related topics.
 - (a) Discuss the adequation of current negotiation models for Ambient Intelligence.
 - (b) Discuss the state-of-the-art of negotiation teams in related topics like the social sciences and point out the relations with agent-based negotiation teams.

2. **Negotiation in Ambient Intelligence Environments:** Ambient Intelligence is a domain that requires special features due to the limited capabilities of the devices that are usually employed. Even though negotiation models in Artificial Intelligence care about computational efficiency, they have not focused on specially limited domains as Ambient Intelligence. Therefore, it is necessary to propose and validate a computational model for negotiation in Ambient Intelligence.

- (a) Propose a general bilateral negotiation model for Ambient Intelligence that can be adapted to several domains.
- (b) Validate the computational efficiency of the proposed mechanism: offers sampled, and number of negotiation rounds.
- (c) Validate the economic efficiency of the proposed mechanism.

3. **Agent-based Negotiation Teams:** As far as we are concerned, the topic of agent-based negotiation teams is introduced in automated negotiation with this thesis. The negotiation process is complex since not only the team should solve the conflict with the opponent, but also the conflict that may arise inside the team. Therefore, due to its novelty, we put a special emphasis on exploring this type of complex negotiation. There may be multiple negotiation team models for the same scenario and their performance may vary depending on several environmental factors. We aim to propose several computational models for negotiation teams and analyze the impact of environment conditions on team performance.

- (a) Identify and analyze the workflow of tasks necessary that may help agent-based negotiation teams to perform successfully in negotiations.

1. MOTIVATION

- (b) Propose and validate computational models for negotiation teams under negotiation domains with *predictable and compatible* issues.
- (c) Propose and validate computational models for negotiation teams under negotiation domains with *predictable and compatible* and *unpredictable* issues.
- (d) Propose and validate computational models for negotiation teams that are capable of guaranteeing unanimously acceptable agreements among team members.
- (e) Analyze the effect of environmental conditions on negotiation team's models and identify those models that work better under specific environmental conditions.

1.2 Contributions

The specific contributions of this thesis are:

- **State of the art.** To achieve sub-goal *1.a* we review the most important works in automated negotiation, and more specifically in bilateral negotiation. We analyze the adequateness of the different models proposed in the literature for Ambient Intelligence, and we identify those mechanisms that may prove more interesting for the aforementioned domain. For sub-goal *1.b*, we review the literature in the social sciences and relate findings with its computational counterpart.
- **A genetic-aided bilateral negotiation model for negotiation in Ambient Intelligence.** We propose a computational negotiation model for bilateral negotiations carried out in Ambient Intelligence domains. The negotiation model aims to work with complex interdependent utility functions

using computational resources as efficiently as possible. It aims to provide solutions for sub-goals *2.a*, *2.b*, and *2.c*.

- **A general workflow of tasks for agent-based negotiation teams.** We propose a complete workflow of tasks that may help agent-based negotiation teams to achieve success in negotiation processes. Each task is analyzed and related to other areas in MAS. Additionally, open challenges that may arise in the specific case of negotiation teams are highlighted. This contribution covers *3.a* and also *1.b* to some extent.
- **Agent-based negotiation team models: Representative, Similarity Simple Voting, Similarity Borda Voting, and Full Unanimity Mediated (and extension).** The models aim to provide solutions for *3.b* and *3.c* while providing different levels of unanimity regarding team decisions (i.e., no unanimity, majority/plurality, semi-unanimity, unanimity). Full Unanimity Mediated and its extension, which guarantee unanimity regarding team decisions, thus, cover sub-goal *3.d*. The experimental evaluation of these models under different environmental conditions covers sub-goal *3.e*.

1.3 Document Structure

The remainder of this thesis is structured as follows. First, a state-of-the-art in automated negotiation and negotiation teams is presented in Chapter 2. With respect to the state-of-the-art in automated negotiation, a special emphasis is put in bilateral models and the adequateness of such models for Ambient Intelligence domains. The state-of-the-art regarding negotiation teams is presented from the perspective of the social sciences, and their findings are related with the implantation of negotiation teams in computer systems. Then, Chapter 3 presents our

1. MOTIVATION

proposal and evaluation of a negotiation model for Ambient Intelligence domains. The next chapter, Chapter 4, describes our proposal for a general workflow of tasks for agent-based negotiation teams. Computational models for negotiation teams that work in domains exclusively composed by *predictable and compatible* issues among team members are presented and evaluated under different environmental conditions in Chapter 5. The proposal of this thesis concludes in Chapter 6 with the extension of one of our computational models to guarantee unanimously acceptable team decisions in domains composed by *predictable and compatible* and *unpredictable issues*. Finally, we present our concluding remarks and possible future lines of work in Chapter 7.

2

State of the Art

2.1 Introduction to Agreement Technologies

Open multi-agent systems are distributed systems where heterogeneous agents, with their own goals, can enter and leave the system during the life of the system (27). For instance, we can think of an electronic commerce platform as an open system where users, human or even automated software, acts according to its own interests: in the case of sellers to maximize their own profits, and in the case of buyers to acquire some goods at relatively good price.

Since agents (humans or software) have different goals, act based on their goals, and they are heterogeneous (i.e., humans and software agents may show different behaviors), it is feasible to find situations where an agent's goals conflict with other agents' goals. If we refer ourselves to the example of electronic commerce, the buyer may want to buy the product at a low price, while the seller may want to maximize its revenue. In these situations, mechanisms that allow groups of agents to coordinate, regulate their behavior, and solve conflict are needed.

Electronic commerce is not the only application where conflict may make act

2. STATE OF THE ART

of presence. For instance, in the last few years grid computing (28) has emerged as a new paradigm of computation where different entities collaborate to accomplish several tasks. In grid computing, entities share several resources: from hardware resources (e.g., computing nodes) to software resources (e.g., services). How should those resources be distributed among the different tasks or users of the system? Presumably, users want the best response time for their tasks, and resource owners want to take the highest profit of their resources. Resource allocation is a delicate matter, especially when collaboration requires cross-boundary relationships. Software mechanisms that solve conflict in these scenarios are needed.

Even purely cooperative applications like rescue applications (29) are not alien to conflict. In multi-robot systems for rescue applications, information is usually distributed among the different robotic agents. Coordination among these entities is a problem itself, which becomes more acute when agents' opinions and information conflict. How should these entities solve conflict and rescue as many persons as possible while making an efficient use of the computational resources? Again, software mechanisms are necessary to tackle conflict.

The term *Agreement Technologies* (7, 8) has been coined in the last few years as an umbrella term for addressing all of those technologies that are envisioned to collaborate, directly or indirectly, to the resolution of conflicts in software systems. Even though which works can be considered agreement technologies is arguable (since the contribution to the resolution of a conflict may be indirect), some authors distinguish between several challenges that need to be solved in the so-called agreement technologies. In this thesis, we position ourselves with the taxonomy/challenges introduced by Sierra et al. (8). Despite the fact that automated negotiation is the focus of this thesis, we think that it is important to briefly describe the role of every other technology involved in agreement technologies, since it should help the reader to gain a broader view of how conflict may

2.1 Introduction to Agreement Technologies

be solved in software systems. Next, we briefly describe each of the challenges mentioned by Sierra et al.:

- **Semantics:** The current trend of service-oriented computing (30) has changed the way in which complex systems are built. Nowadays, software is built by using diverse services offered by very different providers. Given the heterogeneity of service providers, it is logic to think that service information is provided in different communication languages, and even using different terms to address the same concept (i.e., different ontologies). Whenever a software system needs to cooperate or solve a conflictive situation with other systems, it requires of mechanisms that allow to understand other software systems by matching and aligning ontologies and semantic concepts (31, 32, 33).
- **Norms:** Most distributed applications are no longer static but open, and agents can exhibit a varied spectrum of behaviors. One possible way of “solving” conflict is avoiding conflict, establishing mechanisms that preclude agents of reaching a conflictive situation. Normative systems (34, 35) are envisioned with such purpose (among others). The society of agents is regulated by norms, which define which actions/states are to be punished in the system (e.g., to avoid conflict) and which actions/states are to be rewarded (e.g., promote actions that avoid conflict).
- **Organizations:** Agents usually have limited computational capabilities. Therefore, if a complex problem needs to be solved, agents need to join together as a group and coordinate to reach such complex goal. Agent organizations (5, 36, 37) may be seen as large and implicit coordination mechanisms that establish the roles to be played by agents and the interaction protocols to

2. STATE OF THE ART

be carried out among organizational members. In this sense, agent organizations may conflict by strictly defining the structure and interactions of the group.

- **Trust:** Trust mechanisms (38, 39), usually used in concordance with reputation models, are devised to help agents to select whom they should interact with. Trust is formed from one's own past experiences with other agents. To put it simply, positive experiences should bias one agent to collaborate and interact with the other party, whereas negative experiences should bias one agent to avoid interactions with the other party. Reputation is built according to the opinion that agent societies have on individuals. Trust and reputation mechanisms may help to reduce conflict by interacting with good partners.
- **Negotiation:** Finally, the technology at core of agreement technologies is the one that makes possible for agents to solve conflict per se. In this case, imitating how humans solve conflicts, agents negotiate looking for an agreement that is acceptable for all of the involved parties. Without this technology, it would not be possible for agents to solve conflict, at least not in an efficient way. Due to the fact that negotiation is the main topic of this thesis, it will receive a thorough review in the next section.

Despite negotiation being crucial for solving conflicts, it should be highlighted that it requires of semantics, norms, organizations, and trust to help in the resolution of conflicts. Semantics may help heterogeneous agents to form a negotiation domain (e.g., negotiation problem) that is understandable by all of the parties involved in the conflict situation. Then, society's norms may be used to formally force agents to respect established agreements. Otherwise, agents would violate agreements whenever it suits them. Organizations establish a framework where

roles and possible interaction protocols are formalized, giving room negotiations with clear rules of interaction (e.g., negotiation protocols), and helping agents to identify and search conflicting agents based on the information provided by roles (e.g., sellers and buyers are classical roles in conflict). Trust and reputation may guide agents to select negotiation opponents that are more likely to guarantee a good service. Hence, every technology in agreement technologies collaborates along negotiation in leading conflict situations towards good terms.

2.2 Automated Negotiation

Despite being part of a new topic like agreement technologies, automated negotiation has been studied by scholars for a few years. Negotiation can be defined as a process in which a joint decision is made by two or more parties. The parties first verbalize contradictory demands and then move towards agreement by a process of concession-making or search for new alternatives (9). Analogously, automated negotiation consists of an automated search process for an agreement between two or more software parties.

Two different research trends can be distinguished in automated negotiation models: game-theoretic models and heuristic models. Since the decade of the 50's, automated negotiation has been studied in game theory. Game theory researchers focus on reaching optimal solutions under assumptions of unbounded computational resources, complete/partial information regarding the strategies and preferences of other parties. Some of the most important theoretical results come from game theory, like the work of Nash (11), Rubinstein (12), Binmore (13), and more recent studies like Fatima et al. (40) and Serrano et al. (41). Although game theory studies are interesting from a theoretical point of view, most of them make strong assumptions that may not hold in real applications. For instance, compu-

2. STATE OF THE ART

tational resources are of extreme importance for agents since they may be scarce and shared among different tasks. Thus, negotiation should not always assume unbounded computational resources. Additionally, since agents are heterogeneous, not all of the agents know the same strategies. Identifying which strategies are known by each agent may be a hard task that can only be successful after several negotiations. The same goes for the knowledge regarding the opponents' preferences, reservations values, and so forth. Hence, models that tackle uncertainty and limit the use of computational resources are mandatory for some situations.

Heuristic models tackle the problem mentioned above. They do not calculate the optimum agreement, but they obtain results that aim to be as close as possible to the optimum. Heuristic models assume imperfect knowledge regarding the opponent and the environment, and aim to be computationally tractable while obtaining good results. The work carried out in this present thesis can be classified into this category of models. The reader is assumed to have some working knowledge on heuristic models for automated negotiation. In other case, the reading of several introductory texts and reviews like (2, 42, 43, 44, 45, 46) is recommended. The amount of literature in automated negotiation is vast and immense, ranging from bilateral negotiations, to multi-party negotiations. An extensive review of all of the problems in automated negotiation would be an almost non-feasible task.

As a part of this thesis, we decided to work in negotiation models for Ambient Intelligence domains. Thus, in this state-of-the-art we have mainly focused on identifying the adequateness of the most important negotiation models to Ambient Intelligence. Following, we discuss some of the most important works in the area of automated negotiation and bilateral negotiation.

2.2.1 Automated Negotiation and Ambient Intelligence

Nowadays, the number of computational devices present in our everyday life has grown considerably. The use of technology helps us to achieve a better quality of life, to make our life easier and more comfortable. However, due to the increasing number of devices, it is necessary that the technology itself adapts to the needs of the user, instead of the human being the one that adapts to technology. In that sense, Ambient Intelligence (AmI) tries to cover that necessity: it looks to offer personalized services and provide users with easier and more efficient ways to communicate and interact with other people and systems (22, 23).

Agent technology has been appointed as a proper technology for the support of AmI solutions (22, 47, 48). In fact, agents show interesting characteristics for AmI environments since they are reactive, proactive and social (1). Firstly, reactivity allows agents to change their behavior according to some new conditions in the AmI environment (new users, new services, etc.). Secondly, pro-activeness makes it possible for agents to act autonomously according to the user's goals, which results in a smooth and non-intrusive interaction with the AmI user. And lastly, the agent's social behavior allows several heterogeneous entities to cooperate and offer new complex services to the AmI user.

Conflict situations are not alien to AmI applications. For instance, shopping malls may be converted into ubiquitous environments where several vendors offer their products to passing shoppers (49, 50). In many cases, the shoppers know what they want but do not have time to check every shop that offers such products. A possible way of enhancing the customer experience is to automatically negotiate with all of the vendors. A list with the best agreements may be presented to the user through his mobile device. This way, the user does not have to check every possible shop since his mobile device has negotiated with every shop taking into

2. STATE OF THE ART

account the user preferences. Nevertheless, there are also benefits for vendors since automated negotiation allows a more flexible commerce than classic e-commerce. For instance, they may negotiate issues such as price, payment method, discounts, and dispatch dates, which is what often happens in traditional non-electronic commerce. Flexibility in e-commerce may result in client loyalty since the vendor is able to adapt as much as possible to the client preferences. Therefore, automated negotiation is a proper technology for e-commerce-based AmI applications such as shopping malls.

If executed in environments with limited capabilities like Ambient Intelligence domains (i.e., limited CPU, limited bandwidth, energy saving necessities, etc.), negotiation models need new requirements that may have been overlooked in the literature. Those requirements are limiting the number of interactions with opponents (i.e., number of messages sent), using the lesser CPU the better (i.e., reducing the number of offers sampled, efficient learning mechanisms, etc.), and reducing the use of memory (also related with the number of offers sampled). Additionally, economic requirements (e.g., utility of the final agreement) should not be forgotten. As far as we are concerned, classic automated negotiation models have not explicitly concentrated on fulfilling all of these requirements at once.

2.2.2 Concession Strategies

The classic view of artificial intelligence with respect to negotiation in incomplete information settings is that agents need to eventually concede in order to reach an agreement (18, 51, 52). However, agents can concede in very different magnitudes and in different rounds of the negotiation. Concession strategies determine how the agents concede and when these concessions are carried out.

The most influential work regarding concession strategies is, perhaps, the work of Faratin et al. (18). The authors proposed concession strategies that are a mix

of different families of concession tactics. The authors divide concession tactics into three different families:

- **Time-dependent tactics:** These tactics take into account the remaining time in the negotiation to carry out concessions. In this family, we can distinguish between linear tactics, bouldware tactics, and conceder tactics. On the one hand, linear tactics carry out the same amount of concession at each negotiation round until the reservation value is reached. On the other hand, conceder tactics concede very rapidly towards the reservation value in the first interactions, whereas bouldware tactics concede very slowly during the first negotiation rounds, but they concede faster as the negotiation process approaches the deadline.
- **Behavior-dependent tactics:** In the case of behavior-dependent tactics, the concession carried out by the agent depends on the negotiation movements performed by the opponent in the previous rounds. The classic tactic in this family is tit-for-tat, which mimics the concession carried out by the opponent in the previous round. Other variants of tit-for-tat include random absolute tit-for-tat, which performs the absolute concession carried out by the opponent in the last offer plus/minus a small deviation, and averaged tit-for-tat which takes the window of γ past opponent offers and carries out the average concession carried out by the opponent.
- **Resource-dependent tactics:** This family of tactics computes concession based on the scarceness of a resource in the environment and resource consumption (i.e., time, product quantity, messages, etc). In general, the scarcer the resource, the more eager should be the agent to maintain/obtain such resource.

2. STATE OF THE ART

Another classic concession based model for bilateral multi-issue negotiations is the Agent Based Market Place (ABMP) framework proposed by Jonker et al. (21, 53). ABMP is a negotiation framework, based on additive utility functions, where proposed bids are concessions to previous bids. The amount of concession is regulated by the concession factor (i.e., reservation utility), the negotiation speed, the acceptable utility gap (the maximal difference between the target utility and the utility of an offer that is acceptable), and the impatience factor, which governs the probability of the agent leaving the negotiation process. Additionally, the framework includes other remarkable characteristics such as the possibility of sharing preference information with the other party, and guessing heuristics that allows agents to determine the ranking of issues and issue values based on the bid history.

2.2.3 Similarity Mechanisms in Negotiation

One of the traditional mechanisms proposed in the literature for solving conflicts is the use of similarity mechanisms. They can be used to solve a current conflict based on solutions given to previous conflicts or as mechanisms that implicitly approximate offers to opponents' preferences. Basically, the two similarity mechanisms more widely used are Case Based Reasoning (54, 55) and similarity heuristics (56, 57).

Sycara proposed a mediator that uses case based reasoning for solving conflicts in the labor domain (i.e., PERSUADER) (14, 15, 16). PERSUADER takes as input a set of conflicting goals and outputs an agreed plan of actions. The system keeps track of the agreements found in past negotiations and, once a new conflict situation is present, it looks for the most similar past situation. The retrieved agreement is adapted to the present conflict situation, since the rationale behind this heuristic is that similar conflict situations should yield similar solutions.

2.2 Automated Negotiation

Another popular use of similarity mechanisms is implicitly approximating one's own proposals to the preferences of the opponent. This is usually carried out by means of similarity heuristics that look for trade-offs. A trade-off consists in decrementing the benefit obtained from some negotiations issues that are not important for us but are important of the other agent, in order to get the decremented benefit as an equivalent increase in the benefit obtained by other issues that are important for us but are not important for the other agent. Faratin et al.(56) introduced the use of similarity heuristics in bilateral multi-issue negotiations to compute similarity between pairs of offers. Given a certain utility u demanded by one of the agents, this agent proposes the offer with utility u that is the most similar to the previous offer proposed by the opponent. The idea behind this heuristic is that the more similar the offer is to the previous opponent offer, the more acceptable it is for the opponent. For computing the similarity between two offers, a fuzzy similarity criterion between issue values. The main drawback of fuzzy similarity heuristics is that they require domain knowledge regarding the similarity between issue values for the opponent.

The use of similarity heuristics was reintroduced again by Lai et al (57). In this work, a bilateral negotiation protocol for multi-issue negotiations, where agents are capable of sending up to k different offers per round, is presented. The k offers sent by agents are selected from the iso-utility curve, which contains all of the offers with a certain utility. The offer that is selected is the one that is the most similar to the previous opponent offer that reported the most utility. The other $k - 1$ offers are selected randomly from the iso-utility curve. In this case, the similarity heuristic employed is the Euclidean distance. As a similarity measure, Euclidean distance may be less powerful than fuzzy similarity, but it has the advantage of being more general and not requiring domain knowledge.

2. STATE OF THE ART

2.2.4 Bayesian Learning in Negotiation

Bayesian learning is a probabilistic learning method based on Bayes' theorem (58). Given a certain set of hypothesis and some evidence, Bayesian learning attempts to compute the probability that a certain hypothesis is true after observing the evidence. Bayes is not only a general learning technique for problems where no prior information may be available, but it is also provides mechanisms for updating a model as new information becomes available. Negotiation is a process where information is revealed gradually as the process advances. Therefore, new information needs to be incorporated into agents' negotiation models. This characteristic is what makes Bayesian learning a widely used learning mechanism in automated negotiation.

When reviewing the use of Bayesian learning in negotiation, one cannot forget about the seminal work of Zeng and Sycara (59). In this article, the authors argue about the benefits of using Bayesian models in negotiation. They study a bilateral negotiation case where the buyer attempts to learn the reservation price of a seller by updating its beliefs with Bayesian learning. Despite the fact that it introduces the use of Bayesian learning in negotiation, the applicability of the article is limited since it only focuses on single issue models.

Bayesian classifiers have been used to model the preferences of negotiating agents. In Bui et al. (19), the authors propose a multi-party cooperative negotiation mechanism for the distributed meeting scheduling domain. Agents start from an initial set of possible agreements and jointly look for good collective agreements by partitioning the set of possible agreements in a tree until a set with only one agreement (leaf node) that is acceptable by all of the agents is found. From the joint set of possible agreements, each agent proposes a partition of such set where the final agreement will be looked for. Agents decide on which set should

2.2 Automated Negotiation

be explored from all of the partitions that have been proposed. If all of the agents agree on the partition to be selected, the partition becomes the new joint set of possible agreements and the refinement process continues. Otherwise, agents exchange preferences on the proposed partitions and the partition that maximizes the preferences of the group is selected as the next joint set of agreements. In order to save messages exchanged, the agents employ Bayesian classifiers to learn the preferences of other agents according to the information gathered from the current and past negotiations. Intervals of utility are used as classes and partitions represent attributes of the Bayesian model.

Later, Narayanan et al.(60) present a negotiation framework where pairs of agents negotiate over a single issue (i.e., price). The authors assume that the environment is non-stationary in the sense that agents' strategies may change over time. Non-stationary Markov chains and Bayesian learning are employed to tackle the uncertainty in this domain, and eventually converge towards the optimal negotiation strategy. Non-stationary Markov chains are processes where the next state of the process depends solely on the current state and transition probabilities between states. The main difference with classic Markov chains resides in the fact that transition probabilities change over time. In this negotiation framework, states of the non-stationary Markov chain represent possible strategies that the opponent may use to negotiate for the price. Since transition probabilities are unknown for the agent, a set of candidate transition probabilities become hypothesis of the Bayesian learning process, which is updated each time a new offer is received from the opponent during several negotiations with the same opponent. Based on the estimation of which strategy will be used by the opponent, agents choose the best responding strategy.

Another example of the use of Bayesian learning in negotiation is presented by Buffett et al. (61). In this article, a bilateral negotiation framework is pre-

2. STATE OF THE ART

sented. In the negotiation domain, agents negotiate over a set of limited objects that can be included or excluded from the final deal. It is assumed that for one of the agents (i.e., consumer agent), adding objects to a deal always results in higher utility, whereas for the other agent (i.e., producer agent) subtracting objects from the deal results in higher utility. However, how much each agent values each object may vary, leaving room for integrative bargaining. The negotiation protocol forces agent to send offers that are necessarily a subset of the previous offer in the consumer case, and a superset of the previous offer in the producer case. However, which subset/superset should be selected is not trivial. For that purpose, a Bayesian classifier is employed to classify opponent's preferences into classes of preference relations. A preference relation is a strict preference relation over the objects in the negotiation domain. Groups of similar preference relations are grouped according to the k-means algorithm prior to the negotiation process in order to determine such classes. The classifiers are trained prior to the negotiation by generating random offers that pertain to the different classes, and noting the number of violated preference relations and the true class label. During the negotiation, the negotiation history is compared against the different classes and the number of violated preference relations is used to assess which class is more likely to explain the negotiation history.

Hindriks et al. (52) present a negotiation framework for bilateral multi-issue negotiations where agents' preferences are represented by means of additive utility functions. The main goal of this work is learning a model of the opponent's preferences, and Bayesian learning is used for this purpose. The opponent's preference profile is composed of the importance weights given to each negotiation issue, and the type of valuation function (e.g., monotonically decreasing, monotonically increasing, triangular, etc.). The negotiation framework also assumes that the opponent uses time-based concession strategies, gradually conceding towards the

reservation utility along time. The Cartesian product of the possible orderings of issues and the types of valuation functions become the hypothesis of the Bayesian learning process. Hence, the estimated utility by the concession strategy and the estimated utility by each of the Bayesian hypothesis for the bid history become the core of the learning and updating mechanism used to estimate the opponent's preferences. Since the number of hypothesis, and thus the learning cost, may grow exponentially with the number of negotiation issues, a scalable learning algorithm is introduced where the number of possible orderings for issues is reduced, while still obtaining reasonably good results.

2.2.5 Genetic Algorithms in Negotiation

Genetic Algorithms (GAs) (62) have also contributed to the state-of-art in automated negotiation. They are general optimization and learning algorithms based on the evolutionary processes found in the nature. Candidate solutions for a problem form the genetic population of the algorithm, which gradually converges towards high quality solutions by applying genetic operators like mutation and crossover. GAs are general, which means that they do not rely on a specific problem structure. Additionally, they can be used as an implicit learning and adaptation mechanism in environments where dynamics and structure is also uncertain. This is perhaps what makes GA an adequate approach to negotiation problems, since they can be used to learn and adapt both to the opponent and the environment.

The seminal work of GA's in Automated Negotiation is Oliver et al. (63). They focused on evolving negotiation strategies for bilateral multi-issue negotiations where agents' preferences are represented by means of additive utility functions. In the proposed negotiation framework, a negotiation strategy is a set sequential rules (i.e., rules that are applied in sequential order according to the round), where a rule

2. STATE OF THE ART

is a utility threshold that determines if an offer from the opponent is acceptable and a counter-offer to be made to the other party in case that the opponent's offer is not acceptable. A negotiation strategy is coded as a chromosome. A random population of negotiation strategies (e.g., chromosomes) is generated as initial population of a genetic algorithm for a specific negotiation domain. Each negotiation strategy in the pool of candidates is faced against the different types of opponents and the fitness of the strategy is obtained as the average utility obtained by the strategy against the different opponents. With the purpose of evolving the negotiation strategies and looking for near optimal strategies in a negotiation domain, the strategies with the highest fitness are selected as the parents of the new population, which is created through genetic operators like mutation and crossover. This way, the population of negotiation strategies progressively converges towards a good set of negotiation strategies. Even though the genetic algorithm were able to converge towards reasonably good negotiation strategies, the expressivity of the negotiation strategies in this framework (e.g., simple rules) may be far from the complexity needed in real negotiation problems where hundreds of rounds may be possible, leading to huge exploration space for this negotiation framework.

As commented above, Faratin et al. (18) introduced a negotiation framework for bilateral negotiations where agents' concession strategies can be classified into time-dependent strategies, behavior-dependent strategies and resource-dependent strategies. Matos et al. (64) proposed a framework where the concession to be carried out in each negotiation issue is a linear combination of the concession of the families of concessions strategies proposed by Faratin. The main research goal of Matos et al. (64) is determining which the optimal negotiation strategies in different negotiation environments are. For this purpose, an evolutionary process is proposed where the weights given to the concession strategies for each negotiation issue represent a candidate solution in a genetic algorithm. Populations of sellers

2.2 Automated Negotiation

and buyers with different negotiation strategies negotiate in a round robin way. After each round robin round, negotiation strategies are assigned a fitness value which takes into account the utility obtained in the negotiations and the numbers of messages exchanged. The highest fitness negotiation strategies for sellers and buyers become the parents of the next population of negotiation strategies, which is obtained by the application of genetic operators like mutation and crossover. Eventually, the population of negotiation strategies for sellers and buyers converges towards an optimal set of strategies for the environment under study. The advantage of this proposal with respect to Oliver et al (63) is that the evolutionary process does not depend on the number of negotiation rounds but on the number of negotiation issues, which results in a more tractable search space.

Another authors that have studied genetic algorithms as mechanisms for evolving negotiation strategies are Tu et al. (65). However, the representation employed for negotiation strategies is finite state machines (FSM). According to the representation used by the authors, nodes represent states in the strategy and transitions between states have a precondition and an action associated. The precondition indicates a condition that needs to be satisfied in the last opponent's offer, and the action is the proposal to be sent to the opponent. The evolutionary process is, in essence, the same than the one applied by Oliver et al.(63) and Matos et al. (64): an initial population of negotiation strategies, coded as FSM chromosomes, is generated randomly. After that, the evolution process starts by testing the strategies against several opponents and selecting the highest fitness strategies as parents of the next generation. One of the advantages of using FSM is that they allow branching and states represent certain memory of what has happened in the negotiation process.

Other experiments involving GA and negotiation were carried out by Gerding et al. (66). The authors retake the framework introduced by Oliver et al. focuses

2. STATE OF THE ART

on negotiation processes where the utility function is a linear combination of the issues. The domain of the negotiation issues is $[0,1]$ and one of the agents employs monotonically increasing valuation functions, whereas the other uses monotonically decreasing valuation functions. However, importance weights given by each agent to the negotiation issues may differ, leaving room for integrative bargaining. The main difference between this work and Oliver et al. resides in the introduction of the concept of fairness and social awareness. The former relates to an attempt to avoid unbalanced final agreements, which are common in game theory and negotiations without time pressure due to the use of take-it-or-leave-it strategies by agents in the last round. The fairness check consists in an agent checking an offer and accepting such offer, if it exceeds the reservation value, with a probability related to the utility reported by the offer. Thus, low utility offers have high probabilities of being rejected even though they may be acceptable. The latter, social awareness, refers to the fact that each agent may be able to negotiate with multiple opponents in an agent society. Hence, if a negotiation with an opponent fails, it is still possible to find another deal in the society.

Despite the fact that genetic algorithms have been used mostly for evolving negotiation strategies, there are some works that have proposed the use of genetic algorithms a learning mechanism for opponents' preferences during the negotiation process. Here, we may highlight the seminal work of Krovi et al. (67). The authors propose a bilateral multi-issue negotiation framework where each agent uses a different genetic algorithm each time a negotiation round ends. The population of chromosomes is randomly initialized with 90 random offers and 10 heuristically chosen offers: the last offer of the opponent and the nine best offers from the genetic algorithm executed in the previous round. The fitness function employed for evaluating offers may take into account several factors: one's own utility function, the utility function of the opponent, and one's own negotiation attitude (e.g.,

competitive, cooperative, etc.). After the application of genetic operators like mutation and crossover, the offer with the highest fitness is sent to the opponent as counter-offer.

Choi et al.(68) base their negotiation framework on the idea introduced by Krovi et al., using a GA as learning mechanism each negotiation round. The authors propose the use of multiplicative utility functions for representing agents' preferences, the use of stochastic approximation (69) of the opponent's issue importance by observing concessions, and the use of and adaptive mutation rate that prevents abruptly escaping from high fitness search areas. Finally, Lau (70) also proposes the use of genetic algorithms each round to compute the next offer to be sent to the opponent. The author introduces genetic algorithms in a bilateral framework where the fitness function of an offer is computed based on the utility according to one's own utility function, the Euclidean distance to the last opponent's offer, and a factor that represents time pressure.

2.2.6 Offline Learning in Negotiation

By offline learning we refer to a learning process that is carried out after or before the negotiation process starts. Hence, the model is not updated during the negotiation and it requires of several iterations of the negotiation game to learn an educated model. From the works that we have already reviewed, we can highlight some works like Buffett et al. (61) where the learning of the Bayesian classifiers is carried out before the negotiation starts. However, there are also other approaches that have advocated for the use of learning before or after the negotiation process.

For instance, Coehoorn et al. (71) propose the use of kernel density (72) for the estimation of the weights of the opponent's additive utility function. The negotiation model revolves around the idea that a rational agent gradually concedes towards its reservation utility, and a rational agent should concede less on

2. STATE OF THE ART

the most important issues at the start of the negotiation and concede more on the least important issues. Given this assumption, the agent calculates for each pair of consecutive offers the concession carried out in each issue, and an educated guess of the weight based on such concession. Each tuple, composed of the difference between pairs of consecutive offers, the estimated weight, and the probability density of the weight, forms a three dimensional kernel that is used along the other kernels to calculate an estimation of the real issue weight.

Another set of approaches that heavily rely on offline learning are those approaches based on Artificial Neural Networks (73, 74). In (73), Carbonneau et al. propose a neural network that takes as input the negotiation history of a bilateral negotiation with continuous issues and an offer to make an estimation of the opponent's counter-offer. The major drawback of this approach is that it requires that an artificial neural network is trained per negotiation case. Similarly, the same authors propose an improvement over their previous work in (74). It aims to make a predictive model that does not depend on the negotiation case. The model takes pairs of negotiation issues as inputs of the neural network, where one of the issues is considered the primary issue (i.e., independent variable) and the other issue is considered the secondary issue (i.e., dependent variable). The neural network may also take historical information from each issue like the minimum value, maximum value, average value, etc. The output of such neural network is the predicted value for the issues. The fact that the input is partitioned into pairs of issues, allows the neural network to capture relationships between pairs of issues and how these affect the counter-offer to be proposed by the opponent, which is a much more general approach than taking the whole set of issues as input and learning relationships between all of the issues. This model can be adapted to new negotiation cases since the trained networks is independent of the negotiation

scenario. The main drawback of this work is that the model is not able to capture relationships between issues more complex than binary relationships.

2.2.7 Complex Interdependent Utility Functions

Negotiation processes normally consist in the exchange of proposal between the involved parties. One of the key issues in negotiation strategies is the way in which the agents' preferences are represented. This issue strongly affects how proposals are evaluated and how offers should be generated. In processes where just a single issue is involved, it is quite clear how to evaluate and generate proposals: the value of the issue. However, it is not easy to give a valuation when multiple issues are involved. The multi-attribute utility theory (75, 76) comes into play in this case. This theory provides mechanisms for the evaluation of proposals composed of multiple issues. Classic multi-attribute theory has considered that issues are independent. Issue independence means that the value of negotiation issues does not affect the valuation of other issues. Hence, a classic way of representing such preferences is by means of linear additive utility functions.

Despite the fact that linear additive functions perform well in some simple domains, there are scenarios where they become poorly suited (77). Just as an example, we could think of a water market domain where two parties negotiate over the exploitation of several water resources. One of the parties desires to satisfy its water needs whereas the other party has rights over several water exploitations. In this negotiation, the different issues are the water exploitations to be included in the deal. Even though the provider offers a proposal whose amount of water may satisfy the buyer, the value of the proposal may turn into a low utility for the buyer if the water sources are too distant. Thus, some issues have a negative effect over the value of others, and preferences can no longer be represented as classic linear additive utility functions. There is a need to provide complex utility

2. STATE OF THE ART

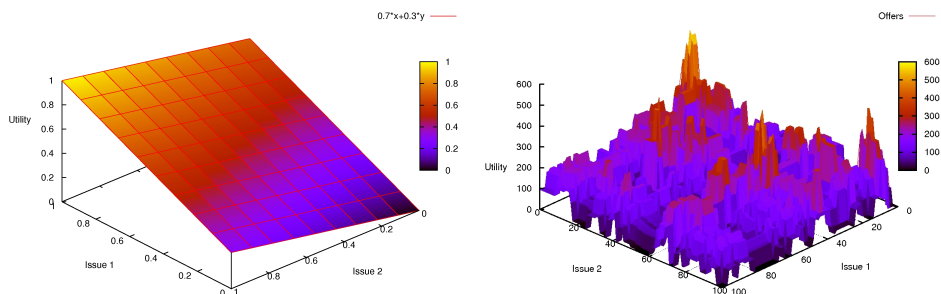


Figure 2.1: An example of linear (left) and complex (right) utility functions

functions that are capable of representing interdependences between negotiation issues. Negotiation strategies that perform well in domains with linear utility functions may not perform equally in the case of complex interdependent utility functions. In fact, the search space for each agent is much more complicated, needing new negotiation strategies adapted to these complex functions. Figure 2.1 shows illustrative examples of the search space in the case of linear utility functions and complex interdependent utility functions. The left figure shows the search space of a two issues linear utility function, whereas the right figure shows the search space of a two issues complex interdependent utility function using the model introduced by Ito et al. (78), which will be reviewed later. As it can be observed, the optimization problem faced by each agent is more complex in the case of interdependent utility functions, leading to the need of new mechanisms that tackle these domains.

Klein et al.(77) constitutes one of the seminal works in negotiation models with complex interdependent utility functions. The authors propose a bilateral negotiation model, which can be extended to the multi-party case, with complex interdependent utility functions for bundles of items that can either be included

2.2 Automated Negotiation

or excluded from the agreement. The preferences of such bundles are modeled by means of preference matrices. The content of a cell matrix represents the increase or decrease in the overall evaluation of the offer when two different issues (i.e., represented by the indices of the matrix cell) are included in an offer. Therefore, this complex utility function can take into account interdependence relations that involve pairs of issues. The negotiation protocol goes as follows. An annealing trusted mediator proposes an offer to all of the parties. Each party can strong/weak reject or strong/weak accept the proposal. If every party accepts the deal, or one of the agents emits a strong accept and the other agent emits a weak reject, the offer becomes an acceptable offer and becomes the base for future offers. Rejected offers can be made acceptable by the mediator due to its annealing mechanism. Then, in subsequent proposals, the annealing mediator mutates the last accepted offer and proposes the modified offer to the agents. The agents may have different strategies: hill climbing or annealing. A hill climber agent only accepts one offer if it is better than the previously accepted offer, whilst annealing agents can accept an offer that is worse than the previously accepted offer with a certain probability that depends on the annealing configuration. The iterated proposal mechanism continues until a fixed number of offers have been proposed.

Robu et al.(79, 80) introduce a bilateral negotiation model where agents represent their preferences by means of utility graphs. The negotiation domain is formed of bundles of items that can be either included or excluded in a final deal. Utility graphs are graphical models that relate negotiation issues that are dependent. Nodes represent negotiation issues whereas arcs connect issues that have some joint effect on the utility function (i.e., positive for complementary issues, and negative for substitutable issues). Hence, utility graphs represent binary dependencies between issues. The authors propose a negotiation scenario where the buyer's preferences and the seller's preferences are modeled through utility graphs.

2. STATE OF THE ART

The seller is the agent that carries out a more thorough exploration of the negotiation space in order to search for agreements where both parties are satisfied. With this purpose, the seller builds a model of the buyer's preferences based on historic information of past deals and expert knowledge about the negotiation domain. This model is updated during the negotiation based on the bids exchanged with the buyer. The model also introduces a proposal strategy based on utility graphs that is capable of selecting which offers are more adequate based on the model of the other agent's preferences and one's own preferences.

Later on, Ito et al. (78, 81) propose a multi-party negotiation model where agents have their preferences represented by means of weighted hyper-cubic constraints. The negotiation domain is composed of multiple issues whose domain is found in the integer domain. A utility function is composed of several constraints. A constraint is composed of n different issues, one value interval for each of the n issues, and a utility value u . The constraint is fulfilled by an agreement when, for each issue present in the constraint, the issue interval contains the issue value of the agreement. In that case, since the constraint is fulfilled, it adds u to the overall utility of the agreement. The utility of the agreement is the sum of the utilities reported by each constraint. The main difference between the preference model proposed by Ito et al. and the previous approaches is that it is able to capture dependencies for more than two issues. The authors also propose a negotiation protocol and negotiation strategies. Since agent preferences are complex and interdependent, each agent samples its own search space to find high utility agreements. These agreements are discovered by a process of offer sampling followed by a simulated annealing started from each of the offers sampled. After the annealing process, each agent forms several constraint bids based on optimized contracts and a bid value which represents the total value reported by fulfilled constraints. A trusted mediator receives bids from all of the agents and attempts

to find the contract that maximizes social welfare. Marsa-Maestre et al. (82, 83) carry out further research in the area of negotiation models for complex utility functions. More specifically, they extend the constraint based model proposed by Ito et al. (78, 81) by proposing different bidding mechanisms for agents. They also propose a negotiation protocol that may not be *one-shot*. In fact, the mediator can suggest the relaxation of some constraint bids in order to increase the probability of finding an agreement.

2.3 Negotiation Teams

The literature in human negotiation led us to the discovery of another potential scenario where complex negotiation can take place: negotiation teams.

A negotiation team is a group of two or more interdependent persons who join together as a single negotiating party because their similar interests and objectives relate to the negotiation, and who are all present at the bargaining table (24, 26). Hence, a negotiation team is a negotiation party that is formed of multiple individuals instead of just one individual. As a negotiation party, the team negotiates with other parties in order to reach a final agreement.

In what kind of scenarios may negotiation teams be involved? There are several scenarios whose importance range from day to day negotiations to crucial negotiations like the ones found in business and politics. For instance, we can think of the following negotiation cases where teams participate in real life:

- Imagine that a married couple want to purchase a car (25). For that matter, the couple has to negotiate with a car seller the purchasing conditions like price, payment method, and extras included in the contract. Clearly, this is an agreement that is signed between two parties: the couple, and the car

2. STATE OF THE ART

seller. However, one of the parties is clearly composed of two individuals (i.e., the couple) that share the same goal (i.e., buy a car).

- Imagine that a group of four friends decides to go on a travel together. If a travel needs to be arranged, the group of friends needs to find an adequate destination, some nice accommodation and flights. Additionally, it may even be interesting to include some pre-arranged social activities like visits to museums, some sport activities, and so forth. There may be several travel agencies that offer such services, and the group of friends may need to negotiate with some of them to get a travel package that satisfies their needs. As in the case of the couple, the group of friends is one single negotiation party that is composed of multiple individuals that share a common objective (i.e., go on a trip together).
- In another scenario, a human organization desires to sell a product line to another company (25). It is usual for each company to send a negotiation team composed of different experts coming from different organizational departments. This team is entrusted with the task of understanding the complex scenario at hand, and taking the most adequate course of action for their principals. It is unnecessary to highlight that, obviously, in this case both parties are also composed of multiple individuals.
- Similarly to the scenario mentioned above, negotiations in politics also involve negotiation teams. We could for instance think of the negotiations carried out between Cambodia, Laos, Thailand and Vietnam for promoting cooperation on water resources (84). In these negotiations, each national party formed a negotiation team that participated actively during the negotiation process. Each team was formed of different specialists.

Thus, it can be appreciated that negotiation teams are common in real life negotiations. Despite their importance in real negotiations, teams have not been studied by social sciences to the same extent as dyadic negotiations (85, 86, 87). However, what are the reasons to send a negotiation team to the negotiation table instead of a single negotiator? The first reason that may come to our minds is that the more individuals may mean more cognitive capabilities and, therefore, better task performance in the search process involved in negotiation. Effectively, this was shown by Thompson et al. (24) where several experiments involving human negotiation teams determined that as long as one of the parties is a negotiation team, better joint outcomes (i.e., integrative outcomes) were obtained. This is partially explained due to the fact that when teams are present at the negotiation table, parties are more inclined to exchange information (24).

Another reason to send a negotiation team is skill distribution and information distribution (10, 25, 88). With this, we mean that different team members may have different and knowledge complementary skills needed to tackle properly the negotiation. Working as a team allows to discover such specializations and learn to take advantage of them (25). Thompson (10) recommends that managers should recruit negotiation teams composed of experts in negotiation, experts in the subject to be negotiated, and individuals with a variety of interpersonal skills. Mannix (88) states that negotiation teams require a diverse set of knowledge, abilities, or expertise in complex negotiations, and points out the correct assessment of such skills as one of the keys for success in a negotiation. Skill distribution and complementary skills are of vital importance when using some classic team negotiation tactics like the good cop/bad cop persuasion tactic (89).

Finally, other authors consider that another reason to send a negotiation team are stakeholders. The entity may be formed of different members whose interests have to be reflected in the final agreement (88, 90). For instance, Mannix

2. STATE OF THE ART

(88) points out union negotiation as an example of negotiations where parties are formed of different interests that have to be represented in the negotiation table. Halevy (90) also remarks the importance that despite negotiation teams being a single negotiation party, they are hardly ever a unitary player. In fact, a negotiation team is usually a multi-player party with different and possibly conflicting interests.

2.4 Conclusions & Discussion

In this section we try to identify some of the open issues in both of our fields of interest: negotiation models for Ambient Intelligence and agent-based negotiation teams.

2.4.1 Ambient Intelligence

As stated in the introduction and motivation of this thesis, one of our objectives is providing computational models for negotiations that are carried out by agents in devices with scarce resources like mobile phones, pdas and smartphones. In these types of devices, an efficient use of the computation time, the memory usage, and the bandwidth is crucial. In this discussion, we analyze how each reviewed work would fit in the Ambient Intelligence domain. The analysis is based on different criteria. The criteria are composed of factors that are interesting for every negotiation domain, and factors that are especially interesting for the Ambient Intelligence domain.

The number of negotiation issues tackled by the model is a very important factor. Negotiation processes are usually complex by nature and they should include several issues. In this sense, we can categorize negotiation models into single issue or multi-issue, preferring the latter when possible. Another important factor in

2.4 Conclusions & Discussion

negotiation models is whether or not they are mediated. It is true that mediators may help to reach better agreements, however, they also require the existence of an entity that is trusted by all of the parties. Non-mediated negotiation models are more interesting from the point of view of Ambient Intelligence due to the fact that users enter and leave the system in an extremely dynamic way. Thus, it may be difficult to find a trusted mediator for every user.

Learning is also a very important issue in any negotiation model. First of all, learning may help agents to find better agreements (either win-win agreements or individually good agreements). Second, by learning agents may end negotiations more quickly, thus reducing the computation time, the energy, and the bandwidth spent by the negotiation. Therefore, learning is a desired feature in any negotiation model and, of course, in negotiation models for Ambient Intelligence. When analyzing different learning mechanisms in negotiation models, one can observe differences in the object of learning. Some models attempt to learn the opponents' preferences (e.g., issue rankings, issue weights, best offers for the opponent, etc.), other models try to learn the optimal negotiation strategy, others aim to predict opponents' responses, and so forth. Another characteristic of learning mechanisms is the data source for the learning process. It can either come from the present negotiation or from a history of negotiations. In an Ambient Intelligence environment, where users enter and withdraw from the system in a very dynamic way, agents may only face opponents once or a few times. Thus, learning mechanisms that rely mainly on the current negotiation are preferred for Ambient Intelligence domains. Finally, another important consideration is whether the learning model can be easily adapted in the presence of new data or it needs to undergo a new learning process. Given the limited computational resources, learning mechanisms that are easily adapted are preferred over learning mechanisms that need new training in the presence of new data.

2. STATE OF THE ART

Complex interdependent utility functions are able to represent richer preferences for users. This should provide with more accurate preference models that, when used by negotiation strategies efficiently, should end up with agreements that are more satisfactory for users. Thus, negotiation domains should benefit from the uses of complex and interdependent utility functions. Regarding interdependence relations among negotiation issues, we pay attention to the cardinality of such relations. Utility functions that are able to capture interdependence relations among more issues should result in more powerful and flexible models since they may be able to explain more complex scenarios. Table 2.1 gathers the analysis on the aforementioned criteria for the negotiation models reviewed in Section 2.2. Next, we discuss which approaches are more adequate for our goals.

Most of the models reviewed consider multiple issues, which fits our goals. Only Zeng et al. (59) and Narayanan et al. (60) focus on single issue negotiations that involve price. From those negotiation models that consider complex interdependent utility functions, we have been able to observe that most of them focus on mediated processes (77, 78, 81, 82, 83), which we argue that should be avoided in Ambient Intelligence domains. Only Robu et al. (79, 80) consider complex interdependent utility functions and their model does not require the presence of a mediator. However, their model only captures dependencies between pairs of issues. Therefore, there seems to be a dearth in non-mediated negotiation models that work with complex interdependent utility functions, which should be our goal for Ambient Intelligence domains. In terms of, interdependence cardinality, the works of Ito et al. (78, 81) and Marsa-Maestre et al. (82, 83) are the only ones to consider interdependence relations more complex than binary ones.

Concerning learning mechanisms, in general it can be observed that those models that cannot be easily adapted are those models that require data from multiple negotiations (71, 73, 74). Other models can be easily adapted but require data

	General		Learning		Adaptive?	Interdependence Cardinality
	# Issues	Mediated?	Learning Object	Data Source		
Faratin et al. (18)	Multi	No	-	-	-	-
Jonker et al. (21, 53)	Multi	No	Opponent preferences	Current negotiation	Yes	-
Sycara et al. (14, 15, 16)	Multi	Yes	Conflict resolution	History of negotiations	Yes	-
Faratin et al. (56)	Multi	No	Opponent preferences	Current negotiation	Yes	-
Lai et al. (57)	Multi	No	Opponent preferences	Current negotiation	Yes	-
Zeng et al. (59)	Single	No	Reservation price	Current negotiation	Yes	-
Bui et al. (19)	Multi	No	Opponent preferences	History of negotiations	Yes	-
Narayanan et al. (60)	Single	No	Optimal strategy	History of negotiations	Yes	-
Buffett et al. (61)	Multi	No	Opponent preferences	Current negotiation	Yes	-
Hindriks et al. (52)	Multi	No	Opponent preferences	Current negotiation	Yes	-
Oliver et al. (63)	Multi	No	Optimal strategy	History of negotiations	Yes	-
Matos et al. (64)	Multi	No	Optimal strategy	History of negotiations	Yes	-
Tu et al. (65)	Multi	No	Optimal strategy	History of negotiations	Yes	-
Gerding et al. (66)	Multi	No	Optimal strategy	History of negotiations	Yes	-
Krovi et al. (67)	Multi	No	Opponent preferences	Current negotiation	Yes	-
Choi et al. (68)	Multi	No	Opponent preferences	Current negotiation	Yes	-
Lau et al. (70)	Multi	No	Opponent preferences	Current negotiation	Yes	-
Coehoorn et al. (71)	Multi	No	Opponent preferences	History of negotiations	No	-
Carbonneau et al. (73, 74)	Multi	No	Opponent response	History of negotiations	No	-
Klein et al. (77)	Multi	Yes	-	-	-	Binary
Robu et al. (79, 80)	Multi	No	Opponent preferences	Current negotiation	Yes	Binary
Ito et al. (78, 81)	Multi	Yes	-	-	-	n-ary
Marsa-Maestre et al. (82, 83)	Multi	Yes	-	-	-	n-ary

Table 2.1: Categorization of computational negotiation models.

2. STATE OF THE ART

from multiple negotiations, which is also not desirable (14, 15, 16, 19, 63, 64, 65, 66) for Ambient Intelligence domains. Thus, similarity heuristics (56, 57), most multi-issue Bayesian learning approaches (52, 61) and genetic algorithms used as an implicit learning mechanism of opponents' preferences (67, 68, 70) seem the most appropriate learning mechanisms for Ambient Intelligence domains. Nevertheless, some considerations have to be taken into account among these learning mechanisms. Most Bayesian learning approaches have been devised for linear utility functions with no dependence among negotiation issues. This same assumption has been used to relax the learning cost in Bayesian approaches. It is expected that if issue dependencies are to be considered, the learning cost of Bayesian approaches will explode compared to other learning mechanisms. Hence, it seems that similarity heuristics and genetic algorithms seem more appropriate. Both are general and implicit learning mechanism of the opponents' preferences, which should be able to handle issue interdependence with relative little effort. On top of that, genetic algorithms are also search and optimization mechanisms, which could help in the exploration of one's own complex utility function.

In conclusion, we have observed that none of the current models perfectly fits the requirements of Ambient Intelligence domains: non-mediated protocols, complex interdependent utility functions, and adaptive learning mechanisms that rely on the current negotiation. Among the latter, we believe that similarity heuristics and genetic algorithms may be considered the most appropriate learning mechanisms. In Chapter 3 we propose a non-mediated bilateral negotiation model with complex interdependent utility functions that aims to cover the necessities of Ambient Intelligence. For that purpose, it relies on similarity heuristics and genetic algorithms.

2.4.2 Agent-Based Negotiation Teams

Up to this point, we have strictly considered human negotiation teams and the advantages for humans to negotiate as a team. But, are negotiation teams also feasible and needed in automated negotiation and electronic applications? We argue that the answer to such a question is positive. Agent-based systems are not alien to negotiation scenarios where it may be interesting to employ negotiation teams. For instance, imagine a tourism e-market application. It is usual for groups of friends/families, or even strangers, to organize their holidays as a group. However, travelers usually have different preferences regarding trip conditions (e.g., cities to visit, hotel location, leisure activities, number of days to spend, budget limitations, etc.). Humans may be extremely slow at coming with a proper negotiated deal that accounts for everyone's preferences. Thus, software agents representing each traveler could form a negotiation team that negotiates with travel agencies in an e-market to obtain a quick and good trip package for the group. The application of negotiation teams is not limited to the aforementioned example. It can be extrapolated to other domains such as electronic farming cooperatives, customer coalitions, negotiation support systems for labor negotiations, and so forth. Thus, there is a need for agent-based negotiation team models.

A trusted mediator with perfect knowledge regarding the group of travelers' preferences or a trusted mediator who can aggregate preferences can be thought of as possible mechanisms to coordinate a negotiation team. Nevertheless, there are several reasons that preclude us from aligning ourselves with this kind of coordination mechanisms. The first reason is that privacy is usually a concern among users in electronic applications. In fact approximately 90% of the users in electronic applications care to some degree about the amount of information that they filtrate in electronic applications, and only 10% do not care about letting others

2. STATE OF THE ART

manage their information (91). Hence, one cannot expect that every team member may be willing to share its full preferences to a mediator. The other important issue is the fact that even though there may be some degree of cooperation among team members, one should not forget that the team is a multi-player party and opportunistic behavior may be present. In that case, preference aggregation is a dangerous mechanism since it is quite prone to being manipulated and exaggerated for one's own benefits. Therefore, other types of mechanisms are needed to coordinate agent-based negotiation teams. These mechanisms, which we have coined as intra-team strategies, should reflect the preferences of team members in the final agreement. For that reason, we think that unanimity regarding team decisions is a very important factor when designing intra-team strategies. Agreements that are unacceptable for a team member should be avoided since they might deteriorate human relationships. Furthermore, technologies that help to form unanimous decisions may provide more user satisfaction, and they can help team members to avoid unexpected outcomes. Despite the fact that several negotiation models have been proposed in MAS, as far as we know, this thesis is the only work that has considered agent-based negotiation team so far.

We can make an analogy between what we have commented regarding human negotiation teams and agent-based negotiation teams. Basically, we can define an agent-based negotiation team as a group of two or more interdependent agents that join as a single negotiation party because of their shared interests in the negotiation with some opponents. The reasons to use an agent-based negotiation team are also analogous to the human case. First, more agents in the team may mean more computation capabilities and, thus, more extensive and parallelized exploration of the negotiation space. Second, we can also assume that different and heterogeneous agents may have different experiences, they may offer different services/skills, they may implement different algorithms, which in the end results

2.4 Conclusions & Discussion

in teams being able to tackle complex negotiation problems more efficiently (in Chapter 4 we will analyze some factors that may play in detriment of negotiation teams). Third and lastly, as in the human case, the team may really represent a multi-player party whose preferences need to be satisfied as much as possible by the final agreement between the team and opponents.

Nevertheless, human negotiation teams do not always guarantee a better outcome than individuals. The performance of the team is directly related to coordination among team members, and a team that is not capable of achieving such coordination may fail at the negotiation. In fact, Behfar et al. (87) study the causes that pose problems for human negotiation teams: logistics and communication problems (e.g., communications inefficiencies), substantive differences (e.g., confusion about goals, conflicting interests), inter-personal and personality differences (e.g., different negotiation styles), and confusion about team roles (e.g., unclear decision rights). The same authors also identify those strategies that help to overcome the aforementioned problems and lead teams towards success: time and logistics management (e.g., coordinating strategies during negotiation by stepping away from the table), team communications (e.g., preparing with teammates), within-team negotiations (e.g., team problem solving, managing conflicting interests), and defining leadership and team roles (e.g., defining decision rights). To put it briefly, communications, coordination, intra-team negotiation, and clear rules of the game lead human negotiation teams to success. We believe that those key elements are also important in agent-based negotiation teams. For that reason, we put a special emphasis in analyzing the tasks to be carried out by successful agent-based negotiation teams in Chapter 4. We describe which tasks should be carried out, relate each specific task to similar research in multi-agent systems, and point out some interesting issues that may arise in due to the nature of agent-based negotiation teams. Intra-team strategies, especially those that

2. STATE OF THE ART

guarantee unanimity regarding team decisions, are proposed in Chapters 5 and 6 of this thesis.

3

Bilateral Negotiation for Limited Devices

3.1 Introduction

In this chapter, a **non-mediated** bilateral multi-issue negotiation model for AmI environments is presented. Its main goal is to optimize the computational resources while maintaining a good performance in the negotiation process. The proposed model is inspired by the seminal work of Lai et al. (57). The work of Lai et al. presents a non-mediated strategy for general utility functions, which obviously includes complex utility functions (one of the necessities identified in Chapter 2). The strategy is based on the calculation of current iso-utility curves and a similarity heuristic that sends offers from the current iso-utility curve that are the most similar to the last offers received from the opponent. However, the entire calculation of the iso-utility curve may require an exhaustive exploration of the utility function, which may not be tractable in the case of a large number of issues. Furthermore, if the exploration of one's own utility function is not per-

3. BILATERAL NEGOTIATION FOR LIMITED DEVICES

formed in an intelligent way, the result may be that most of the offers sampled are of no use for the negotiation process since they might not interest the opponent. This behavior is not desirable in AmI environments and we tackle these problems in the negotiation model proposed in this chapter. The three main differences between this present work and the work of Lai et al. are: (i) The present approach assumes that it is not possible to exhaustively search the utility function. Before the negotiation process starts, each agent samples its own utility function by means of a niching genetic algorithm (GA) (92, 93). The effect of this sampling is that offers obtained are highly fit and significantly different; (ii) A few additional samples are obtained during the negotiation process by means of genetic operators that are applied over received offers and one's own offers. The heuristic behind this sampling is that offers obtained by genetic operators have genetic material from one's own agent and the opponent's offers. Thus, these new offers may be interesting for both parties. (iii) Genetic operators and similarity heuristics act as a learning mechanism that implicitly guides the offer sampling and selection of which offers must be sent to the opponent.

This chapter is organized as follows: section 3.2 describes an example of application where automated negotiation and Ambient Intelligence can be combined in order to offer a useful service for the user; section 3.3 describes the negotiation model, explaining the chosen protocol and the new negotiation strategy in detail. In Section 3.4, the experimental setting and the results obtained are discussed. Finally, the conclusions are explained in Section 3.5.

3.2 Motivating Scenario: Product Fairs

In this section we introduce an example of application where automated negotiation may be used along with well-known AmI technologies in order to provide

3.2 Motivating Scenario: Product Fairs

a profitable service for users. The example focuses on product fairs. Fairs are public events where sellers exhibit their products to a wide range of consumers. At this kind of events there are usually a large number of exhibitors and products. Therefore, it is extremely difficult to explore the whole fair or find interesting deals for one's interests. It is also difficult for sellers to attract interesting clients. Thus, both consumers and sellers would benefit from a tool which allows to attract/search prospective deals quickly. Furniture fairs are very popular, especially in Valencia's region. Even though our negotiation model has not been specifically designed for furniture fairs, we use it in the examples that describe this chapter.

At this point, automated negotiation in an AmI environment may come in handy. Let us suppose the following scenario at a furniture fair: each vendor has been assigned a booth where he attends to clients. As well as setting up the typical equipment, a hardware device with Bluetooth wireless communication is provided (e.g. a personal computer). An agent, which can be downloaded and configured by the vendor prior to the fair, is installed in this hardware device. These agents should be provided with information regarding its owner's preferences by means of user modeling methods such as questionnaires, past experiences, and so forth.

Additionally, consumers are allowed to download an agent to their mobile devices prior to the fair event. The only requirement for the mobile device is Bluetooth wireless capabilities. The consumer's agent can be configured similarly to the vendor's agent. More specifically, the agent may ask what products its owner would be interested in buying and general questions about the preferences regarding possible negotiation issues.

When consumers and vendors enter the fair, they should start the execution of their respective agents. Each consumer agent offers a negotiation service which can be invoked by vendor agents. Whenever this service is invoked by a vendor agent, a negotiation process starts between the vendor agent and the consumer agent.

3. BILATERAL NEGOTIATION FOR LIMITED DEVICES

The negotiation process continues until a deal has been found or the consumer has exited the Bluetooth coverage area of the vendor. If the deal is considered as interesting by both parties (i.e. utility of the deal higher than a certain threshold or reservation utility) and the deal is among the best ones for the consumer in that specific area (determined by which vendors can be reached by Bluetooth in that space point), the consumer agent and the vendor agent notify their respective owners regarding the possible deal. However, deals discovered by this automatic process are not to be considered as binding but as recommendations. If the deal is considered as interesting enough by the consumer, it may result in the consumer approaching the vendor's booth. At that point, both parties may decide to renegotiate or polish the deal which has been found by their agents.

Since Bluetooth technology has coverage limitations, the service can usually only be discovered by vendor agents that are nearby. Usually, the range of communications for Bluetooth devices goes from 5 to 10 meters (some devices may be able to reach 100 meters, but they consume more energy). Therefore, negotiation processes help consumer and vendor agents to find prospective deals as consumers walk around the fair. These negotiations have to be as quick as possible to avoid the consumer from exiting the covering range of vendors, and they also have to save mobile devices' energy by limiting the number of communications. An illustrative and simplified example of this application can be observed in Figure 3.1. Consumer 1 is in Bluetooth range of vendors 2 and 5, whereas consumer 2 is in range of vendors 4 and 7. Thus, consumer 1 agent can negotiate with vendor 2 agent and vendor 5 agent, and consumer 2 agent may negotiate with vendor 4 and 7.

The possible benefits of the proposed application can be summarized in: (i) it allows consumers to save physical time by filtering the vendors that seem more suitable for their needs as they walk around the fair; (ii) it also helps vendors

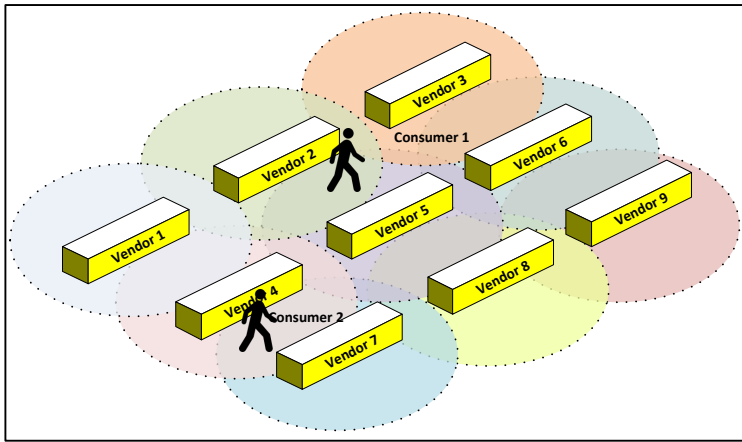


Figure 3.1: An example of the furniture fair application

since it attracts consumers that will probably be interested in buying their goods, instead of losing time with clients with whom the possibilities of making a good deal are very low.

3.3 Negotiation Model

As it can be appreciated in the motivating scenario, the application is collaborative in nature. This will be reflected in the negotiation model employed by agents. It is very important for sellers and buyers to find good deals quickly given the dynamic nature of the negotiation.

Negotiation models are composed of a negotiation protocol and a negotiation strategy. On the one hand, the negotiation protocol defines the communication rules to be followed by the agents that participate in the negotiation process. More specifically, it states at which moments the different agents are allowed to send messages and which kind of messages the agents are allowed to send. For instance, the Rubinstein alternating protocol specifies (94) that agents are allowed

3. BILATERAL NEGOTIATION FOR LIMITED DEVICES

to send one offer in alternating turns. Basically, the negotiation protocol acts as a mechanism for the coordination and regulation of the agents that take part in the negotiation process.

On the other hand, the negotiation strategy defines the different decisions that the agent will make at each step of the negotiation process. It includes the opponent's offers acceptance rule, the selection of which offers are to be sent to the opponent, the concession strategy, the decision of whether the agent should continue in the negotiation process or not, and so forth. Therefore, the negotiation strategy includes all the decision-making mechanisms that are involved in the negotiation process.

The negotiation protocol used can be categorized as an alternating protocol for bilateral bargaining (94). More specifically, the protocol used is the *k-alternating protocol* proposed by Lai et al. (57). In our setting, we assume that agents do not know other agents' preferences, nor they know the strategies carried out by agents and the exact conditions of the negotiation environment (i.e., incomplete information setting). Additionally, the special characteristics of the devices where agents are executed define a tightly bounded computational environment (i.e., bounded rationality). The proposed negotiation strategy is composed by a time-based concession strategy and an offer proposal strategy that belongs to the family of negotiation strategies that use a similarity heuristic in order to propose new offers to the opponent (56, 57).

3.3.1 Negotiation Protocol

As mentioned above, the negotiation protocol belongs to the family of alternating protocols for bilateral bargaining. In this kind of protocols, two different agents negotiate without the need of a mediator. Non-mediated strategies are more adequate for AmI applications since users enter and leave the AmI system in a very

dynamic way. Thus, it may not be feasible to find a trusted mediator for every possible pair of agents. Furthermore, in some AmI domains such as shopping malls, where there are different competing vendors and lots of potential users, it is difficult to determine who will mediate the negotiation process.

The protocol used is the *k-alternating protocol* proposed by Lai et al. (57). This protocol is composed of several rounds where the agents exchange offers in an alternating way. One of the agents, called the *initiator*, is responsible for starting the current round. He can accept one of the offers received from the opponent in the last round, exit from the negotiation process, or send up to k different offers to the opponent agent. Once the *initiator* has performed one of the possible actions, the opponent agent is able to accept one of the offers he has just received, exit from the negotiation process or propose up to k different offers to the *initiator*. Then, the round ends and a new round is initiated by the *initiator* agent. The negotiation process ends when one of the agents accepts an offer (the negotiation succeeded) or one of the agents decides to abandon the negotiation (the negotiation failed).

Some of the properties of the *k-alternating protocol* proposed by Lai et al. are:

- The protocol is adequate for situations where both agents are equal in power (e.g. none of them has the monopoly over a resource).
- Each agent is capable of sending up to k different offers, making it more probable that one of the proposed offers satisfies the requirements of the opponent agent.
- Since k different offers are proposed in each turn more information about the opponent preferences can be inferred. This may produce faster agreements, which is inherently interesting for every domain but particularly for AmI domains since it may reduce the number of messages exchanged and thus

3. BILATERAL NEGOTIATION FOR LIMITED DEVICES

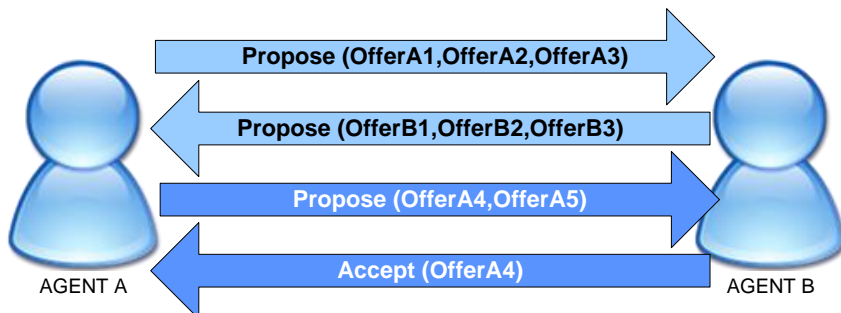


Figure 3.2: An example of two agents negotiating in the k -alternating protocol.

the bandwidth consumption. Additionally, learning the preferences from a complex utility function with dependences between issues is a hard task that requires more information.

An example of two agents negotiating with a 3 -alternating protocol ($k = 3$) can be observed in Figure 3.2. Agent A is the initiator of the negotiation round, whereas Agent B is the responding agent. The first round starts with 3 offers proposed by the initiator. Once the offers reach Agent B, he decides whether he should accept one of them or not. Since the 3 offers are not interesting for Agent B, he decides to counteroffer 3 different offers. Due to the fact that none of the offers proposed by Agent B are of interest to the initiator, he decides to send 2 offers. The 2 offers from the initiator reach Agent B, who analyzes the offers in order to determine whether they are interesting. Since he found OfferA4 to be interesting, he decides to accept it and the protocol thus ends with an agreement.

3.3.2 Negotiation Strategy

Agents follow a negotiation strategy based on a time-based concession strategy and a proposal strategy that employs similarity heuristics to propose new offers

to the opponent (56, 57). Both types of agents employ this type of collaborative approach due to the characteristics of the motivating example.

The negotiation strategy complements some of the benefits introduced in the inspiring work of Lai et al. (57), making it especially interesting for AmI environments. The goal is to optimize the computational resources while maintaining a good performance in the negotiation process. The main traits of the proposed model are twofold. Firstly, it is not necessary to sample the entire utility function. Secondly, the proposed strategy provides an implicit learning mechanism that guides the offer sampling and which of the offers sampled are to be sent to the opponent.

The different mechanisms of the negotiation strategy can be grouped according to the period during which they are applied: pre-negotiation and negotiation. In the pre-negotiation, since utility functions are complex and dependencies exist between negotiation issues, even an agent does not know which offers entail good quality for itself without previous exploration. Since the negotiation is carried out in limited devices, it is not feasible to completely explore the whole set of possible offers. Therefore, each agent samples its own utility function by means of a niching GA (*self-sampling*).

The mechanisms used during the negotiation include the acceptance criteria for opponent's offers, the concession strategy, the *evolutionary sampling*, and the selection of which offers are sent to the opponent. The most remarkable part is introduced with *evolutionary sampling*: genetic operators are carried out over received offers and one's own good quality offers in order to sample new offers that may be of interest to both parties. *Evolutionary sampling* acts as an implicit learning mechanism of the opponent's preferences. The result of evolutionary sampling may be used afterwards when the offers to be sent to the opponent are selected. A brief outline of the proposed strategy can be observed in Algorithm

3. BILATERAL NEGOTIATION FOR LIMITED DEVICES

1. A more detailed outline of the strategy used before the negotiation process and during the negotiation process can be observed in Algorithms 2 and 3.

Negotiation Strategy;

Pre-negotiation;

1. Self-sampling;

Negotiation Process;

2. Receive opponent offer(s) if there are any offers;

3. Acceptance criteria: accept an offer and end the negotiation, or reject all of them and continue the negotiation process;

4. Concession strategy;

5. Evolutionary sampling;

6. Select which offers to send;

7. Send offer(s) and go to step 2;

Algorithm 1: A brief outline of the evolutionary negotiation strategy.

3.3.2.1 Pre-negotiation: Self-sampling

When an agent uses complex utility functions with issue dependencies, it may be difficult to find offers with good utility for oneself. When the negotiation domain is not large, a complete sampling of the utility function may be feasible. However, when the domain is large, which is usually the case in real negotiations, a complete sampling may be an extremely expensive process. For instance, a complete sampling of a negotiation domain formed by 10 integer issues from 0 to 9 requires sampling 10^{10} offers. The cost associated to this sampling can be exorbitant, especially if agent preferences change with a frequency that is greater than the time invested in the sampling. Furthermore, this sampling is unacceptable for AmI domains. Not only does it take too much computational time and power, but it would also need too much storage for the limited devices usually found in

AmI.

A possible solution to this problem is to use mechanisms that enable an agent to sample good offers for the negotiation process and skip those of low quality. Due to the highly non-linear nature of complex utility functions, non-linear optimizers are required for this task. The main goal is to sample a set of different offers that have good utility and are significantly different, because these offers may point to different regions of the negotiation space where a good deal may be found for the agent.

In this work, a genetic algorithm (GA) was used to solve this problem. GA's are general search and optimization mechanisms based on the Darwinian selection process for species (62, 92). Genetic operators such as crossover, mutation, and selection are employed in order to find near-optimal solutions for the required problem. Nevertheless, the problem posed by classic GA's is that the entire population converges to one optimal solution. As already stated, different interesting offers for the negotiation process need to be explored. Niching methods are introduced to confront problems of this kind (93, 95). These methods look to converge to multiple, highly fit, and significantly different solutions.

A possible family of niching methods for GAs is the crowding approach (95). Crowding methods achieve the desired result by introducing local competition among similar individuals. One advantage of crowding methods is that they do not require parameters beyond classic GAs. Euclidean distance is usually used to assess the similarity among individuals. Probabilistic Crowding and Deterministic Crowding (95) are two of the most popular crowding methods. They only require a special selection rule with respect to classic GAs. Both rules are employed to select a winner given n different individuals. On the one hand, Deterministic Crowding selects the individual that has the highest fitness value, resulting in an elitist selection strategy. On the other hand, Probabilistic Crowding allows lower

3. BILATERAL NEGOTIATION FOR LIMITED DEVICES

fitness value individuals to be selected as winners with a certain probability. This probability is usually proportional to the fitness of each individual. In general, Probabilistic Crowding is more exploratory than Deterministic Crowding. In both cases, the niching effect is achieved by applying either of the two methods to those individuals that are similar. Each parent is usually paired with one of its children in such a way that the sum of the distances between pair elements is minimal. For each pair, one of the two crowding methods is employed to determine which individuals will form the next generation. In this work we define D_C as our Deterministic Crowding rule and P_C as our Probabilistic Crowding rule. Both rules can be observed in more detail in Equations 3.1 and 3.2, respectively.

$$D_c(X_1, X_2) = \begin{cases} X_1 & U(X_1) > U(X_2) \\ X_2 & U(X_1) < U(X_2) \\ X_1 \vee X_2 & \text{other} \end{cases} \quad (3.1)$$

$$P_c(X_1, X_2) = \begin{cases} X_1 & U(X_1) > U(X_2) \wedge \text{random}() \leq p_1 \\ X_2 & U(X_1) > U(X_2) \wedge \text{random}() > p_1 \\ X_2 & U(X_1) < U(X_2) \wedge \text{random}() \leq p_2 \\ X_1 & U(X_2) < U(X_1) \wedge \text{random}() > p_2 \\ X_1 \vee X_2 & \text{other} \end{cases} \quad (3.2)$$

$$\text{with } p_i = \frac{U(X_i)}{U(X_i) + U(X_{i'})}$$

where $\text{random}() \in [0, 1]$, $U(\cdot)$ is the fitness function of our genetic algorithm, which in our case corresponds to the utility of the offer. X_1 and X_2 are two offers, and p_1 and p_2 are the probability of acceptance of both offers by Probabilistic Crowding.

Self-sampling uses a GA that employs crowding methods to find significantly different good offers. This GA is individually executed by the agent before the negotiation process begins. The chromosomes of this GA represent possible offers in the negotiation process, whereas the fitness function used is one's own utility

function. A portfolio with D_C and P_C is used. The population has a fixed number of individuals and the whole population is selected to form part of the genetic operator pool. Pairs of parents are selected randomly and multi-point crossover or mutation operators are applied over them. In both cases, the result is two children. To apply crowding methods, each parent is paired with the child that is the most similar to it according to Euclidean distance. P_C or D_C is applied to each of the similar pairs according to an established probability p_{dc} (i.e., probability for Deterministic Crowding to be applied) and $1 - p_{dc}$ (i.e., probability for Probabilistic Crowding to be applied) respectively. Those individuals that are selected as winners by crowding replace the whole current population. The stop criterion was set to a specific number of generations. At the end of the process, the whole population should have converged to different good offers that are to be used in the negotiation process as an approximation to the real set of good deals for the agent. This population, called P , is used as an input for the negotiation process. A more detailed outline of the proposed GA can be observed in Algorithm 2.

3.3.2.2 Negotiation: Concession strategy

A concession strategy determines the aspirations of the agent at each negotiation time instant. The agent usually proposes offers that have a utility equal or above the utility level defined by its current aspirations. In this work, we assume a time-dependent tactic, where the utility demanded by each agent depends on the remaining negotiation time. This kind of concession strategies are adequate for environments such as AmI, where time is a very important limitation (e.g., limited power devices, dynamic environments, real-time environments, etc.). Some examples of concession tactics linear (same concession rate at each step), bouldware (18) (no concession until the last rounds, where it quickly concedes to the reservation value), and conceder (18) (at the start, it quickly concedes to the reservation

3. BILATERAL NEGOTIATION FOR LIMITED DEVICES

P : Explored preferences, good quality offers;
 D_c : Deterministic crowding rule;
 P_c : Probabilistic crowding rule;
 p_{cr} : Probability of crossover operator;
 p_{dc} : Probability of DC;
 n : Current number of generations;
 n_{max} : Maximum number of generations;
 $pair_i$: Pair of solutions;
 $P = \text{initialize}()$;
 $n = 0$;
while $n < n_{max}$ **do**
 shuffle P ;
 $P_{aux} = \emptyset$;
 $i = 1$;
 while $i \leq |P| - 1$ **do**
 $O_1 = P_i$;
 $O_2 = P_{i+1}$;
 if $\text{Random}() \leq p_{cr}$ **then**
 $(X_1, X_2) = \text{crossover}(O_1, O_2)$;
 else
 $X_1 = \text{mutate}(O_1)$;
 $X_2 = \text{mutate}(O_2)$;
 end
 $(pair_1, pair_2) = \underset{O_i \neq O_j \wedge X_k \neq X_l}{\text{argmin}} \ ||O_i - X_k|| + ||O_j - X_l||$;
 if $\text{Random}() \leq p_{dc}$ **then**
 Add($P_{aux}, D_c(pair_1)$);
 Add($P_{aux}, D_c(pair_2)$);
 else
 Add($P_{aux}, P_c(pair_1)$);
 Add($P_{aux}, P_c(pair_2)$);
 end
 $i = i + 2$;
 end
 $P = P_{aux}$; $n = n + 1$;
end
Return P ;

Algorithm 2: Pre-negotiation: Genetic algorithm with niching mechanism. Its goal is to sample the agent utility function

value).

One of the traits of similarity heuristics is that they are usually independent of the underlying concession strategy. In our motivating application, it is very important for both parties to appeal the other part. It is a more collaborative relationship that does not give room for such competitive strategies. In bouldware and conceder strategies, agents may invest too many rounds exchanging high/low utility offers that are not good for one of both parties. Since in our application it is very important to appeal the other part while maintaining a good utility, we assume a more exploratory concession strategy like linear concession.

In each negotiation round, the agents concede according to their strategy until a private deadline is reached. The utility that an agent a demands for a negotiation round t (i.e., concession strategy) can be formalized as follows (57):

$$s_a(t) = 1 - (1 - RU_a) \left(\frac{t}{T_a} \right)^{\frac{1}{\beta_a}} \pm \delta \quad (3.3)$$

where $s_a(t)$ is the concession strategy itself, which defines the demanded utility level for agent a at negotiation round t . RU_a is the reservation utility, and T_a is the private deadline of the agent, and β_a represents the concession speed of the agent. Since a linear concession speed is assumed, $\beta_a = 1$. δ is a small correction factor that allows demands of the agent to be more flexible in a negotiation round.

3.3.2.3 Negotiation: Acceptance criteria

An opponent offer is accepted if it provides a utility that is equal or greater than the demanded utility for the next negotiation round. Consequently, given the set of offers $X_{b \rightarrow a}^t = \{X_{b \rightarrow a}^{t,1}, X_{b \rightarrow a}^{t,2}, \dots, X_{b \rightarrow a}^{t,k}\}$ received by agent a from agent b at instant t , the acceptance criteria for agent a can be formalized as depicted in the following expression:

3. BILATERAL NEGOTIATION FOR LIMITED DEVICES

$$acc_a(X_{b \rightarrow a}^{t,best}) = \begin{cases} \text{accept} & U_a(X_{b \rightarrow a}^{t,best}) \geq s_a(t+1) \\ \text{reject} & \text{otherwise} \end{cases} \quad (3.4)$$

where $X_{b \rightarrow a}^{t,best} = \underset{X \in X_{b \rightarrow a}^t}{argmax} (U_a(X))$ is the best offer received in the last negotiation round in terms of one's own utility function, $U_a(\cdot)$ is the utility function of the agent, and $s_a(t+1)$ is the utility demanded for the next negotiation round.

3.3.2.4 Negotiation: Evolutionary sampling

One of the keys of the proposed strategy is the *evolutionary sampling*. This provides an implicit mechanism for learning opponent preferences and making an intelligent sampling of the negotiation domain. Basically, it is based in the application of genetic operators to offers received from the opponent in the last negotiation round and one's own good offers from P . The idea behind the *evolutionary sampling* is that offers generated by this method have genetic material from the opponent and one's own agent. Therefore, these offers may yield a greater probability of being accepted by the opponent that offers that have been sampled in the pre-negotiation without considering the opponent's preferences. The new offers sampled in this mechanism are added to a special population called P_{evo} .

Let us consider $X_{b \rightarrow a}^t = \{X_{b \rightarrow a}^{t,1}, X_{b \rightarrow a}^{t,2}, \dots, X_{b \rightarrow a}^{t,k}\}$, which is the set of offers sent by agent b to agent a at negotiation round t , and $s_a(t)$ the current demands of agent a . For each offer received from the opponent $X_{b \rightarrow a}^{t,i}$, a total of M offers are selected from the current iso-utility curve (i.e., offers with a utility in $s_a(t)$) defined in the population P . These M offers minimize the expression:

$$C = \underset{\{X_1, X_2, \dots, X_M \in P | U_a(X_j) \in s_a(t)\}}{\operatorname{argmin}} \sum_{j=1}^M \|X_{b \rightarrow a}^{t,i} - X_j\| \quad (3.5)$$

where C is the set of M different offers, and $\|X_{b \rightarrow a}^{t,i} - X_j\|$ is the Euclidean distance between one of the offers in C and the offer received from the opponent. Thus, these M offers are the ones most similar to $X_{b \rightarrow a}^{t,i}$ from the current iso-utility curve in P and they will be involved in the evolutionary process. The M selected offers are the most similar since applying crossover operators over offers that are too different may disrupt the quality of the solution for both agents (the resulting offer is too far from both agents' offers).

Once the M closest offers have been selected, a total of n_{cross} crossover operations are performed for each pair $(X_{b \rightarrow a}^{t,i}, X_j)$, where $X_j \in C$. The crossover operator takes two parents and generates one child. More specifically, the number of issues that come from $X_{b \rightarrow a}^{t,i}$ is chosen randomly from 1 and $N - 1$, with N being the number of issues. The rest of the issues come from X_j . Which particular issues come from each parent is also decided randomly. This way, each agent's preferences are taken into account in a statistically equal manner. Each child is added to a special pool, called P_{evo} , that contains new offers sampled during the different *evolutionary sampling* phases. An example of a crossover operation can be observed in Figure 3.3.

A total of n_{mut} mutation operations are carried out for each generated child by crossover operations. The mutation operator changes issue values randomly, according to a certain probability of mutating individual issues (p_{attr}). When p_{attr} is low, mutated offers are close to the original offer, so the effect is the exploration of the neighborhood of the offer. The operator is applied n_{mut} times to each child that is produced by crossover operations and to the original offers from the opponent. Mutation also generates new children that are added to the special pool P_{evo} .

Note that no offer from P_{evo} is discarded even though their utility may be considered too low for the current negotiation round. The reason for this mechanism

3. BILATERAL NEGOTIATION FOR LIMITED DEVICES

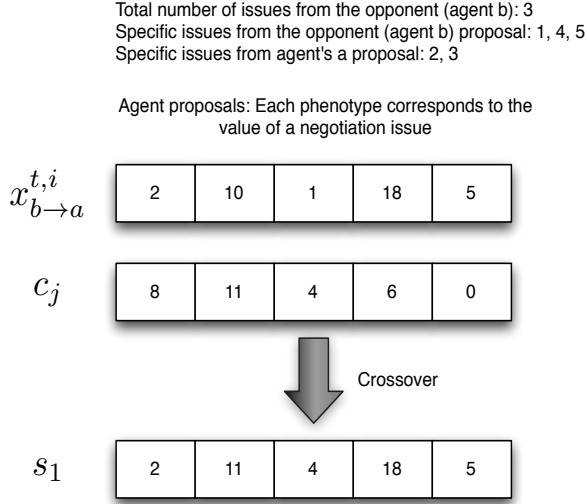


Figure 3.3: An example of a crossover operation.

is that offers that are not currently acceptable may be interesting in future rounds. Furthermore, since they have genetic material from the opponent's offers, they are more likely to be accepted.

As can be observed in Algorithm 3, if the negotiation process lasts t rounds, the Evolutionary Sampling explores a total number of offers that is equal to:

$$\begin{aligned}
 Samples_{evo} &= t \times ((k \times M \times n_{cross}) + (k \times M \times n_{cross}) \times n_{mut} + k \times n_{mut}) \\
 &= t \times k \times (M \times n_{cross} \times (1 + n_{mut}) + n_{mut})
 \end{aligned}$$

Then, the number of offers sampled during the negotiation process depends on the number of rounds that the negotiation lasts, k , M , and the number of genetic operators that are performed per offer selected from the iso-utility curve.

3.3.2.5 Negotiation: Select which offers to send

The next step in specifying the negotiation strategy consists of defining the mechanism to propose new offers. In this case, it is necessary to devise a mechanism that is capable of proposing up to k different offers to the opponent taking into account its preferences. The proposed heuristic takes into account the k offers received from the opponent and the offers in P and P_{evo} .

In our proposal strategy, k offers from the current iso-utility curve are sent. More specifically, two different iso-utility curves are calculated. The first one is the iso-utility curve calculated using offers in P . The second one is the iso-utility curve calculated using offers in P_{evo} . Basically, the first iso-utility curve has offers that were generated during the *self-sampling* (only taking into account one's own preferences), whereas the second iso-utility curve only has offers that were generated in the *evolutionary sampling* (they may take into account both agents' preferences). The negotiation strategy defines a proportion of p_{pevo} offers to come from the iso-utility curve in P_{evo} . The rest of the offers come from the iso-utility curve in P .

The $k_1 = \lceil p_{pevo} \times k \rceil$ offers selected from the iso-utility curve in P_{evo} are those that minimize the distance to any offer received from the opponent in the previous negotiation round. This selection may be formalized as:

$$\operatorname{argmin}_{\{X_1, X_2, \dots, X_{k_1} \in P_{evo} | U_a(X_j) \in s_a(t)\}} \left(\sum_{j=1}^{k_1} \min_{X \in X_{b \rightarrow a}^t} \|X_j - X\| \right) \quad (3.6)$$

On the other hand, offers are also selected from the iso-utility curve in P . The total number of offers corresponds to $k_2 = k - k_1$. In this case, offers that are the closest to any offer received from the opponent in the previous negotiation round

3. BILATERAL NEGOTIATION FOR LIMITED DEVICES

are selected. This selection can be formalized as:

$$\operatorname{argmin}_{\{X_1, X_2, \dots, X_{k_2} \in P | U_a(X_j) \in s_a(t)\}} \left(\sum_{j=1}^{k_2} \min_{X \in X_{b \rightarrow a}^t} \|X_j - X\| \right) \quad (3.7)$$

The parameter p_{pevo} determines the degree of relevance of the new offers sampled during the *evolutionary sampling* with respect to the offers sampled before the negotiation process. When $p_{pevo} = 0$, the strategy ignores the results that come from P_{pevo} . Consequently, only offers that were sampled in the pre-negotiation phase (*self-sampling*) are sent to the opponent. In this particular case, the strategy is equivalent to a negotiation strategy that only samples before the negotiation process and does not take into account the opponent's preferences. In contrast, when $p_{pevo} = 1$, the offers sampled during the *evolutionary sampling* are the only ones taken into account. In any case, p_{pevo} is a parameter to be adjusted.

3.3.2.6 Negotiation Trace

We prepared a very simple case based on the product fair example. To be more specific, it depicts a purchase in a furniture fair where one buyer is interested in buying chairs and tables from a seller. It should be pointed out that the goal of this case study is not to test the performance of the model, which will be thoroughly studied in Section 3.4, but to show a trace of the negotiation model from the point of view of one of the agents. In this case, we will focus on the buyer.

We use the weighted constraint model proposed by Ito et al. (78) to represent the utility functions of the buyer and the seller. The weighted constraint model is introduced as a complex utility function to model agent preferences. Let us consider a negotiation model where the number of issues is N , x_i represents the value of the i -th issue, each issue has an integer domain, and $X = \{x_1, x_2, \dots, x_N\}$

P : Offers from *self-sampling* ;
 P_{evo} : Offers from *evolutionary sampling* ;
 k : Number of offers of the protocol;
 M : Number of selected offers for genetic operations;
 n_{cross} : Number of times to crossover;
 n_{mut} : Number of times to mutate;
 p_{pevo} : Proportion of offers from P_{evo} ;
Receive ($X_{b \rightarrow a}^t \leftarrow b$);
 $X_{b \rightarrow a}^{t,best} = \underset{X \in X_{b \rightarrow a}^t}{argmax} (U_a(X))$;
if $acc_a(X_{b \rightarrow a}^{t,best}) = accept$ **then** Send (Accept $X_{b \rightarrow a}^{t,best} \rightarrow b$) ;
 $t = t + 1$;
/*Evolutionary sampling*/;
foreach $X_{b \rightarrow a}^{t,i} \in X_{b \rightarrow a}^t$ **do**
 $C = \underset{\{X_1, X_2, \dots, X_M \in P | U_a(X_j) \in s_a(t)\}}{argmin} \sum_{j=1}^M \|X_{b \rightarrow a}^{t,i} - X_j\|$;
 foreach $X_j \in C$ **do**
 repeat
 $X1 = Crossover(X_{b \rightarrow a}^{t,i}, X_j)$;
 if $X1 \notin P_{evo}$ **then** $Add(P_{evo}, X1)$ **repeat**
 $X2 = Mutate(X1)$;
 if $X2 \notin P_{evo}$ **then** $Add(P_{evo}, X2)$
 until n_{mut} times;
 until n_{cross} times;
 end
 repeat
 $X1 = Mutate(X_{b \rightarrow a}^{t,i})$;
 if $X1 \notin P_{evo}$ **then** $Add(P_{evo}, X1)$
 until n_{mut} times;
end
/*Select which offers to send*/;
 $k_1 = \lceil p_{pevo} * k \rceil$; $k_2 = k - k_1$;
 $X_{a \rightarrow b} = \underset{\{X_1, X_2, \dots, X_{k_1} \in P_{evo} | U_a(X_j) \in s_a(t)\}}{argmin} \left(\sum_{j=1}^{k_1} \min_{X \in X_{b \rightarrow a}^t} \|X_j - X\| \right)$;
 $X_{a \rightarrow b} = X_{a \rightarrow b} \cup \underset{\{X_1, X_2, \dots, X_{k_2} \in P | U_a(X_j) \in s_a(t)\}}{argmin} \left(\sum_{j=1}^{k_2} \min_{X \in X_{b \rightarrow a}^t} \|X_j - X\| \right)$;
Send ($X_{a \rightarrow b}^{t+1} \rightarrow b$);

Algorithm 3: Negotiation strategy during the negotiation process

3. BILATERAL NEGOTIATION FOR LIMITED DEVICES

represents a particular offer. These settings make up an N -dimensional space for the utility function.

In the weighted constraint model, a constraint c_l represents a specific region of the negotiation space. Any point of the space enclosed in that region is said to satisfy the constraint c_l . Basically, the term *constraint* represents an interdependence relationship among the negotiation issues. Each constraint c_l has a certain value $v(c_l, X)$ that is added to the utility of X when the constraint is satisfied by the offer X . For instance, a constraint defined as $c_l = (1 \leq x_1 \leq 10 \wedge 3 \leq x_2 \leq 4)$ and $v(c_l, X) = 10$ would hold a utility of 10 for the offer (2,3).

A utility function in the weighted constraint model is formed by l constraints whose values are summed up whenever the constraints are satisfied. The utility of a point X given l constraints can be defined as:

$$U(X) = \sum_{c_l \in L} v(c_l, X) \quad (3.8)$$

where X is the offer, c_l is a constraint, L is the set of constraints, and $v(c_l, X)$ is the value of the constraint if it is satisfied (0 otherwise).

As stated in (78), although the expression seems linear, it produces a non-linear utility space due to the interdependence among the issues. Furthermore, the utility function may generate spaces with several local maxima, which makes the problem highly non-linear and very difficult to optimize. Additionally, the agents do not have any knowledge about the possible constraints of the opponent, thus making the problem of negotiation still more difficult.

This example consists of 3 different issues: price (P) [0-9] \times 100 euros, chair color in a chromatic scale (CC) [0-9], and table color in a chromatic scale (TC) [0-9]. Next, we introduce the utility functions employed to represent the preferences of both consumer and seller:

3.3 Negotiation Model

Buyer Utility Function	Seller Utility Function
$(v_1 = 100) (0 \leq P \leq 1)$	$(v_1 = 80) (8 \leq P \leq 9)$
$(v_2 = 50) (2 \leq P \leq 4)$	$(v_2 = 60) (6 \leq P \leq 7)$
$(v_3 = 25) (5 \leq P \leq 7)$	$(v_3 = 45) (4 \leq P \leq 5)$
$(v_4 = 30) (0 \leq CC \leq 3) \wedge (0 \leq TC \leq 3)$	$(v_4 = 20) (1 \leq P \leq 3)$
$(v_5 = 10) (0 \leq CC \leq 3) \wedge (6 \leq TC \leq 9)$	$(v_5 = 15) (1 \leq CC \leq 2)$
$(v_6 = 50) (0 \leq CC \leq 3) \wedge (5 \leq TC \leq 6)$	$(v_6 = 10) (0 \leq CC \leq 1)$
$(v_7 = 30) (4 \leq CC \leq 6) \wedge (0 \leq TC \leq 3)$	$(v_7 = 10) (2 \leq CC \leq 5)$
$(v_8 = 20) (4 \leq CC \leq 5) \wedge (4 \leq TC \leq 5)$	$(v_8 = 5) (5 \leq CC \leq 9)$
$(v_9 = 10) (4 \leq CC \leq 5) \wedge (8 \leq TC \leq 9)$	$(v_9 = 20) (8 \leq CC \leq 9)$
$(v_{10} = 50) (7 \leq CC \leq 9) \wedge (2 \leq TC \leq 4)$	$(v_{10} = 60) (0 \leq TC \leq 1)$
$(v_{11} = 20) (7 \leq CC \leq 9) \wedge (6 \leq TC \leq 8)$	$(v_{11} = 30) (1 \leq TC \leq 4)$
	$(v_{12} = 5) (4 \leq TC \leq 6)$
	$(v_{13} = 20) (6 \leq TC \leq 9)$
	$(v_{14} = 10) (8 \leq TC \leq 9)$

The consumer shows issue interdependences relating the two types of furniture (e.g. some pairs of colors fit better than other pairs). In the case of the seller, no interdependences are found but he may present preferences regarding which models to sell (e.g. some of them need to be manufactured; some models only have a few units, etc.).

As for the parameters of the self-sampling phase, they were set to $|P| = 16$, $n_{max} = 100$, $p_{dc} = 80\%$ and $p_{cr} = 80\%$. The rest of parameters of the negotiation model were set to $\delta = 0.05$, $k = 2$, $T = 10$, $p_{pevo} = 100\%$, $n_{cross} = 2$, $n_{mut} = 2$, and $M = 2$.

The next table shows the 16 offers found by the self-sampling process carried out by the buyer. It depicts the value for each issue and the utility of the offer. In this case the utility has been scaled to $[0,1]$ for the sake of simplicity.

3. BILATERAL NEGOTIATION FOR LIMITED DEVICES

P=Self-sampling results for the buyer	
($u = 1.00$) 1 1 6	($u = 0.81$) 1 3 0
($u = 1.00$) 0 1 6	($u = 0.81$) 1 5 3
($u = 0.93$) 1 7 3	($u = 0.62$) 0 2 4
($u = 0.93$) 1 7 4	($u = 0.62$) 1 9 1
($u = 0.93$) 1 9 2	
($u = 0.93$) 1 2 5	
($u = 0.93$) 1 8 3	
($u = 0.93$) 1 8 4	
($u = 0.93$) 0 7 3	
($u = 0.93$) 1 9 3	
($u = 0.93$) 0 1 5	
($u = 0.81$) 1 5 0	

Round 1 $s_s(1) = [0.95, 1]$ $s_b(1) = [0.95, 1]$ Once the self-sampling phase has finished, the negotiation process starts with the buyer acting as initiator. Since there are no opponent offers to evaluate, evolutionary sampling is skipped and the agent directly proposes offers to the opponent. Due to the fact that no evolutionary sampling has been carried out, P_{evo} is empty and only the iso-utility curve that comes from P can be calculated. $X=(1\ 1\ 6)$ and $Y=(0\ 1\ 6)$ are selected since there is no opponent offer to compare with. The opponent rejects the offers since they yield a utility of 0.35 and 0.25 respectively. The opponent makes a counteroffer which contains $W=(8\ 1\ 1)$ and $Z=(9\ 1\ 1)$. Both of them are rejected since their utilities (0.18 for both of them) are lower than 0.85.

Round 2 $s_s(2) = [0.85-0.95]$ $s_b(2) = [0.85-0.95]$ Two offers have been received from the opponent. Thus, the evolutionary sampling phase is carried out. The iso-utility curve from P ($s_s(2) = [0.85 - 0.95]$) is shown in the following tables. It shows the offers and the Euclidean distance to W and Z. For both W and Z, the $M = 2$ offers which are more similar are selected. The offers selected from the iso-utility curve become one of the parents for genetic operations, which are also shown in the following tables. For the sake of simplicity, genetic operations which produced children that were already in P_{evo} are not included (nor are they stored

3.3 Negotiation Model

more than once). All of the offers generated during this phase are added to P_{evo} .

Iso-utility curve (P)			Genetic Operations				
Offer	d(W)	d(Z)	Crossover			Mutation	
			Parent 1	Parent 2	Child	Parent 1	Child
1 2 5	0.90	1.00	8 1 1	1 2 5	(u=0.31) 8 2 5	8 2 5	(u=0.34) 6 2 1
0 1 5	0.99	1.09	8 1 1	1 2 5	(u=0.81) 1 1 1	1 1 1	(u=0.68) 1 1 7
1 7 3	1.04	1.13	8 1 1	0 1 5	(u=0.81) 0 1 1	1 1 1	(u=0.18) 8 1 1
1 7 4	1.07	1.16	9 1 1	1 2 5	(u=0.31) 9 1 5	8 1 1	(u=0.31) 2 1 4
1 8 3	1.12	1.20	9 1 1	1 2 5	(u=0.31) 9 2 5	8 1 1	(u=0.15) 5 7 1
0 7 3	1.13	1.22	9 1 1	0 1 5	(u=0.31) 9 1 5	9 1 5	(u=0.46) 6 1 5
1 8 4	1.14	1.22	9 1 1	0 1 5	(u=0.81) 0 1 1	9 1 5	(u=0.62) 1 8 5
1 9 2	1.18	1.26				9 2 5	(u=0.81) 1 2 3
1 9 3	1.20	1.27				9 2 5	(u=0.15) 7 6 5
						9 1 1	(u=0.50) 4 0 1
						9 1 1	(u=0.37) 9 2 6

Next, it is necessary to select which offers to send to the opponent. Since $p_{pevo} = 100\%$, if possible, all of the offers will come from the iso-utility curve calculated using P_{evo} . If it is not possible, it will take as many offers as possible from the iso-utility curve from P_{evo} and take the rest from the iso-utility curve from P . In this case, $X=(1\ 2\ 5)$ and $Y=(0\ 1\ 5)$ are selected from P since P_{evo} does not contain elements to form a current iso-utility curve. The opponent receives the offers X and Y . Since they yield a utility of 0.25 and 0.15 respectively, both are rejected. The seller sends $W=(6\ 1\ 1)$ and $Z=(9\ 4\ 1)$ as counteroffers. Both of them are rejected since their utilities (0.34 and 0.18 respectively) are lower than 0.75.

Round 3 $s_s(2) = [0.75-0.85]$ $s_b(2) = [0.75-0.85]$ Two offers have been received from the opponent. Thus, the evolutionary sampling phase is carried out. The iso-utility curve from P ($s_s(2) = [0.75 - 0.85]$) and genetic operations are shown in the following tables.

3. BILATERAL NEGOTIATION FOR LIMITED DEVICES

Iso-utility curve (P)		
Offer	d(W)	d(Z)
1 3 0	0.60	0.90
1 5 0	0.72	0.90
1 5 3	0.74	0.92

Genetic Operations				
Crossover			Mutation	
Parent 1	Parent 2	Child	Parent 1	Child
6 1 1	1 3 0	(u=0.81) 1 3 1	1 3 1	(u=0.00) 8 8 1
6 1 1	1 3 0	(u=0.34) 6 1 0	1 3 1	(u=0.62) 1 6 7
6 1 1	1 5 0	(u=0.81) 1 1 0	6 1 1	(u=0.34) 6 2 1
9 4 1	1 3 0	(u=0.18) 9 4 0	6 1 1	(u=0.21) 6 1 8
9 4 1	1 5 0	(u=0.81) 1 5 1	1 1 0	(u=1.00) 0 1 6
			1 8 5	(u=0.62) 1 8 5
			9 4 0	(u=0.18) 9 5 0
			9 4 0	(u=0.81) 1 4 1
			9 4 1	(u=0.18) 8 4 1
			9 4 1	(u=0.18) 9 6 1
			1 5 1	(u=0.62) 1 7 0
			1 5 1	(u=0.31) 4 7 1

Next, it is necessary to select which offers to send to the opponent. The table below shows the iso-utility curve calculated from P_{evo} . In this case, $X=(1 1 1)$ and $Y=(1 1 0)$ are selected from P_{evo} . The opponent receives the offers X and Y. Since they yield a utility of 0.69 and 0.53 respectively, both are rejected. However, in this round, the seller sends $W=(4 1 1)$ as counteroffer. The offer is rejected because its utility is equal to 0.5, and is thus lower than 0.65. From this point on we will overlook the inner steps of the model due to the fact that the way it works has already been described.

Iso-utility curve (P_{evo})		
Offer	d(W)	d(Z)
1 1 0	0.56	0.95
0 1 1	0.66	1.05
1 4 1	0.64	0.88
1 3 1	0.59	0.89
1 5 1	0.71	0.89
1 1 1	0.55	0.94
1 2 3	0.60	0.94

Round 4 $s_s(2) = [0.65 - 0.75]$ $s_b(2) = [0.65 - 0.75]$ In this round, the buyer sends $X=(1 1 7)$, which yields a utility of 0.33 for the seller. Therefore, the offer is rejected. Then, the opponent sends $W=(1 1 1)$ and $Z=(1 2 1)$, Z being accepted by

the buyer since its utility is equal to 0.81. The negotiation process ends with the deal $(U_b = 0.81, U_s = 0.69)$, which is the Nash Bargaining Point for this negotiation case.

This section has described the main traits of the proposed negotiation model for AmI environments. More specifically, it has explained the protocol employed, and the negotiation strategy that is adapted to AmI domains thanks to the intelligent sampling provided by genetic operators during the negotiation process. Additionally, we have also shown how the proposed model works in a small case study. In the next section the proposed model is tested in several scenarios to check its performance.

3.4 Experiments

The performance of the devised strategy is studied in this section. The proposed negotiation model is tested in the weighted constraint model proposed by Ito et al. (78). This model makes it possible to represent unrestricted interdependence relationships among issues. Furthermore, if the number of constraints is large, it can represent highly non-linear utility functions. Therefore, it represents a proper testbed for the proposed strategy. Nevertheless, as in the work of Lai et al. (57), the proposed negotiation model is general and does not depend on a particular utility function. The model of Ito et al. was selected as a testbed because it provides a well studied utility function (78, 82, 83) that holds enough complexity to study the real performance of the negotiation model.

Firstly, the negotiation setting employed in the experiments is briefly described. After this, the different experiments and their results are presented.

3. BILATERAL NEGOTIATION FOR LIMITED DEVICES

3.4.1 Negotiation Setting

The aim of these experiments was to evaluate whether or not the proposed model is capable of working in domains where the agents' utility functions are highly non-linear. For that purpose, different negotiation cases were randomly created:

- Number of issues $N = [4-7]$.
- Integer issues. $x_i \in [0, 9]$.
- $L = N*5$ uniformly distributed constraints per agent. There are constraints for every possible interdependence cardinality. For instance if $N=4$, there are 5 unary constraints, 5 binary constraints, 5 ternary constraints and 5 quaternary constraints.
- $v(c_i, \cdot)$ for each n -ary constraint drawn randomly from $[0, 100 * n]$.
- For every constraint, the constraint width for each issue x_i is uniformly drawn from $[2, 4]$. For instance, if the constraint width for issue x_1 is 3, then $(0 \leq x_1 \leq 3)$, $(1 \leq x_1 \leq 4)$, $(2 \leq x_1 \leq 5)$, $(3 \leq x_1 \leq 6)$, $(4 \leq x_1 \leq 7)$, $(5 \leq x_1 \leq 8)$ and $(6 \leq x_1 \leq 9)$ are all of the possible configurations for issue x_1 in the constraint (just one is used in the constraint).
- Agent deadline was set to a maximum of 10 rounds. This represents a total of 20 messages exchanged between both agents (offers and counteroffers). Agents do not know their opponent's private deadlines.
- Agent reservation utility $RU = 0$. It is set to zero in order to find a deal, if possible. Should this be the case, the deal is checked against certain thresholds which will determine whether the application notifies its owner of the possible deal.

- Agents do not know their opponent's utility functions

For each number of issues, a total of 100 negotiation cases were generated with the above settings. The execution of each case was repeated 30 times in order to capture differences between executions of the negotiation model.

In order to evaluate the quality of final agreements, some measures were gathered at the end of each negotiation.

- Euclidean distance to the closest Pareto frontier point (51). This is a measure of economic efficiency for agreements. If an offer is not in the Pareto frontier, it means that one of the two parties can improve its utility without decreasing the utility of the other agent. The closest an agreement is to the Pareto frontier, the better.
- Euclidean distance to the Nash Product (51). Since the proposed model is collaborative in essence, it is worth to study the distance to the Nash product. This is the point that maximizes the product $U_1 * U_2$ in the Pareto Frontier, where U_1 is the utility of agent 1, and U_2 is the utility of agent 2.
- Number of negotiation rounds. Faster agreements are preferred since a lesser number of messages are exchanged, less bandwidth is needed, and limited devices need less power to send messages.

Additionally, some experiments were also devised in order to test the computational performance of the proposed model in a real environment. Measures such as the time spent in decision making tasks before the negotiation process (self-sampling) and during the negotiation process (opponent offer acceptance phase, evolutionary sampling, and offer proposal) were gathered. For that purpose, the proposed model was implemented using a HTC Desire (1 Ghz, 576MB RAM, Android Operating System) as one of the parties and a PC (2 Ghz, 4096MB RAM,

3. BILATERAL NEGOTIATION FOR LIMITED DEVICES

Ubuntu Operating System) as the other party. A total number of 30 negotiations were carried out in order to measure the computational cost of the proposed model.

In this work, we employ confidence intervals (95%, $\alpha = 0.05$) to study possible differences in the averages. If confidence intervals for both data samples do not overlap, we can claim that there are statistical and significant differences between both data samples.

3.4.2 Results

The proposed strategy, which will be named as *Evolutionary Sampling* or *ES*, was compared with two different negotiation models. The first strategy is an implementation of the general framework proposed by Lai et al. (57). This model is provided with the whole sampling of the utility function, so that it can completely calculate iso-utility curves. It is used as a measure of how close the proposed strategy is to the ideal case where all of the offers are available. The second model assumes that it is not possible to completely sample all of the offers. Therefore, it samples before the negotiation process by means of a niching GA (*self-sampling*) and uses the similarity heuristic ($p_{pevo} = 0$) during the negotiation process, which will be named as *Non Evolutionary Sampling* or *NES* model. The number of samples explored by the *NES* model before the negotiation process is set equal to the number of samples explored by the *ES* model ($|P| + Samples_{evo}$). Consequently, both the *NES* and *ES* model yield the same sampling cost in every experimentation.

Five different experiments were carried out in order to test the proposed model. In the first experiment, the three different models are compared as the number of issues is increased. The second experiment, studies the impact of the proportion of offers (p_{pevo}) that are sent from the special pool P_{evo} in the *ES* model. Next, the three models are compared as the number of proposals k increases. Finally,

the *ES* and the *NES* model are compared as the size of the population ($|P|$) provided by the *self-sampling* increases. Finally, we studied the time consumed by the proposed method in a realistic environment involving limited devices.

3.4.2.1 Experiment 1: Performance Study on the Number of Issues

The goal of this experiment is to study how the proposed strategy behaves for negotiations with a different number of issues $N = \{4, 5, 6, 7\}$. It is important for the proposed model to be capable of properly handling negotiations with multiple issues. A negotiation setting where agents are limited to $k = 3$ proposals per negotiation round is used. The three different models were tested during this experiment.

The parameters of the *self-sampling* were set to $n_{max} = 100$, $p_{dc} = 80\%$ and $p_{cr} = 80\%$. The number of samples optimized before the negotiation process was set to $|P| = 128$ for the *ES* model and to $|P| = 128 + Samples_{evo}$ for the *NES* model.

The parameters of the *ES* were set to $M = 5$, $n_{cross} = 4$, $n_{mut} = 4$, $p_{attr} = 30\%$, and $p_{pevo} = 100\%$. Therefore, all the offers are sent from the samples generated by the *evolutionary sampling* carried out during the negotiation process.

The distance to the Nash Product, the distance to the closer Pareto Frontier Point and the number of negotiation rounds were measured for the three models. The results for this experiment can be found in Figure 3.4. The figure shows the average and its associated confidence intervals (95%, $\alpha = 0.05$). Intuitively speaking, since the number of offers sampled remains constant and the number of issues increases, the performance of the *NES* and the *ES* model should be worsened with respect to the results achieved by the model of Lai et al. However, the results for the *ES* do not comply with this intuitive hypothesis. As can be observed, even though the proposed model and the *NES* model explore the same

3. BILATERAL NEGOTIATION FOR LIMITED DEVICES

number of offers, the *NES* obtains worse results than the other two models. This is particularly true as the number of issues increases, since the performance of this method drastically decreases. On the contrary, the *ES* model is capable of achieving statistically equal results to the model of Lai et al., which can access the whole iso-utility curve. Nevertheless, the proposed model explores far fewer offers than the complete sampling of the utility function, especially for larger number of issues. For instance, when $N = 6$, Lai et al. has access to 10^6 offers, whereas the proposed model has only sampled an average of 1510 samples (128+ average *Samples_{evo}*). Only when the number of issues is equal to 7, there are significant differences between Lai et al. and *ES*, which highlights the obvious fact that, as the negotiation domain gets larger, more sampling is necessary.

The *ES* model has been able to achieve similar results to the case where the full iso-utility curve can be calculated, while maintaining the offers sampled to a small number. This result is particularly interesting for AmI domains where agents may be executed in devices with low computational and storage capabilities. Therefore, fewer samples mean less power consumption and less capacity needed to store them. Moreover, it must also be highlighted that the number of rounds was also lower than that obtained by *NES*, which, consequently means fewer messages sent, less bandwidth needed and, of course, less power consumption by the limited devices.

The reason for this improvement is the intelligent sampling achieved by the use of genetic operators during the negotiation process. On the contrary, sampling only before the negotiation process leads to worse results since it is not capable detecting which offers will be interesting for the negotiation. Both, the *ES* and the *NES* model, have the same computational cost, but the *ES* is obviously preferred since it is capable of achieving a better performance un all aspects.

3.4.2.2 Experiment 2: Performance Study on p_{pevo}

In this case, the experiment's goal is to study how relevant the proportion of offers that are sent from the offers sampled during the negotiation process (governed by the parameter p_{pevo}) in the *ES* model is. Since all of the configurations sample new offers during the negotiation process, all of them yield a very similar computational cost. In fact, it may only be different if one of the configurations obtains a significantly different number of negotiation rounds. Consequently, the main subject of study in this scenario is the economic efficiency (distance to Nash and Pareto Frontier), although some improvements in the computational cost may be observed due to a lower number of rounds.

The same conditions from the previous experiment were set ($k = 3$ and $N = \{4, 5, 6, 7\}$), and the same configuration parameters were set for the *ES* ($M = 5$, $n_{cross} = 4$, $n_{mut} = 4$, and $p_{attr} = 30\%$). However, in this scenario we compare the *ES* model results when 1 out of 3 offers ($p_{pevo} = 30\%$), 2 out of 3 offers ($p_{pevo} = 50\%$), and 3 out of 3 offers ($p_{pevo} = 100\%$) come from the offers sampled during the *evolutionary sampling* phase.

The results for this second scenario can be observed in Figure 3.5. The graphic shows the averages and their associated confidence intervals (95%, $\alpha = 0.05$). It can be observed that the three different configurations yield similar results for the distance to the Nash Product, the distance to the closest Pareto Frontier Point, and the number of negotiation rounds. This similarity is explained due to the fact that, on most occasions, the offer accepted by the opponent is the closest one from the *evolutionary sampling* population (P_{evo}). Therefore, it is always sent, as long as the results from the *evolutionary sampling* are not ignored. Nevertheless, it seems that higher values of p_{pevo} have a slightly (and significantly) better economic and computational performance than lower ones. The reason for

3. BILATERAL NEGOTIATION FOR LIMITED DEVICES

this slight improvement is that, in some cases, the offer preferred by the opponent may be the second or third closest from P_{evo} . Due to this small improvement, higher values of p_{pevo} are preferred in practice.

3.4.2.3 Experiment 3: Performance Study on k

The next experiment aims to study the performance of the three different models (Lai et al., *NES*, and *ES*) as k is increased. The number of offers sent may help to reach agreements faster since more negotiation space is explored. This is very important in AmI environments where devices have limited power and their running time must be optimized. Lai et al. (57), demonstrated how higher values of k helped to reach better agreements. In this scenario, the experiment is repeated in order to evaluate whether the differences between the three models still hold for different values of k .

The studied values of k were 1, 3, 5, and 7. The rest of the negotiation setting was configured to use negotiation cases with $N = 6$ issues. The parameters of the *self-sampling* were set to the values employed in the previous tests except for $|P| = 256$. The parameters of the *ES* were set to the same conditions described in Experiment 1.

The results for this experiment are shown in Figure 3.6. The graphic shows the averages and their associated confidence intervals (95%, $\alpha = 0.05$). As it can be observed, the three models achieve better results as k increases. These results agree with those presented in (57). Although all of the models improve, the differences observed in Experiment 1 still hold for this scenario. The *NES* model gets worse results than Lai et al. and the proposed model. On the contrary, the *ES* obtains results that are statistically equivalent to the case when the full iso-utility curve can be calculated for small values of k . For higher values of k the proposed model gets slightly better results than Lai et al.

It must be noted again that the number of offers sampled for *ES* and *NES* is the same and it is much lower than the complete sampling of the utility function. For instance, in this scenario, the complete sampling consists of 10^6 offers, whereas the other two methods sampled an average of 773 samples for $k = 1$, 1653 for $k = 3$, 2497 for $k = 5$, and 3357 for $k = 7$.

3.4.2.4 Experiment 4: Performance Study on $|P|$ and Memory Performance

This last experiment was designed to assess the influence of the population optimized by the *self-sampling* on the performance of the *ES* model and the *NES* model. It is especially relevant to see how many samples the *NES* model needs to achieve similar results to those ones obtained by *ES*. Obviously, more population means more storage needed and more computational cost since it needs to optimize more samples.

The average number of samples explored was analyzed for a negotiation setting where $N = 6$ and $k = 3$. The settings used for the *self-sampling* and the *ES* in previous experiments were repeated for this scenario. The number of sampled offers was increased by allowing more offers to be optimized in the *self-sampling* ($|P| = \{128, 256, 512, 1024, 2048, 4096\}$).

The results for this experiment can be observed in Figure 3.7. The figure contains the averages and their associated confidence intervals (95%, $\alpha = 0.05$). The x axis of the graphics show the average number of offers sampled by both models, thus it shows $|P| + rounds * Samples_{evo}$. In the case of the *NES* model all of the samples were produced before the negotiation process started. Several observations can be made from the data shown in the graphics. On the one hand, it seems that the size of $|P|$ does not have too much of an effect on the performance of the *ES* model, since it is more dependent on the exploration carried out during

3. BILATERAL NEGOTIATION FOR LIMITED DEVICES

the negotiation process and does not need as much sampling to get results similar to the case where the full iso-utility curve can be accessed. Therefore, the behavior of the model remained almost constant for different configurations of $|P|$. Again, this behavior is very adequate for AmI environments since the model can properly work with configurations that do not require too many computational resources. On the other hand, the *NES* model performance increased along with the number of offers sampled. It must be noted, that when the number of samples for both methods was 5506, the two of them obtained very similar, almost equivalent, results. Therefore, the *NES* needed 5506 samples to achieve similar results to the same results obtained by the *ES* model for 1510 samples. It can be concluded that *NES* needs $\frac{5506}{1510} = 3.64$ times more samples to achieve similar results to *ES*.

It is possible to approximately analyze the total amount of memory employed by both methods when they achieve statistically equivalent results. As has been suggested by the previous experiment (Experiment 4), the *NES* model needs 5506 samples to achieve statistically equivalent results to those the *ES* model with 1510 samples. If we assume that the underlying platform is a 32 bit platform, where integers usually need 32 bits to be stored, we can approximately calculate the memory needed by both models as follows:

$$Memory(KB) = |Samples| * N * 32 * \frac{1}{8} * \frac{1}{1024} \quad (3.9)$$

where $|Samples|$ is the number of samples, N is the number of issues of the negotiation process, 32 is the size of an integer, $\frac{1}{8}$ converts from bits to Bytes, and $\frac{1}{1024}$ converts from bytes to KBytes. Taking into account the formula above, the *NES* model would take 129 KB to store the data needed for the previous type of negotiation process ($N = 6$, $|Samples| = 5506$), whereas the *ES* model would take 35 KB ($N = 6$, $|Samples| = 1510$). Depending on the underlying device, this difference may be important (e.g. devices with a few MB of storage available).

3.4 Experiments

Number negotiations	Memory(KB) <i>NES</i>	Memory(KB) <i>ES</i>
1	129	35
3	387	105
5	645	175
7	903	245
10	1290	350

Table 3.1: Approximate amount of memory needed by the *NES* and *ES* model.

However, this difference may be still more important if we consider that in some scenarios it may be necessary to perform several negotiations at the same time (e.g. the fair scenario). For instance, Table 3.1 shows the approximate amount of memory necessary (in Kilobytes) for *NES* and *ES* to carry out several negotiations at the same time. As the number of negotiation issues is larger, the amount of space needed to store offers is bigger. Thus, storing a lesser number of offers is preferred to larger numbers.

3.4.2.5 Experiment 5: Time Performance

As introduced earlier, it was also interesting to test the computational performance of the model in a real environment. Thus, the proposed model was implemented using a HTC (1 Ghz, 576MB RAM, Android Operating System) as one of the parties and a PC (2 Ghz, 4096MB RAM, Ubuntu Operating System) as the other party. The *self-sampling* parameters were set to $n_{max} = 100$, $p_{dc} = 80\%$ and $p_{cr} = 80\%$. The number of samples optimized before the negotiation process was set to $|P| = 128$. As for the parameters employed during the negotiation process, these were set to $k = 3$, $M = 5$, $n_{cross} = 4$, $n_{mut} = 4$, $p_{attr} = 30\%$, and $p_{pevo} = 100\%$. The number of issues of the negotiation process was $N = 5$. The time spent in the whole negotiation process (t_t), the time spent in sending/waiting for offers (t_m), the time spent in *self-sampling* (t_s), and the time spent in decision-

3. BILATERAL NEGOTIATION FOR LIMITED DEVICES

t_t (s)	t_s (s)	t_{dm} (s)	t_m (s)
0.773	0.264	0.358	0.415

Table 3.2: Average time performance of the ES model for 30 negotiations.

making during the negotiation process (t_{dm}) were measured. Table 3.2 shows the average negotiation time in seconds for the 30 negotiation cases that were studied.

As can be observed, the time spent for a negotiation process t_t was reasonably good (less than a second) and it enables negotiations to be carried out in environments where real-time responses are needed (e.g. Ambient Intelligence). Moreover, it can also be observed that the time spent in decision-making tasks t_{dm} does not take as much time as other tasks such as sending/waiting for offers t_m . This leaves room for more negotiation processes to be carried out in parallel during CPU idle time (e.g. waiting for offers). Again, carrying out multiple negotiation processes simultaneously proves especially interesting again for AmI environments. For instance, in the fair scenario, it makes it possible to negotiate simultaneously with those vendors who are available in the area where the consumer is walking at that moment. The time spent in self-sampling is the least problematic since it is a process to be carried out only once until agent preferences change. In some AmI environments, such as the fair, we may consider preferences to be static during the fair event. Thus, *self-sampling* would only be needed once. Despite all those facts, it must be remarked that the time spent in *self-sampling* is reasonably good.

3.5 Conclusions

A multi-issue bilateral bargaining model for Ambient Intelligence domains that deals with complex interdependent utility functions has been presented in this chapter. This work complements the inspiring work of Lai et al. (57) and provides

a negotiation model that is adequate for Ambient Intelligence applications. The main contribution of this chapter has been achieving efficient agreements while maintaining the use of computational resources low.

The proposed model uses a negotiation protocol where agents are allowed to send up to k different offers in each negotiation round. Before the negotiation process starts, each agent samples its own utility function by means of a niching genetic algorithm. This genetic algorithm gets highly interesting and significantly different offers for one's own utility function (*self-sampling*). After the negotiation process starts, the agents apply genetic operators over the last offers received from the opponent and those offers that are most similar from the current iso-utility curve (*evolutionary sampling*). The desired effect is to sample new offers that are interesting for both parties. Therefore, the opponent's preferences guide the sampling process during the negotiation process. The offers that are sent to the opponent are selected from the current iso-utility curve, being those that are the most similar to the last offers received from the opponent. An additional mechanism is introduced that allowing priority to be given to those offers that come from the *evolutionary sampling* iso-utility curve. The results obtained by the proposed model, while maintaining good economic performance, cope with the problems found in AmI environments. The results of the experiments can be summarized as:

- The proposed model needs very few computational resources and storage to obtain results statistically equivalent to the ideal case where the all of the offers are available (57). It obtained similar results in economic performance (distance to Nash, distance to Pareto Frontier) and number of negotiation rounds.

3. BILATERAL NEGOTIATION FOR LIMITED DEVICES

- When the proposed model and the *NES* model sample the same number of offers, the first obtains better results. In fact, the *NES* model needs to sample 3.64 times more offers to obtain similar results.
- The proposed model needs less negotiation rounds to achieve better results than the *NES* model. Therefore, the environment bandwidth is optimized since it needs fewer messages to be sent in order to reach agreements.
- We have also been able to appreciate that, in an environment involving limited devices, negotiations are executed in a very reasonable time (i.e., less than a second for a maximum deadline of 10 rounds, 20 messages exchanged). This is also a very important factor for AmI devices and the motivating application.

Consequently, the proposed model fits perfectly for the conditions needed by AmI environments, since it needs less computational resources and it obtains economically efficient results.

In this chapter we were able to cover our goals with regards to negotiation in Ambient Intelligence: a negotiation model that is capable of reaching good agreements while being computationally efficient. Hence, for this point on, we drove our main research efforts towards a novel topic like agent-based negotiation teams. This effort is reflected in the following chapters, where we propose a general workflow of tasks for agent-based negotiation teams and we propose several negotiation models for a wide variety of negotiation scenarios (i.e., scenarios exclusively composed by *predictable* and *compatible* issues, and scenarios that also include *unpredictable* issues).

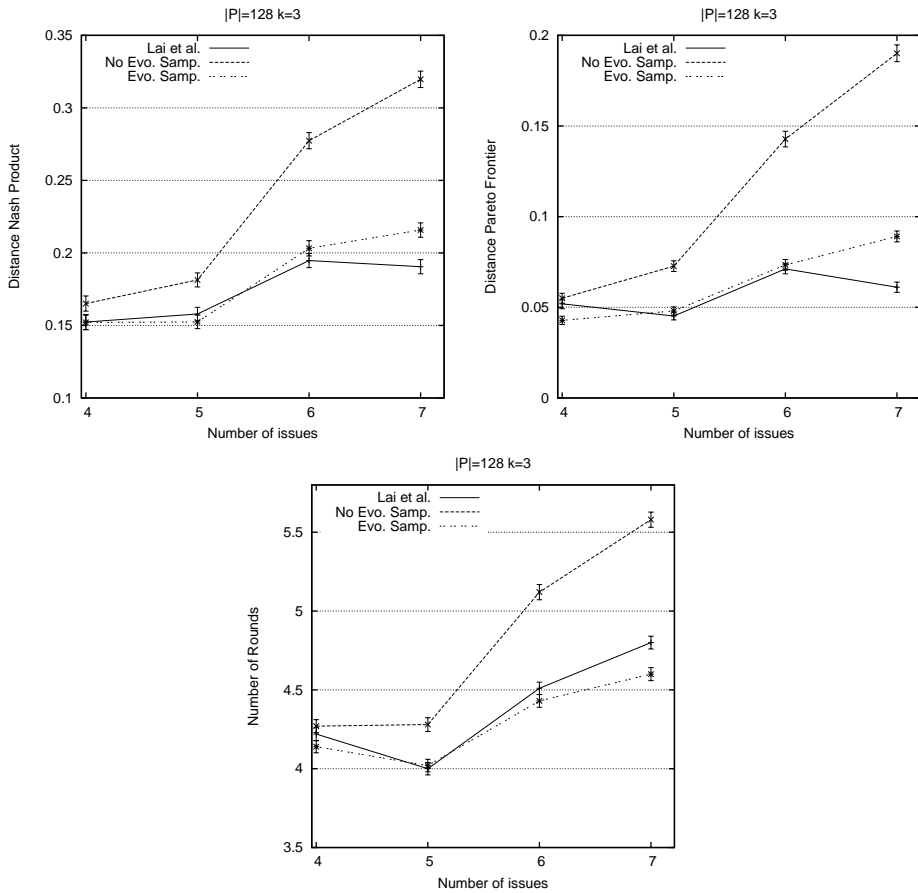


Figure 3.4: Distance to the Nash Product, distance to the closest Pareto Point, and number of negotiation in Experiment 1.

3. BILATERAL NEGOTIATION FOR LIMITED DEVICES

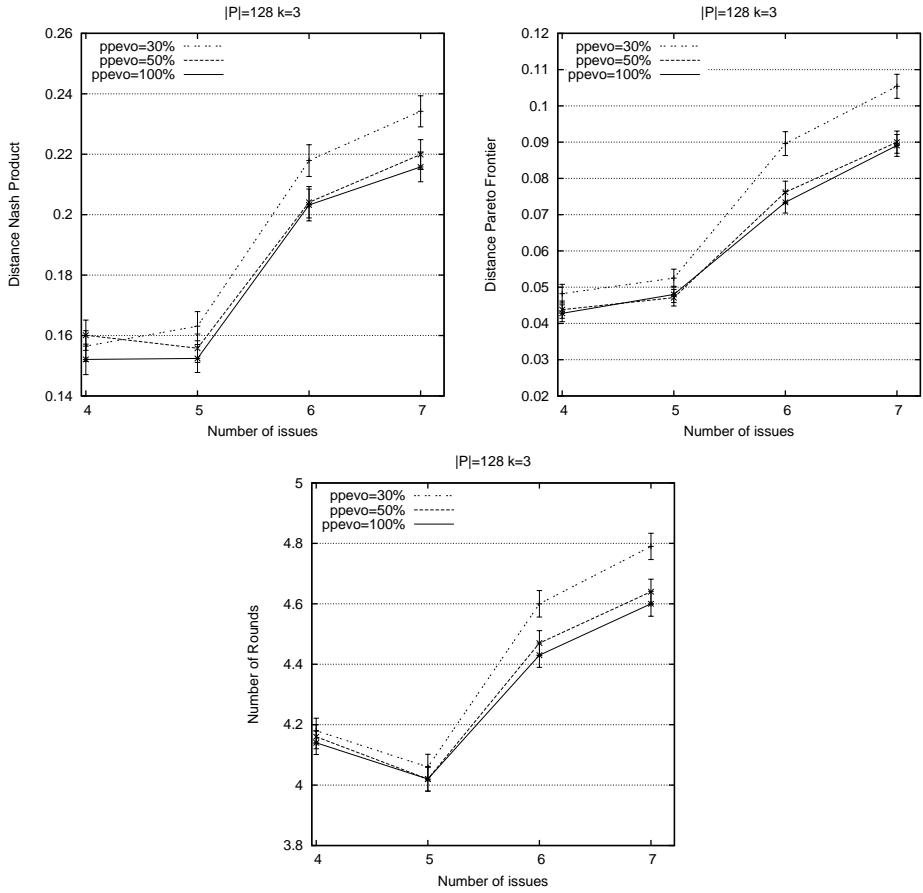


Figure 3.5: Distance to the Nash Product, distance to the closer Pareto Point, and number of negotiation rounds in Experiment 2.

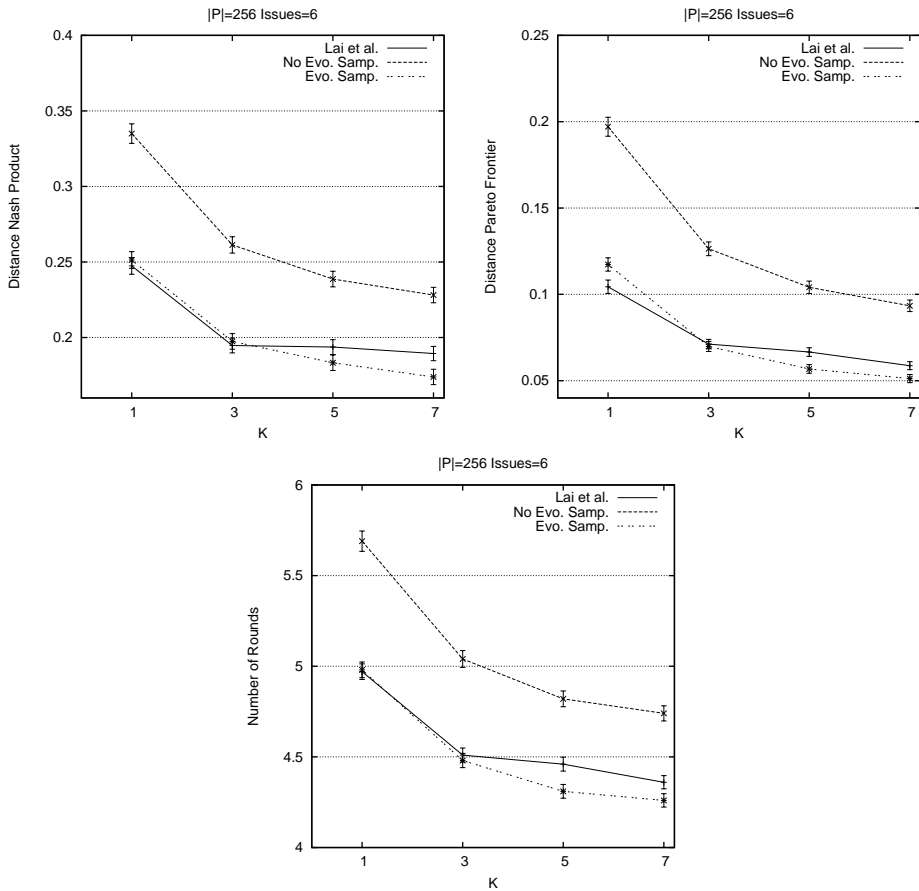


Figure 3.6: Distance to the Nash Product, distance to the closest Pareto Point, and number of negotiation rounds in Experiment 3.

3. BILATERAL NEGOTIATION FOR LIMITED DEVICES

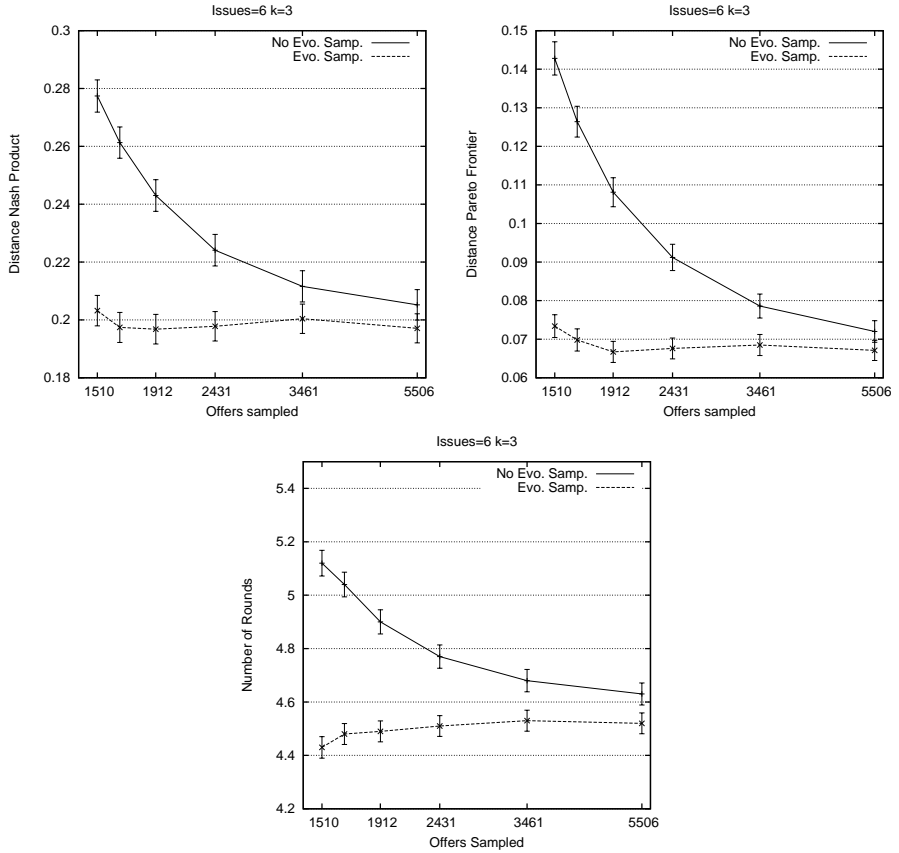


Figure 3.7: Distance to the Nash Product, distance to the closest Pareto Point, and number of negotiation rounds in Experiment 4.

4

A General Workflow of Tasks for Negotiation Teams

4.1 Introduction

In a negotiation, the steps that both parties have to take in order to implement an agreement have been studied by different scholars. Whenever parties engage in a negotiation, the following steps are usually necessary to finally implement an agreement (46): Identify social conflict, identify negotiation parties, structuring personal information, analysis of the opponents, define a protocol and select a negotiation strategy, negotiation (i.e., exchange of offers, argumentation, learning, etc.), and re-negotiation. In this chapter, we focus our study to negotiations between a negotiation team and an opponent party. However, the steps introduced in the workflow are general and could be easily adapted to multiparty negotiation where negotiation teams participate. The steps that are necessary to implement an agreement in such setting are similar to the steps proposed in the literature (46). However, some special considerations have to be taken into account since

4. A GENERAL WORKFLOW OF TASKS FOR NEGOTIATION TEAMS

at least one of the parties is composed of more than a single individual. In this chapter, we propose a general workflow of tasks for a negotiation team involved in a negotiation with an opponent party which refines previous general workflows (46). The workflow is designed from the perspective of an initiator agent that identifies a conflict situation and may need to form a team to negotiate with one or several opponents. The proposed schema aims to be general to be potentially adaptable to a wide range of domains.

The tasks that have been included in the workflow are thoroughly described in this chapter. For each workflow task, we attempt to identify which factors may be important, which problems may arise in each of the tasks, and which related work may help to efficiently solve the task at hand. The analysis is qualitative and descriptive.

The general workflow of tasks can be observed in Figure 4.1. In the proposed schema, we distinguish between tasks that are carried out with opponents (task with opp.), tasks that mainly concern interactions with team members (team task), tasks that only involve one individual agent (individual task), and tasks that involve team members, opponents, and the individual. The workflow is therefore divided into *Identify Negotiation* (individual task), *Team Formation* (team task), *Opponent Selection* (team task), *Understand Negotiation Domain* (team task), *Agree Negotiation Issues* (Task with Opponents), *Plan Negotiation Protocol* (team task), *Agree External Negotiation Protocol* (task with opponents), *Decide Intra-team Strategy* (team task), *Select Individual Strategy* (individual task), and *Negotiation & Adaptation* (team, individual and opponents). In the graph, the flow of tasks seems to follow a linear path. Nevertheless, it must be taken into account that we consider that in each step the team of agents may agree to re-plan a previous step according to new information acquired. The planning must

be seen as a continuous process where the team adapts itself to deal with unexpected events. However, if these unexpected events do not occur, the workflow is expected to follow the path depicted in Figure 4.1. Additionally, it should be also considered that since the workflow aims to be general, some of the tasks may not be necessary for some application domains. For instance, in some negotiation scenarios like the traveling friends domain, the negotiation team may be formed from the start of the problem since it was created by the users. If the workflow is to be used as a basis for a negotiation support system, the opponent selection phase may not be necessary since there is only one and it is already known by users. These are just some examples where a part of the workflow is skipped due to domain special features. Next, we describe all of the workflow tasks.

4.2 Identify Negotiation

The first step consists of identifying a conflict situation that requires negotiation. The agent has to analyze its environment and determine whether or not a conflict exists, the number of involved parties, potential partners, and whether or not it is convenient to form a team. As commented by Lopes et al. (46), most artificial intelligence researchers have focused on how to reach an agreement, but very few have studied the problem of detecting conflict. In multi-agent literature, one can identify works where conflict detection mechanisms are designed for specific domains like cooperative planning or air traffic management (96, 97, 98, 99). However, research in domain independent conflict detection mechanisms is a topic that needs further research. This research is especially important if one attempts to design general negotiators that are able to work in different domains. Some researchers like Lopes et al. (100) have employed libraries of axioms that allow agents to compare their own plans and intentions with those expected plans and

4. A GENERAL WORKFLOW OF TASKS FOR NEGOTIATION TEAMS

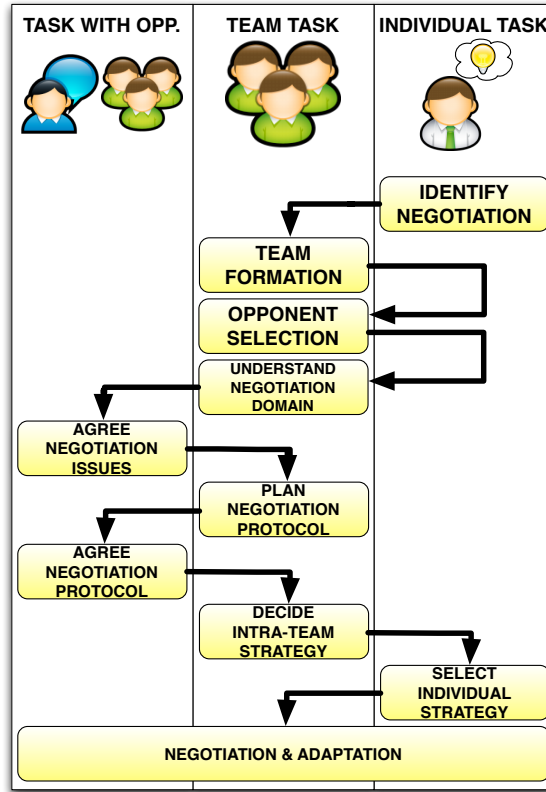


Figure 4.1: A general workflow of tasks for agent-based negotiation teams.

intentions of other agents to detect potential conflict. Libraries of axioms have the disadvantage that they are static unless they are provided with additional learning mechanisms. In that sense, case-based reasoning (54, 55) may help to have a library of conflict detection axioms that learns from the environment and evolves as the environment does.

Once a potential conflict has been identified by an agent, we propose that the agent needs to answer the following additional questions:

4.2 Identify Negotiation

- *Which agents are affected by the detected conflict?* This question aims to identify possible participants in the negotiation process.
- *Which agents do not share common goals with me?* Related to the previous question, the purpose of this question is to determine which potential participants may represent opponent parties. These potential opponents may form a unique opponent party, or they may be considered as individual opponents. In the latter case, it may be possible to conduct several negotiations threads in parallel and treat negotiation threads as outside options (101, 102, 103, 104).
- *Which agents share common goals with me?* The purpose is to determine the agents that may form a common party with the agent. If some of these agents cannot form a common party with the agent but they still may want to purchase the same product, they may be considered as competitors (101). For instance, in an e-commerce site, those agents that want to buy a unique and exclusive good compete with other agents that attempt to buy the exact same good.

Libraries can be complemented by search mechanisms. The technologies employed for looking for partners/opponents may vary. In fact, it heavily depends on the application where negotiation teams are deployed. One interesting technology is searching in social networks and markets (105, 106). For instance, in the traveling domain, one agent may look for travel companions based on its own social network and its extended social network (e.g., which may include the social network of its friends). Similarly, one agent could look for service providers or travel agencies based on the same type of algorithms.

4. A GENERAL WORKFLOW OF TASKS FOR NEGOTIATION TEAMS

In any case, the final product of this workflow task is a list of potential team members, a list of potential opponents, and, possibly, a list of potential competitors. With these lists, the agent faces the problem of selecting its team in the next workflow task.

4.3 Team Formation

Once the agent has studied which agents may be considered as potential partners, and which agents may be considered as opponents or competitors, the agent faces the challenge of determining whether benefits arise from forming a negotiation team (107). In some situations, it may be mandatory for the agent to be part of a negotiation team. In fact, the team may even be static (i.e., a married couple negotiating with a seller over an apartment). If that is the case, identifying negotiation partners, and forming a negotiation team are tasks that can be skipped from this workflow. Nevertheless, some scenarios may be less rigid and the agent may be able to form a negotiation team from the list of potential partners. Thus, the agent should analyze which team he expects to be the optimal negotiation team according to the list of potential partners, the list of opponents, and the list of competitors. If it is expected that no team reports more benefits for the agent than negotiating individually, the agent should decide to negotiate as a single individual party.

Traditionally, allocating agents into optimal groups has been a field of study for coalition formation (108, 109, 110, 111, 112, 113, 114, 115, 116). Many coalition formation algorithms focus on optimally dividing coalitional payoffs (112, 113, 114, 115, 116), which are the resulting benefits from carrying on a task as a group. In a team negotiation process, such benefits may be difficult to anticipate since it requires solving the problem of the negotiation with the opponent when the group

has not been formed yet. On top of that, the result of the negotiation may be an object whose payoff may not be divided among team members. For instance, in the case of the traveling friends, the final result of the negotiation is the travel. Even though the cost of the travel may be divided among team members, how do you expect to divide the benefits of other factors of the negotiation such as the payment method, arranged foods, hotel location, and so forth?

Another trend of research in coalition formation are buyers coalitions (108, 109, 110, 111): groups of buyers that join together in order to take advantage from volume discounts. However, most works in group buying have focused on single issue transactions where only price is involved and coalitional benefits can be shared. Therefore, complex multi-issue negotiations faced by negotiation teams are not supported by current group buying approaches.

Additionally, it should be noted that even though every groups of buyers may be considered a type of negotiation team, not every negotiation team is a group of buyers. For instance, let us imagine the negotiation between a union and the manager of an enterprise. The union may send a negotiation team formed by different experts or different stakeholders (i.e., representatives for different types of workers). In this case, the goal is not obtaining volume discounts as group buyers' case. Moreover, group of buyers are highly dynamic formations that may change when better coalitional options arise. That is not the case of the group of traveling friends or the union, where once the team has been formed, it usually remains static during the negotiation process.

Another field relevant to this workflow task is classic team formation (117, 118, 119). When teams are formed, agents with different skills are sought. Team members with different skills/expertise may be desirable for complex negotiation domains. Nevertheless, team literature in multi-agent systems usually assumes that team members share the same goal. Therefore, they are fully cooperative with

4. A GENERAL WORKFLOW OF TASKS FOR NEGOTIATION TEAMS

each other. Team members in a negotiation team may have different sub-goal or preferences despite sharing a common goal. Hence, team members may not be so cooperative with other fellow team members, especially when information sharing is involved. The ideal formula for negotiation formation may be a mixed approach between coalition formation and team formation. In any case, we propose that the following factors may be interesting to be considered when forming a negotiation team:

- Electronic commerce has given more social power to consumers, which now can find new sellers at a relatively low cost (120, 121). Not only that, but trust and reputation models (38, 39) and gossiping (122) may give an additional coercive power to consumers over sellers. This may produce sellers that are more willing to act cooperatively. Thus, it is expected that, the larger the negotiation team, the greater social power it will be able to exert, and the more cooperative the seller will be.
- Even though from the previous rule it seems intuitive that the larger a negotiation team is, the better, this may not be necessarily true in every case. If the preferences of the team members are compatible and very similar, adding new team members to a negotiation team may only result in greater social power. However, if the preferences of the team members are very dissimilar, adding new team members may result in greater intra-team conflict. Attempting to satisfy more preference profiles may considerably reduce the agreement space of the negotiation, and result in lower utilitarian outcomes. Thus, despite the fact that larger teams may be able to bring together more social power, intra-team conflict may deteriorate the quality of the final agreement to a point that greater social power does not compensate. Generally, it is not possible to exactly know the preferences of potential team

members prior to the negotiation itself. Nevertheless, past experiences with negotiation partners (123) and recommender techniques like collaborative filtering (124) may help to accomplish the task of assessing which negotiation partners are more similar.

- Related to the previous issue, more team members or stakeholders may bring additional negotiation issues to the negotiation table. The first effect over the negotiation is that the negotiation domain becomes larger and, possibly, computationally harder to work with. Despite this computational disadvantage, it may introduce negotiation issues that are only interesting to a sub-group of the agents that participate in the negotiation. If the opponent is not interested in these issues, it may make trade-offs easier. In contrast, if a negotiation issue is introduced to satisfy a sub-group of the team members, and it results in high conflict with the opponent, it may be difficult finding an agreement. Thus, additional issues in the negotiation are double-edged swords that can report both benefits and disadvantages.
- One of the problems that negotiation teams may face is tackling negotiation domains that are inherently complex. This means that the nature of the domain is hard to understand and it requires the expertise of different persons. For instance, when an organization negotiates in a complex negotiation, it sends a negotiation team composed of different experts. These experts may come from the different departments of the organization (e.g., marketing, human resources, research & development, etc.) and have different backgrounds that enrich the understanding of the problem. Information regarding agent identities (39) may come handy to determine which potential team members are more fit for the negotiation problem.

4. A GENERAL WORKFLOW OF TASKS FOR NEGOTIATION TEAMS

- Different agents may provide different social relationships to the team's social network. Social networks may directly impact upon the performance of teams (125) since it can provide with extra information for the team.

4.4 Opponent Selection

Once the team has been formed, it is necessary to find suitable opponents from the list of prospective opponents. The team should decide which opponents they are going to face. If enough computational resources are available, all of the opponents can be selected and negotiations can be carried out in parallel. However, if computational resources are scarce, a subset of the opponents has to be selected. If team members are rational, they should select the opponents that are expected to satisfy more one's own demands and the demands of the team members. In this sense, evaluating negotiation opponents based on the expected utility calculated from a set of past negotiation experiences (123) may prove an appropriate strategy. However, if no negotiation experiences are available about the different opponents, or there is not enough data to make conclusions about which opponents should be chosen, teams may resort discuss about the different opponents via argumentation in groups (126, 127, 128, 129, 130). Once the evaluation of the different opponents has been carried out, the selection of a subset of negotiating opponents may be carried out by means of classic social choice techniques like voting (131).

In any case, the final product of this workflow phase should be the list of opponents with whom the team pretends to negotiate.

4.5 Understand the Negotiation Domain

Understanding together the negotiation domain is a task of extreme importance. Not only does it allow team members to get a grasp of other team members' prefer-

4.5 Understand the Negotiation Domain

ences, but it makes it possible for team members to tackle correctly the negotiation when the domain is complex and requires expertise in different knowledge areas. For example, imagine that a negotiation team representing a software company negotiates with a client the development of a new software product. The negotiation team is composed of the manager of the software company, a representative from the R&D department, and a representative from the economic department. The manager of the software company and the representative from the economic department know that the price is a very important issue for the software company. Nevertheless, they may not be able to identify which technologies are viable for the product, which services are viable in the final product, and the development time for the application. Thus, they require the knowledge of the R&D representative during the negotiation. It is important that team members share knowledge about the negotiation domain prior to the negotiation, especially when team members have very different expertise. Otherwise, the negotiation may end up with an inefficient agreement (i.e., a high price but deadlines that are far beyond what is realistic).

Even assuming that team members have similar backgrounds, it is still important to understand the negotiation domain together. Let us imagine that a group of friends (e.g., Alice, Bob and Charlie) decides to go on a travel together and have fun. What is the meaning of “having fun”? Clearly, it may be different for each friend: Alice thinks that a city that offers lots of adventure sports is fun, Bob thinks that having fun also involves finding a place with a considerable night life, whereas Charlie is happy with any plan as long as it does not involve much money. From this situation, it can be inferred that the price, adventure activities included in the travel package, and the night life activities are relevant negotiation issues for the team in the negotiation at hand.

4. A GENERAL WORKFLOW OF TASKS FOR NEGOTIATION TEAMS

But identifying negotiation issues that are relevant for the team is not the only task necessary to completely understand the negotiation domain. Identifying which issues are predictable and compatible for the team, and which issues are not predictable is also crucial. On the one hand, a negotiation issue is *predictable and compatible* among team members if the preferences of all of the team members over issue values are known and compatible. For instance, in a team of buyers, it is logical that all team members prefer low prices over high prices. In this type of negotiation issues, there is full potential for cooperation among team members since increasing the utility for a team member (i.e., decreasing the price) results in other team members staying at the same utility or increasing their utility. On the other hand, a negotiation issue is not predictable among team members if nothing can be inferred about which issue values are preferred by team members. The issue may be compatible among team members (i.e., same ranking of preferences over issue values) or not, but it is not possible to know the nature of the negotiation issue unless team members are willing to share information. For example, in the team of traveling friends, it is not known whether team members prefer Rome to London, London to Rome, or they all prefer Berlin to London and Rome. Using information regarding which issues are *predictable and compatible* and *unpredictable* among team members may be useful for deciding on which negotiation strategy is used among team members.

The technologies that can give support to these processes are varied. As in any phase that involves deliberation and discussion, argumentation in groups (126, 127, 128, 129, 130) may be a useful technique to discuss regarding the negotiation domain and reach an agreement over which negotiation issues are relevant to the negotiation process. Other technologies like formation of shared expert mental models (132) and belief merging of multiple knowledge bases (133, 134, 135) may also prove useful for obtaining a shared model of the negotiation domain. However,

it should be noted that most belief merging methods are not strategy-proof (135, 136). A belief merging method is strategy-proof when it is robust against attempts of manipulation by agents. An agent may try to manipulate the belief merging process if it expects to increase its utility. Another interesting issue is whether or not agents have incentive to share all of the information regarding the negotiation domain. An agent may be willing to share a piece of information only if it expects that it is going to report higher utility than hiding the piece of information (i.e., selective information disclosure (137)). For example, taking up the example of the traveling friends, if Alice likes Rome, but she knows that Rome may not be a good place for night activities, she may hide this information from Bob in order to avoid making Rome less likeable by the group.

The final result of this phase should be a list of negotiation issues that are relevant to the team, and, ideally, an understanding of which issues are *predictable and compatible* and *unpredictable* among team members.

4.6 Agree Negotiation Issues

Since the previous stage produced a list of issues which is relevant for the team members, the next stage consists of agreeing a final list of negotiation issues with the opponent. The opponent may have its own list of issues relevant to the negotiation. Thus, a final list of issues to be negotiated should be agreed between both parties.

From the initial set of negotiation issues proposed by the team, some of the issues may not be negotiable since the opponent does not offer that service. For example, if the team members had originally concluded that negotiating packages of adventure activities is a relevant issue to the team but a travel agency does not work with such packages, the issue cannot be included in the negotiation.

4. A GENERAL WORKFLOW OF TASKS FOR NEGOTIATION TEAMS

Additionally, some negotiation issues that were not included in the list proposed by the team may be included in the final list since they are relevant to the opponent.

As for those negotiation issues present in the lists proposed by team members and the opponent, it may also be necessary to agree on the issue domain (i.e., the values that the negotiation issue can take). Similarly to the agreement on the list of issues, the final domain value may not contain all of the values proposed by both parties (i.e., Rome cannot be a value for the city of destination if the travel agency does not offer flights to Rome).

Despite being an important process in the pre-negotiation, very little attention has been paid to agreeing negotiation issues between parties. In fact, most researchers in negotiation assume that the list of negotiation issues and their domains are already agreed in their negotiation models. Faratin (20) mentions in his thesis the possibility of adding and removing non-core issues during the negotiation. While core negotiation issues remain static during the negotiation process, involved parties may be able to add or remove non-core negotiation issues as the negotiation process advances. However, the list of non-core issues is assumed to be known by both parties and the development of an issue-manipulation algorithm was appointed as future work. We acknowledge that this is a process that needs to be researched in the future.

4.7 Plan Negotiation Protocol

After the list of issues is set, the parties have to agree a negotiation protocol. There are different negotiation protocols that may be applied for a specific situation. For instance, if the negotiation team engages with an opponent in a bilateral negotiation, both parties could employ the classical alternating offers protocol

4.8 Agree Negotiation Protocol

(12), extensions of such protocol like the k -alternating offers protocol (57), or more complex protocols like (138, 139).

The team as a whole may have different opinions and knowledge about the available protocols. In fact, some of the team members may not even know some of those protocols. In that case, those protocols cannot be used by the team since some of its players do not know the rules and decision making strategies to face such games.

In this phase, if more than an applicable protocol is known by all of the team members, they should decide as a group which protocols are preferred by the team (i.e., a ranking of the known protocols). This may be based on the expertise of each agent in the aforementioned protocol, computational efficiency, decision making mechanisms known by team members, and so forth. Given the assumption that the set of known protocols for a specific situation is limited and probably small, team members may employ argumentation techniques (126, 127, 128, 129, 130) followed by a voting mechanism (131) to decide on a ranking of the available protocols. As far as we are concerned, very little work has been carried out with respect to evaluating negotiation protocols. The only exception is presented in Miller et al. (140). The authors propose a framework where protocols are not imposed at design time. Instead, protocols are inside a dynamic library at runtime. Agent are able to analyze, instantiate and reason regarding the possible outcomes that the protocol may entail.

4.8 Agree Negotiation Protocol

Considering that the team has already decided on which negotiation protocols are preferred by team members (i.e., some sort of ranking over the negotiation protocols), they should agree with the opponents on the negotiation protocols that

4. A GENERAL WORKFLOW OF TASKS FOR NEGOTIATION TEAMS

are to be used for interacting. Again, opponents may not know how to play some games, making some of the options not feasible. Some protocols known by the opponent may not be known by all of the team members. Over the list of protocols that are known by both parties, both parties may have different preferences and knowledge regarding the different protocols. This decision between both parties is going to involve some kind of simple negotiation (i.e., we do not expect the number of possibilities to be large) or discussion among both parties. In some cases, besides the negotiation protocol, some parameters of the protocol have to be decided also by both parties (i.e., who is the initiating party in the alternating offers protocol (12), the number of offers allowed in the k-alternating offers protocol (57), who acts as trusted mediator in mediated protocols like (138, 139), etc.). Some authors like (141, 142) have started to tackle the problem of negotiating over negotiation protocols. In (141), a formal framework is presented for expressing and constructing dynamic negotiation protocols in open environments. The construction of negotiation protocols is based on basic pieces named as dialogue acts (i.e., basic communication particles) that aim to solve one of the specific goals of the negotiation protocol. However, how a group of agents may agree on such protocols is not specifically covered. Reed et al. (142) can be considered as a complement to the previous approach. Basically, the authors present a framework where agents agree on the semantics of communicative particles, which may be later used to construct dynamic communications in open multi-agent systems.

The final result of this phase is a set of negotiation protocols that will be followed in the negotiations with the different opponents.

4.9 Decide Intra-team Strategy

Negotiation protocols define the rules of interaction to be followed by the different parties. For instance, it indicates when the different parties can make offers/arguments, which kind of messages they are expected to send/receive, etc. Generally, these interactions involve one of the parties taking a particular decision (i.e., which offers is sent in the alternating offers protocol, which type of argument is sent to convince about one's own position, the information that is leaked to the other parties, etc.). In a single player party, how these decisions are taken are up to the agent. However, when the party is formed by multiple individuals, which is the case of the negotiation team, the team has to decide on *how*, *when*, and *what* decisions are taken, and *who* takes those decisions. This is what we termed as an intra-team strategy or team dynamics.

For example, in the case of the alternating offers protocol (12), each party should decide on which offer is sent to the opponent party, whether or not to accept the offer proposed by the other party, and when one should withdraw from the negotiation process. Thus, any intra-team strategy for teams participating in the alternating offers protocol should decide on those issues.

For the same negotiation protocol, there may be different intra-team strategies. In the case of the alternating offers protocol, the team may delegate on one of the team members to take all of the decisions or some/all of the may involved in the decision-making processes of the team. The list of offers to be sent to the opponent may be decided prior to the negotiation with the opponent, or it may be dynamically constructed as the opponent gives its feedback in the negotiation process by means of its counter-offers. As for the acceptance of opponent's offers, the team may decide on using voting mechanisms with different unanimity rules.

4. A GENERAL WORKFLOW OF TASKS FOR NEGOTIATION TEAMS

In fact, the spectrum of intra-team strategies for a certain negotiation protocol may be as large as to be considered infinite.

Obviously, for different negotiation protocols, different decisions have to be taken and an intra-team strategy that has been proposed for a particular negotiation protocol may not be directly applied to other type of negotiation protocol. Thus, usually negotiation protocols and intra-team strategies are tightly coupled.

Another interesting problem that may arise in several intra-team strategies is role/task allocation. For instance, if an intra-team strategy relies on the selection of a representative that will act on behalf of the team, who plays such role? If the intra-team strategy requires that an agent coordinates voting processes, who acts as a trusted mediator? The role of trusted mediator may be taken by an agent from inside the team or by an external agent. It is also known that negotiation may involve several tasks like looking for outside options, seeking information, and monitoring the market. How are these tasks divided between the team members? Traditionally, multi-agent teamwork literature has been especially fruitful in the area of task/role allocation (36, 143, 144, 145, 146). Nevertheless, as far as we know, none of these approaches have been explicitly applied to teamwork in negotiation teams. Partly, this may be explained by the fact that most studies in multi-agent teamwork have focused on teams where all of the team members share a joint goal. Despite the fact that negotiation teams have a common joint goal (otherwise, they would not collaborate), each team member may have its own individual goals. Teams with mixed motives have not been as extensively studied in multi-agent literature with the exceptions of (147, 148). Therefore, it is necessary to study to what extent team members would fully collaborate in negotiation team's tasks in spite of their own utility. For example, how interesting is it for a team member to look for new outside options for the team when current ones

4.10 Select Individual Strategy

report high utility for himself? A self-interested team member may decide to neglect its search tasks and continue with present outside options if it considers that new outside options will not increase its current welfare.

One of the hypotheses of this thesis is that there is not a single intra-team strategy that is capable of outperforming the rest of intra-team strategies for every possible scenario. Depending on the goal of the team (e.g., social choice performance measure), and depending on the conditions of the negotiation environment (e.g., team size, similarity among team members' preferences, opponent concession, deadline length, etc.), some intra-team strategies will perform better than others. Some researchers have proposed the use of extensive simulation in the laboratory to assess which strategies would work better in certain specific conditions (63, 64, 65, 66, 149, 150). The results of these simulations can provide profitable knowledge to be used when the agents face the challenge of selecting an appropriate strategy. Despite the fact that these simulations have been carried out in the bilateral setting for single individual parties, no study exists for the team case.

4.10 Select Individual Strategy

Each team member should plan its individual strategy before heading into the negotiation. An intra-team strategy defines mechanisms for team decision-making but they do not define how individual team members behave when playing those mechanisms. It is up to the agent to decide how to act inside the team: it can be more or less cooperative. The agent should also decide its attitude with the opponent, which may be classified as competitive, conceding, matcher or inverter (151). The two aforementioned factors will define the initial negotiation strategy of each team member.

4. A GENERAL WORKFLOW OF TASKS FOR NEGOTIATION TEAMS

Generally, the selection of the initial negotiation strategy is based on what is expected about the opponent and teammates. As stated in the previous section, one of our hypotheses is that the state of the negotiation environment plays a key role in selecting which intra-team strategies are more appropriate for each specific situation. Thus, team members should also decide on their individual strategy based on the knowledge about the negotiation environment.

4.11 Negotiation & Adaptation

The final phase is the negotiation itself. During this phase, team members should follow the planned intra-team strategies, individual strategies, and negotiation protocols. However, negotiation is a dynamic process that may not go as planned (e.g., opponents not behaving as one initially thought, team members performing below/above one's expectations, members leaving the team, etc.). Therefore, it may be necessary that each team member adapts its own negotiation strategy, and that the team replans some of the aspects related to team composition and team dynamics. More specifically, we argue that it may be interesting to study the following adaptation problems:

- **Team membership:** As stated, team membership may be dynamic. In fact, how dynamic a negotiation team is may depend on the application domain. Domains where team members are more self-interested and less bonds exist between team members (e.g., team of buyers) may be more dynamic than domains where team members are more cooperative and there are human bonds (e.g., group of travelers, human organizations, etc.). In any case, in both situations the problem of dynamic membership may arise. For instance, new buyers may appear in the electronic market and they may be added to the team to take advantage of larger price discounts. Similarly, a new

traveler may decide to travel during holidays and his user may state the desire of joining the pre-existing group of travelers. Cases of team member's withdrawal are also possible. For instance, one of the buyers participates in other buyers' coalitions and it decides to close a deal, making its membership in the rest of the buyers' coalitions no longer necessary.

- **Negotiation issues:** Initially, both parties agreed to negotiate over some initial issues. For some reasons (e.g., computational issues, computational tractability, etc.), they may have decided to leave some less relevant issues out of the negotiation. However, at some points an impasse (152, 153) may occur in the negotiation. A negotiation impasse occurs when the parties are unable to reach an agreement and the perspectives of reaching one are very negative. They are in a deadlock. A possible solution for such problematic situation is what is known as issue linkage (154, 155, 156). Basically, when parties negotiate on one issue, adding another issue and linking its value to the value of the initial issue can increase the probability of finding an agreement. The new issue may be added to reduce intra-team conflict (e.g., how costs are split in the team), or they may be added to reduce conflict with the opponent (e.g., include a payment method issue and maximize the preferences of the opponent in the new issue). This adaptation heuristic may be positive for cases that are prone to fail. However, as suggested by (155, 156), issue linkage may also have negative effects since it may also reduce the agreement space. As of today, issue linkage is an area that has not been widely studied in automated negotiation, where it has been assumed that issues remain static during the negotiation process. Hence, it is an area that requires further exploration, especially for the team case since conflict may appear at the team level and the opponent level.

4. A GENERAL WORKFLOW OF TASKS FOR NEGOTIATION TEAMS

- Intra-team strategy and individual strategy adaptation: As stated, intra-team strategies define *what* decisions are taken by the team, *how* decisions are taken, and *when* those decisions are taken. Assuming that team members chose the best intra-team strategy according to the initial negotiation conditions and expectations, it may be possible that one of such conditions and expectations changed during the negotiation process, precluding the initial intra-team strategy from being the best choice. In that case, it may be wise for team members to change their current intra-team strategy in order to match the new changes in the negotiation environment. Obviously, changes in the intra-team strategy and environment's condition also call for an adaptation in team members' individual strategies. In this sense, there have been some works that advocate for a change in individual agents' strategies in bilateral negotiations (60, 101, 102, 103, 149, 157, 158, 159). All of these works show the benefits of adapting one's behavior during the negotiation to achieve better results. We can distinguish between works where individual agents adapt their behavior attending to environmental conditions like outside options and competitors (101, 102, 103) and works where individual agents adapt their behavior during the negotiation attending to the attitude of the opponent (60, 149, 157, 158, 159). However, as far as we are concerned, these techniques have not been extrapolated to the team case.

4.12 Conclusions

In this chapter, we have proposed a general workflow of tasks that may help agent-based negotiation teams to perform successfully in their correspondent applications. Next, we conclude by overviewing these tasks in order of appearance, and outlining the major challenges that may arise in each of them.

- **Identify Negotiation:** The first step consists of being able to perceive when conflict is present and who may be involved in the conflict (i.e., opponents, prospective teammates, and competitors). Conflict detection in open systems may be of extreme importance, and adaptive mechanisms may be needed to detect conflict in such systems. We argue that case based reasoning may be a useful technology since it may be able to detect and anticipate conflict based on past experiences, and it may be able to learn from new situations. Additionally, agents may employ search mechanisms in networks to be able to discover new partners, opponents, and competitors. It should be stated that most work in similar areas has focused on domain specific conflict detection or static general rules for detection of conflict. Thus, this area of work remains largely unexplored.
- **Team Formation:** If the agent thinks that it may be beneficial to form a negotiation team, it should attempt to select its teammates. Closely related research areas are coalition formation and cooperative team formation. The former has focused on forming optimal groups of agents and how to divide the payoffs of the group task. However, the result of the negotiation may be difficult to anticipate and, while some negotiation issues like price may be naturally divisible, others may be hard to be divided (e.g., payment method). Furthermore, negotiation teams may not be able to be disbanded or join other teams when better coalitional options appear. Cooperative team formation aims to form teams based on complementary skills for a certain task. Nevertheless, team members may not be fully cooperative since they have their own and possibly conflicting individual interests. An ideal solution to this task may inherit features from both cooperative team formation and coalition formation and, additionally, it may need to

4. A GENERAL WORKFLOW OF TASKS FOR NEGOTIATION TEAMS

take into account factors specific to negotiation teams like (i) the relationship between team size and social power; (ii) the relationship between team size, team similarity, intra-team conflict and conflict with the opponent; (iii) requiring different knowledge expertise; (iv) the social network provided by each team member. In conclusion, even though some related research has been carried out, it may need to consider additional issues that are specific to negotiation teams.

- **Opponent Selection:** The next step consists of selecting the opponents with whom the team will negotiate. In this task, related research exists that could be directly employed like selecting opponents based on past experiences and arguing or using social choice to assess the best available options based on other information sources.
- **Understand the Negotiation Domain:** The general idea behind this task is creating a shared knowledge model of the negotiation domain at hand. It includes identifying negotiation issues, merging different points of views and expertise, clarifying team goals that may be abstract in essence, and identifying the nature of prospective negotiation issues (e.g., predictable, unpredictable, compatible, etc.). Some related research in this area may be argumentation and belief merging. However, it should be taken into account that some team members may be self-interested and they may show opportunistic behaviors (e.g., manipulating belief merging, hiding relevant information for one's own interest, etc.).
- **Agree Negotiation Issues:** The next part consists of agreeing with the opponent which negotiation issues should be considered in the negotiation. Despite its importance, negotiation models usually assume that the negotiation domain as given and they do not provide mechanisms that allow forming

or negotiating the domain. This is potentially one of the most interesting research challenges since, as far as we know, related literature is almost nil.

- **Plan Negotiation Protocol:** Given a specific situation, there may be different negotiation protocols that may be used to negotiate with the opponent. Team members should argue about which protocols are preferred according to his experiences, strategies known, and so forth. Although the problem has not been explicitly covered by the literature, its solution may not pose exceptional efforts compared to classic argumentation and social choice problems.
- **Agree Negotiation Protocol:** Similarly, once team members have discussed about the available negotiation protocols, they should negotiate a proper negotiation protocol and its parameters with the opponent. The problem may not be different from any other negotiation.
- **Decide Intra-team Strategy:** Intra-team strategies define team dynamics for a specific negotiation protocol. This refers to the coordination and negotiation protocol carried out within the team to decide on the steps to be carried out in the negotiation with the opponent. If the intra-team strategy requires role differentiation, techniques from role/task allocation may be employed. However, it should be considered that agents may not be fully cooperative. Thus, the problem slightly differs from classic role/task allocation. Additionally, team members may employ information regarding the current environment state (e.g., deadline length, number of competitors, team size, beliefs regarding the opponent, etc.) in order to decide on the most appropriate intra-team strategy. We have identified that even though some studies exist that identify good practices and good strategies for single individual parties, the area remains largely unexplored for intra-team strategies.

4. A GENERAL WORKFLOW OF TASKS FOR NEGOTIATION TEAMS

- **Select Individual Strategy:** The intra-team strategy defines team dynamics, but it does not define the individual behavior of team members per se. The next step consists of each team member deciding on its own the most appropriate individual behavior for the negotiation at hand. This task may not pose additional difficulties compared to the selection of the individual negotiation strategy in classic negotiations.
- **Negotiation & Adaptation:** The final task of this workflow consists of carrying out the negotiation and adapting some of the decisions taken in order to properly face unexpected events. We have identified three main aspects that may be adapted in a negotiation team. The first of them is team membership since team members may join and leave the team during the negotiation. In that sense, the mechanisms needed for this adaptation may not differ so much from the ones employed in team formation. The second aspect that may be adapted is negotiation issues. Parties can solve impasses in the negotiation and better off other parties by including other issues that were not initially included in the negotiation. As far as we are concerned, this problem has not been widely studied in multi-agent literature. The third and final factor that we consider is the intra-team strategy and the individual strategy. Usually, team members have planned on using an intra-team strategy and an individual strategy based on some initial prediction of the negotiation environment and teammates' behavior. However, based on new evidence, initial predictions may prove wrong and adjustments need to be done in order to properly tackle the negotiation. In automated negotiation, some works exist that allow single individual parties to adapt themselves to changes in the negotiation environment and new information. These mechanisms could be employed as long as they were adapted to the negotiation team case.

It must be stated that, in this thesis, our goal is not to detail how to carry out each task. Each of the proposed workflow tasks may give room for an individual PhD thesis by itself and needs to be studied in-depth. In this thesis we focus on solving the tasks related to the negotiation, although some of our models also cover tasks in the pre-negotiation (e.g., information sharing). In any case, we think the type of analytic studied carried out in this chapter may help the reader to understand the complexity involved in negotiation teams and make clear some of the problems for future researchers in the field. Out of the tasks proposed in this workflow, we decided to focus on the study of intra-team strategies in this thesis. The reason behind this decision is that given that intra-team strategies govern team dynamics and negotiaton decisions, they are expected to have a greater impact on team performance. Thus, in Chapters 5 and 6 we propose and validate different intra-team strategies. The intra-team strategies in Chapter 5 are evaluated in domains where negotiation issues are *predictable and compatible* among team members, whereas intra-team strategies in Chapter 6 are evaluated in domains with *predictable and compatible* and *unpredictable issues*.

4. A GENERAL WORKFLOW OF TASKS FOR NEGOTIATION TEAMS

5

Intra-Team Strategies for Negotiation Teams in Predictable Domains

5.1 Introduction

As mentioned, among the different tasks that a negotiation team has to face, we decided to focus on intra-team strategies. Intra-team strategies define team dynamics during the negotiation process. We consider that due to the fact that intra-team strategies govern the decision making of the team, they directly affect team performance. Hence, the direct effect on team performance drove our research towards intra-team strategies.

In this thesis, we are interested in intra-team strategies for teams whose members may have different preferences regarding the negotiation issues. In this chapter, we start our study on the subject by focusing on models for electronic markets where negotiation issues are *predictable and compatible* among team members. By

5. INTRA-TEAM STRATEGIES FOR NEGOTIATION TEAMS IN PREDICTABLE DOMAINS

issue predictability and compatibility among team members we refer to the fact that the preferences of the team members over issue values are known and compatible. For instance, in a team composed by buyers, it is quite reasonable to assume that all of the buyers prefer low prices to high prices, high quality to low quality, short dispatch time to large dispatch time, and so forth. This will be translated to the fact that even though the exact shape of the valuation function is not known, the type of valuation function for negotiation issues is *predictable and compatible* among team members: a ranking for issue values is known for each issue and it is common among team members. This kind of assumption is common in some electronic commerce scenarios where team members share the same role (e.g., buyers). For instance, buyers usually share the same type of valuation function for attributes such as the price (i.e., they prefer lower prices to higher prices), product quality (i.e., they prefer higher quality to lower quality), and the dispatch time (i.e., they prefer shorter dispatch times to longer dispatch times).

In this scenario, we propose four intra-team strategies for a negotiation team that negotiates with a single opponent by means of the alternating offers protocol (12): representative (RE), similarity simple voting (SSV), similarity borda voting (SVB) and full unanimity mediated (FUM). These strategies are designed to cover four minimum levels of unanimity regarding team decisions: no unanimity guaranteed (i.e., representative), plurality/majority (i.e., similarity simple voting), semi-unanimity (i.e., similarity borda voting) and unanimity (i.e., full unanimity mediated). Among these intra-team strategies, we put a special emphasis on full unanimity mediated since it is able to guarantee unanimity regarding team decisions (i.e., full unanimity mediated). Our belief is that, when possible, unanimity among team members is a very important feature for negotiation teams models. Agreements that are unacceptable for a team member should be avoided since they might deteriorate human relationships. Technologies that help to form unanimous

decisions may provide more user satisfaction, and they can help team members to avoid unexpected outcomes. Hence, a special interest from our part is put into intra-team strategies that are able to guarantee unanimity regarding team decisions.

As stated in the introduction of this thesis, one of our initial hypothesis is that environmental conditions may affect the performance of intra-team strategies. It has been documented that environment conditions such as the deadline, concession speed, and reservation utility may affect the impact of single-individual bilateral strategies (18). However, in the team case, new conditions like the number of team members, team preferences' diversity, and the emergent effect of aggregating team members' behaviors/actions may also end up affecting team performance. Prior to the negotiation process, negotiation teams face the challenge of selecting which intra-team strategy should be employed. If environmental conditions have an effect on the performance of the different intra-team strategies, the intra-team strategy for the negotiation at hand should be selected accordingly to the current environmental conditions inferred by team members. One of our research goal is identifying how these environmental conditions may affect the different intra-team strategies presented in this chapter. Due to the large amount of variables that may affect the negotiation, we employ an empirical approach to study the behavior of the four intra-team strategies.

The chapter is organized as follows. First, we describe the assumptions of our framework (Section 5.2). Then, we illustrate a motivating negotiation scenario for the strategies presented in this chapter. The motivating scenario is based on a group booking domain. After that, the details of the four intra-team strategies are thoroughly described in Sections 5.4, 5.5, 5.6 and 5.7. Then, in Section 5.8, we describe in depth some of the empirical evaluations carried out to study the behavior of full unanimity mediated. In Section 5.9, the experiments carried out to

5. INTRA-TEAM STRATEGIES FOR NEGOTIATION TEAMS IN PREDICTABLE DOMAINS

analyze the different intra-team strategies in different negotiation environments. Finally, we briefly state the conclusions of this chapter in Section 5.10.

5.2 Negotiation Setting

- The team A is formed by M different agents a_i , $1 \leq i \leq M$ ($A = \{a_1, a_2, \dots, a_M\}$). It should be stated that team membership is considered static during the negotiation process. Dynamic agent-based negotiation teams are not considered in this thesis, and they are appointed as future lines of work.
- The common goal of the team A is negotiating a successful deal with the opponent op . Thus, in this case we assume an implicit representation of the teams' goal.
- It is assumed that information is private, even among team members. Therefore, agents do not know other agents' utility functions, strategies, reservation utilities, or deadlines. We also assume that agents have limited computational resources (i.e., bounded rationality). Thus, we take a heuristic approach which seeks good quality results while being computationally tractable.
- It is assumed that the team A and the opponent op communicate following an alternating bilateral protocol (12). One of the two parties acts as the initiator, and is entitled to propose the first offer. The other party receives the offer and can respond with two different actions: accept the offer (successful negotiation), or propose a counteroffer. If a counteroffer is proposed, the initiator party receives the offer and it can either accept the counteroffer or propose another offer, starting a new negotiation round. Depending on the intra-team strategy, one of the team members or a team mediator is

responsible of the communications with the opponent. In this setting, the fact that one of the parties is a team is not directly stated to the other party, although in some domains it may be logical to think that the party is formed by multiple individual (i.e., group booking).

- Additionally, it is also assumed that the negotiation is time-bounded, and each party has a private deadline T_A (team deadline), T_{op} (opponent deadline). When its deadline is achieved, the party leaves the negotiation and it is considered a failed negotiation. In the case of T_A , it is considered a joint deadline for all of the team members, who have agreed upon this deadline prior to the negotiation at hand.
- The team mediator, if present, is never a perfect mediator that aggregates the utility functions of all the team members. This assumption is taken due to the fact that, depending on the application, some team members may not be completely trustable and they may attempt to exaggerate/change their preferences to manipulate the negotiation process. This mischievous behavior is easily carried out when aggregating utility functions.
- The negotiation domain is comprised of n issues. A complete offer is represented as $X = \{x_1, x_2, \dots, x_n\}$, where x_i is a specific instantiation of issue i . Additionally, we use the notation $X_{i \rightarrow j}^t$ to denote that offer X was sent at round t from party i to party j .
- Every agent i (team member or opponent) has its preferences represented by means of additive utility functions in the form:

$$U_i(X) = w_{i,1} V_{i,1}(x_{i,1}) + w_{i,2} V_{i,2}(x_2) + \dots + w_{i,n} V_{i,n}(x_n) \quad (5.1)$$

where X is a complete offer, x_j , is the value given to the j -th issue, $V_{i,j}(\cdot)$ is the valuation function for issue j used by agent i to normalize the issue

5. INTRA-TEAM STRATEGIES FOR NEGOTIATION TEAMS IN PREDICTABLE DOMAINS

value to $[0, 1]$, and $w_{i,j}$ is the weight/importance given by agent i to issue j in the negotiation process. Several observations should be made regarding these utility functions:

- Weights are normalized so that $\sum_{j=1}^n w_{i,j} = 1$.
- Issues are assumed to be independent from each other. Thus, the valuation of one of the issues does not alter the others issues' valuation.
- Negotiation issues are *compatible and predictable* among team members. An issue j with domain D_j is compatible among team members if for each pair of team members $a, b \in A$, and for each pair of issue values $v_1, v_2 \in D_j$, the following expression is true:

$$V_{a,j}(v_2) \geq V_{a,j}(v_1) \longleftrightarrow V_{b,j}(v_2) \geq V_{b,j}(v_1). \quad (5.2)$$

Hence, an issue is compatible among team members if changing its value (v_1) with another value (v_2) increases/decreases a team member's utility, then v_2 would also increase/decrease the utility for other members. Thus, there is potential for cooperation among team members. Examples of functions that are *compatible and predictable* are monotonically increasing valuation functions or monotonically decreasing valuation functions. Moreover, we assume that team members share the same type of monotonic function (i.e., increasing or decreasing) for each $V_{i,j}(\cdot)$. As for the opponent, it is assumed that the monotonic function for $V_{i,j}(\cdot)$ is the opposite type to that of team members. It is reasonable to assume this model for valuation functions in e-commerce scenarios. Buyers usually share the same type of valuation function for issues such as the price (monotonically decreasing), product quality

(monotonically increasing), and the dispatch time (monotonically decreasing), whereas sellers usually use the opposite type of monotonic functions (monotonically increasing for price, monotonically decreasing for product quality, and monotonically increasing for dispatch time).

- Issue weights $w_{i,j}$ are different for team members. This way, we are able to represent the fact that some team members may be more interested in some issues whereas other team members may be more interested in other issues (e.g., some team members prefer price over quality, while others give a higher priority to the product quality). The weights of the opponent's utility function may be different from those of team members.
- The opponent has a reservation utility RU_{op} . Any offer whose utility is lower than RU_{op} will be rejected. Each team member a_i has a private reservation utility RU_{a_i} . This individual reservation utility is not shared among teammates. Therefore, a team member a_i will reject any offer whose value is under RU_{a_i} . In this setting, reservation utilities represent the individual utility of each agent if the negotiation process fails.
- In our negotiation model, we define that a final decision X (i.e., final agreement) is unanimously acceptable among team members when the utility reported by such decision for each team member is equal to or greater than the reservation utility of each team member:

$$\forall a_i \in A, U_{a_i}(X) \geq RU_{a_i} \tag{5.3}$$

5.3 Motivating Example: Group Booking

In this section we present a scenario which can be modeled by means of a negotiation between a team and an opponent. Its purpose is to illustrate the behavior of the different intra-team strategies when they are described in Sections 5.4, 5.5, 5.6 and 5.7. The scenario involves a group of persons that need to book an hotel in a city for a group vacation. We have named this scenario as *Group Booking*.

In this scenario, a group of friends (e.g., Alice, Bob and Charlie), who have decided to spend their holidays together, has to book accommodation for their stay. Their destination is Rome, and they want to spend a whole week. Each friend is represented by his/her electronic agent (a_a, a_b, a_c), who acts semi-automatically on behalf of its user. This agent has previously elicited the preferences of its user regarding booking conditions. Each group member has different preferences regarding possible booking conditions. Thus, the final agreement with the hotel should satisfy every friend as much as possible. This leads to the use of an agent-based negotiation team model. The group of agents, engages in a negotiation with a well-known hotel in their city of destination, which is also represented by an electronic agent (op). During the pre-negotiation, both parties have decided to negotiate the following issues:

- Price per person (pp): The price per person is the amount of money that each friend will pay to the hotel for the accommodation service. The issue domain goes from 210\$, which is the minimum rate (30\$ per night), to 700\$, which is the maximum rate (100\$ per night). A realistic assumption in the group of friends is that friends prefer to pay lower prices to higher prices (monotonically decreasing valuation function), whereas the seller prefers to charge higher prices to lower prices (monotonically increasing valuation function).

5.3 Motivating Example: Group Booking

- Cancellation fee per person (*cf*): When a booking is cancelled, the hotel deletes the reservation but it charges a fee to compensate for losses. The issue domain goes from 0\$ (no cancellation fee) to 150\$. A realistic assumption in the group of friends is that friends prefer to pay lower prices to higher prices (monotonically decreasing valuation function), whereas the seller prefers to charge higher prices to lower prices (monotonically increasing valuation function).
- Full payment deadline (*pd*): The full payment deadline indicates when the group of friends has to pay the full price to confirm their reservation. The domain goes from “Today”=0 days (the date time when the final agreement has been signed) to “Departure Date”=30 days, which indicates that the team should only pay when leaving the hotel. A realistic assumption in the group of friends is that friends prefer to pay as late as possible (monotonically increasing valuation function), whereas the seller prefers to charge as soon as possible (monotonically decreasing valuation function).
- Discount in bar (*db*): As a token of respect for good clients, the hotel offers nice discounts at the hotel bar. The issue domain goes from 0% (no discount) to 20%. A realistic assumption in the group of friends is that friends prefer higher discounts to lower discounts (monotonically increasing valuation function), whereas the seller prefers to offer lower discounts prices to higher discounts (monotonically decreasing valuation function).

For illustrative purposes thorough this chapter, the users’ preference profiles of each friend have been modeled by means of additive utility functions as follows:

- Alice is sure that she wants to go on a travel, but she is short on budget right now. Thus, she gives more importance to the price of the hotel and the

5. INTRA-TEAM STRATEGIES FOR NEGOTIATION TEAMS IN PREDICTABLE DOMAINS

payment date. She thinks that the hotel bar is very expensive, so having food outside the hotel is a better option than using discounts at the bar. After the elicitation process, Alice utility function was elicited with the following weights:

$$U_{a_a}(X) = 0.5 V_{pp}(x_{pp}) + 0.0 V_{cf}(x_{cf}) + 0.5 V_{pd}(x_{pd}) + 0.0 V_{db}(x_{db}) \quad (5.4)$$

- Bob is not very sure about being able to travel in the agreed dates since he may have to attend a conference. Thus, one of his main priorities is minimizing the cancellation fee. Additionally, he is moderately worried about the hotel price but he does not give much importance to the payment date and discounts at the bar. Bob utility function was elicited with the following weights:

$$U_{a_b}(X) = 0.3 V_{pp}(x_{pp}) + 0.6 V_{cf}(x_{cf}) + 0.05 V_{pd}(x_{pd}) + 0.05 V_{db}(x_{db}) \quad (5.5)$$

- On the one hand, Charlie is equally worried about the hotel price and the payment deadline, but he is sure about going on a travel. On the other hand, Charlie is a fan of good food, and he has heard very good reviews about the hotel bar. He thinks that the discounts are a good opportunity to taste the culinary specialties of the hotel bar. His utility function was elicited with the following weights:

$$U_{a_c}(X) = 0.35 V_{pp}(x_{pp}) + 0.1 V_{cf}(x_{cf}) + 0.35 V_{pd}(x_{pd}) + 0.2 V_{db}(x_{db}) \quad (5.6)$$

Alice, Bob and Charlie share the same type of monotonic functions for the valuation functions $V_j(\cdot)$. In this example, monotonically increasing valuation functions have been modeled as $V_j(x) = \frac{x-x_{min}}{x_{max}-x_{min}}$ and monotonically decreasing functions have been modeled as $V_j(x) = 1 - \frac{x-x_{min}}{x_{max}-x_{min}}$. For instance, the valuation

5.4 Representative (RE)

function for the price would look like $V_{pp}(x_{pp}) = 1 - \frac{x-210}{700-210}$, and valuations for 210\$, 400\$, 500\$, and 700\$ would result in $V_{pp}(210) = 1$, $V_{pp}(400) = 0.61$, $V_{pp}(500) = 0.40$, and $V_{pp}(700) = 0$ respectively.

As for the reservation utility, the three friends also have different options. While Alice may buy a new TV in case that the negotiation fails ($RU_{a_a} = 0.3$), Bob can go to his parents' apartment ($RU_{a_b} = 0.2$), and Charlie's only alternative vacation plan is going camping ($RU_{a_c} = 0.3$). The electronic agents of the three friends have decided that they can attempt to interact up to 100 times with the opponent ($T_A = 100$), and they think that a linear concession would be adequate for the negotiation at hand ($\beta_A = 1$) given the current conditions.

It should be pointed out that all of the preference profiles created for all of the experiments carried out in this chapter correspond to our motivating scenario. Once we have described the general assumptions of our negotiation model for *compatible and predictable* domains, we describe each of the intra-team strategies thoroughly.

5.4 Representative (RE)

The representative strategy (RE) is perhaps the simplest intra-team strategy. Basically, one of the team members is selected as representative a_{re} for the team during the negotiation. This agent acts on behalf of the team during the negotiation, making it responsible of selecting which offers are sent to the opponent, and whether or not opponent's offers are accepted. The only communications are those carried out between the representative agent a_{re} and the opponent a_{op} . Therefore, this strategy is equivalent to a classic bilateral strategy.

The representative agent negotiates according to its own utility function $U_{a_{re}}(\cdot)$ since it does not know the utility function's form of the other participants. The

5. INTRA-TEAM STRATEGIES FOR NEGOTIATION TEAMS IN PREDICTABLE DOMAINS

two decisions that have to be taken during the negotiation are which offers are sent to the opponent, and whether or not the opponent's offer is accepted.

5.4.1 Offer proposal

Being a time-bounded negotiation, the representative employs a time-based concession strategy $s_{a_{re}}(\cdot)$ to negotiate with the opponent. It is based on a team deadline T_A and a concession speed β_A , which have been agreed upon prior to the negotiation start. This time-based tactic is formalized as (18, 57):

$$s_{a_{re}}(t) = 1 - (1 - RU_{a_{re}})\left(\frac{t}{T_A}\right)^{\frac{1}{\beta_A}} \quad (5.7)$$

The time-based strategy defines the aspiration level (utility demanded) by the agent at a specific round t . The utility is demanded from the point of view of the representative. Thus, any offer $X_{A \rightarrow op}^t$ proposed by a_{re} at round t will obey the following condition:

$$U_{a_{re}}(X_{A \rightarrow op}^t) \geq s_{a_{re}}(t) \quad (5.8)$$

This means that any offer sent to the opponent by the representative reports a utility for the representative which is greater than or equal to the level of demand marked by the time based concession tactic. Since there is a large number of offers (possibly infinite) that may obey the equation above, we aimed to satisfy the opponent's preferences as much as possible. The representative selects the offer that is the most similar to the previous offer received from the opponent using a similarity heuristic (56, 57) based on the Euclidean distance.

$$X_{A \rightarrow op}^t = \max_{X | U_{a_{re}}(X) = s_{a_{re}}(t)} Sim(X, X_{op \rightarrow A}^{t-1}) \quad (5.9)$$

5.4.2 Evaluation of Opponent's Offer

A common acceptance criterion is that an opponent's offer is accepted if it reports a utility which is higher than or equal to the utility that is to be demanded in the next negotiation step. In the case of the representative, it will accept the opponent's offer $X_{op \rightarrow A}^t$ at round t if it reports a utility $U_{a_{re}}(X_{op \rightarrow A}^t)$ greater than or equal to $s_{a_{re}}(t + 1)$. This can be formalized as follows:

$$ac_{a_{re}}(X_{op \rightarrow A}^t) = \begin{cases} \textit{accept} & \text{if } s_{a_{re}}(t + 1) \leq U_{a_{re}}(X_{op \rightarrow A}^t) \\ \textit{reject} & \text{otherwise} \end{cases} \quad (5.10)$$

5.4.3 Discussion

It is clear that since the representative negotiates according to its own utility function and reservation utility, it cannot guarantee any kind of unanimity regarding team decisions. Decisions taken by the representative may only be acceptable to himself, but nothing can be assured about the rest of team members. Next, we illustrate several examples with the example introduced in Section 5.3. In these examples, we assume that Alice is selected as representative.

In the first example, Alice has to propose an offer to the opponent at round $t = 30$ with a demand of $s_{a_{re}}(30) = 1 - (0.7)(\frac{30}{100}) = 0.79$. The last opponent offer has been $X_{op \rightarrow A}^{29} = (700, 150, 30, 0)$. One of the offers with $U_{a_{re}}(A \rightarrow op) \geq 0.79$ that minimizes the distance to $X_{op \rightarrow A}^{29}$ is $X_{A \rightarrow op}^{30} = (494.2, 150, 30, 0)$. This offers yields $U_{a_b}(X_{A \rightarrow op}^{30}) = 0.176$ and $U_{a_c}(X_{A \rightarrow op}^{30}) = 0.518$. If the offer were to be accepted, it would be acceptable for Alice and Charlie, but not for Bob. Thus, the team decision in this case would only guarantee $\frac{2}{3}$ unanimity.

In another round $t = 100$, Alice receives $X_{op \rightarrow A}^{100} = (406, 150, 0, 0)$ so that $U_{a_{re}}(X_{op \rightarrow A}^{100}) = 0.3 \geq s_{a_{re}}(101)$. Alice would accept the offer, but it does not reach the desired utility level for Bob $RU_{a_b} = 0.2 > U_{a_b}(X_{op \rightarrow A}^{100}) = 0.18$, and Charlie

5. INTRA-TEAM STRATEGIES FOR NEGOTIATION TEAMS IN PREDICTABLE DOMAINS

$RU_{a_c} = 0.3 > U_{a_c}(X_{op \rightarrow A}^{100}) = 0.21$. In this case, no one but the representative would get an acceptable deal from the negotiation.

Finally, in another negotiation at $t = 100$, Alice receives $X_{op \rightarrow A}^{100} = (406, 0, 0, 0)$ so that $U_{a_{re}}(X_{op \rightarrow A}^{100}) = 0.3 \geq s_{a_{re}}(101)$. Alice would accept the offer and it would be an acceptable offer also for Bob $RU_{a_b} = 0.2 < U_{a_b}(X_{op \rightarrow A}^{100}) = 0.78$, and Charlie $RU_{a_c} = 0.3 < U_{a_c}(X_{op \rightarrow A}^{100}) = 0.31$. Hence, the final offer would be acceptable for every team member.

One could think that if no consensus can be guaranteed, this strategy is not worth being used. However, when team members tend to be very similar this strategy is expected to yield acceptable results with communication costs equivalent to a bilateral negotiation process.

Another issue that has to be taken into account when using this strategy is security. In our example, we selected the representative randomly. Nevertheless, this should not be the way to proceed in a real application. Due to the fact that the representative makes all the decisions for the team, this strategy is highly prone to be manipulated by malicious agents. By malicious agents, we mean agents that supplant the identity of party members or agents that falsely allege a certain identity when their real identity is of an exact opposite nature (e.g., sellers posing as buyers). It is acknowledged that the representative should be an agent trusted among team members. For this matter, trust and reputation (38, 39) and social choice (131) mechanisms could be employed to determine the representative.

5.5 Similarity Simple Voting (SSV)

The second intra-team strategy relies on a trusted team mediator that helps team members to participate in the negotiation process. Its main tasks involve coordination of voting processes and communications with the opponent. It should

be highlighted that the team mediator communicates team’s decisions to the opponent, and broadcasts opponent’s decisions among team members. Thus, the fact that every team member participates in the negotiation process remains unknown for the opponent. As for intra-team communications, it should be noted that team members do not communicate among them, but they only communicate anonymously with the team mediator.

The decision rule used for voting processes is plurality/majority. More specifically, a plurality rule is used in the voting process employed to decide which offer is sent to the opponent, and a majority rule is used in the voting process employed to decide opponent’s offer acceptance. A detailed view of the intra-team strategy can be observed in Algorithm 4, which describes the whole process from the point of view of the mediator. Messages are represented as (*Body direction agents*). Therefore, (Accept $\rightarrow op$) means that the agent sends an accept message to *op*, whereas (Reject $\leftarrow op$) describes a message from *op* with the content “Reject”.

5.5.1 Offer proposal

Whenever a new offer has to be proposed to the opponent at round t , the mediator opens a call for proposals among team members. Each team member a_i is allowed to communicate anonymously one offer $X_{a_i \rightarrow A}^t$ to be proposed to the opponent. Once every proposal has been gathered, the mediator opens a voting process where offers proposed $XT^t = \bigcup_{i=1}^M X_{a_i \rightarrow A}^t$ are made public among team members. Then, each agent a_i anonymously sends a multi-vote $Vote_{a_i}$ to the mediator. A multi-vote has votes for every offer in XT^t . We use the notation $Vote_{a_i}(j)$ to denote the vote given by agent a_i to the offer j -th from XT^t . The votes can be either positive (1), if the offer j -th is acceptable for a_i at round t , or negative (0), if the offer j -th is not acceptable for a_i at round t . Once all votes have been gathered, the mediator sums up the number of positive votes and the most supported offer

5. INTRA-TEAM STRATEGIES FOR NEGOTIATION TEAMS IN PREDICTABLE DOMAINS

$X_{A \rightarrow op}^t$ is selected, made public among team members, and sent to the opponent. When a tie is produced, the tie-breaker rule consists in randomly selecting one of the most supported offers. The following Equation describes the selection rule of the previous mechanism:

$$X_{A \rightarrow op}^t = \operatorname{argmax}_{X_j \in XT^t} \sum_{a_i \in A} \text{Vote}_{a_i}(j) \quad (5.11)$$

We assume that, since the negotiation is time-bounded, team members follow a time-based concession strategy where the concession speed β_A is common and agreed by teammates prior to the negotiation process:

$$s_{a_i}(t) = 1 - (1 - RU_{a_i}) \left(\frac{t}{T_A} \right)^{\frac{1}{\beta_A}} \quad (5.12)$$

For proposing an offer to team members, the member a_i proposes an offer $X_{a_i \rightarrow A}^t$ from the iso-utility curve defined by $s_{a_i}(t)$. Since there may be more than a single offer with such utility, the agent has to choose one of the multiple offers. If the agent a_i wants its offer $X_{a_i \rightarrow A}^t$ to be accepted it should maximize the probability of being the most supported proposal by team members and the probability of being accepted by the opponent:

$$X_{a_i \rightarrow A}^t = \operatorname{argmax}_{X | U_{a_i}(X) = s_{a_i}(t)} p_{op}(X) \times p_A(X) \quad (5.13)$$

where $p_{op}(X)$ is the probability for X to be accepted by the opponent, and $p_A(X)$ is the probability for X to be selected by team members. One way to approximate these probabilities, which can be very costly to calculate, is by means of similarity heuristics. We incorporated agents with a similarity heuristic based on the Euclidean distance. It takes into account the last offer proposed by the opponent $X_{op \rightarrow A}^{t-1}$ and the offer sent by team members in the previous negotiation round $X_{A \rightarrow op}^{t-1}$. The most similar an offer is to $X_{op \rightarrow A}^{t-1}$, the more probabilities for the offer to be accepted by the opponent. Analogously, the most similar an offer

5.5 Similarity Simple Voting (SSV)

is to $X_{A \rightarrow op}^{t-1}$, the more probabilities for the offer to be the most supported option in the voting process and, therefore, to be sent to the opponent. Thus, Equation 5.13 can be approximated by similarity heuristics as follows:

$$X_{a_i \rightarrow A}^t = \underset{X|U_{a_i}(X)=s_{a_i}(t)}{\operatorname{argmax}} p_{op}(X) \times p_A(X) \approx \underset{X|U_{a_i}(X)=s_{a_i}(t)}{\operatorname{argmax}} \operatorname{Sim}(X, X_{op \rightarrow A}^{t-1}) \times \operatorname{Sim}(X, X_{A \rightarrow op}^{t-1}) \quad (5.14)$$

Finally, for determining the acceptability of offers proposed by team members at round t , we use a rational criterion so that an agent a_i emits a positive vote $\operatorname{Vote}_{a_i}(j) = 1$ for the j -th offer if it reports a utility that is greater than or equal to the utility marked by the concession strategy $s_{a_i}(t)$. Otherwise, the offer is not supported and a negative vote is emitted. This process can be formalized as:

$$\operatorname{Vote}_{a_i}(j) = \begin{cases} 1 & \text{if } U_{a_i}(XT^t(j)) \geq s_{a_i}(t) \\ 0 & \text{otherwise} \end{cases} \quad (5.15)$$

where $XT^t(j)$ represents the j -th offer in XT^t .

5.5.2 Evaluation of Opponent's Offer

Whenever the mediator receives an offer $X_{op \rightarrow A}^t$ from the opponent at round t , it broadcasts the offer among team members. Then, the mediator opens up a majority voting process where each agent a_i states whether or not the opponent's offer is acceptable $ac_{a_i}(X_{op \rightarrow A}^t)$. The mediator counts the number of acceptances, and if the offers is supported by the majority ($> \frac{|A|}{2}$) then it is accepted by the team. Otherwise, the offer is rejected. If the number of team members is even and a tie has been produced, a random decision is taken by the mediator. This mechanism can be described as follows:

$$ac_A(X_{op \rightarrow A}^t) = \begin{cases} \text{accept} & \text{if } \sum_{a_i \in A} ac_{a_i}(X_{op \rightarrow A}^t) > \frac{|A|}{2} \\ \text{reject} & \text{if } \sum_{a_i \in A} ac_{a_i}(X_{op \rightarrow A}^t) < \frac{|A|}{2} \\ \text{random} & \text{otherwise} \end{cases} \quad (5.16)$$

5. INTRA-TEAM STRATEGIES FOR NEGOTIATION TEAMS IN PREDICTABLE DOMAINS

How team members a_i decide the acceptability of the opponent's offer $ac_{a_i}(X_{op \rightarrow A}^t)$ follows the rational mechanism that we have employed so far. Basically, the offer is acceptable if it yields a utility which is greater than or equal to the utility demanded by the concession strategy in the next negotiation round $s_{a_i}(t+1)$. Otherwise, the offer is not considered acceptable. The following Equation formalizes the acceptance criterion:

$$ac_{a_i}(X_{op \rightarrow A}^t) = \begin{cases} 1 & \text{if } U_{a_i}(X_{op \rightarrow A}^t) \geq s_{a_i}(t+1) \\ 0 & \text{otherwise} \end{cases} \quad (5.17)$$

5.5.3 Discussion

The proposed method is capable of guaranteeing team decisions that are supported by a plurality/majority of the participants. More specifically, plurality is assured in the case of the offer proposed to the opponent, and majority is assured when deciding opponent's offer acceptance. Exceptions for this minimum level of team consensus are ties. For instance, the most extreme case is present when team members propose offers to the team, but they only support their own offers. In that case, each proposal sums up exactly 1 positive vote and there is not a clear plurality winner. In our illustrative example, the trusted mediator has received (700, 120, 1.5, 8) from the opponent at $t = 89$, and the last team offer was (602, 142.5, 3, 4). Now, $t = 90$, the mediator opens a new call for proposals. Alice demands a utility $s_{a_a}(90) = 0.37$, and the offer with such utility that maximizes Equation 5.14 is $X_{a_c}^{90} = (509.65, 133.04, 10.54, 5.68)^1$. Bob demands a utility $s_{a_b}(90) = 0.28$, and the offer with such utility that maximizes Equation 5.14 is $X_{a_b \rightarrow A}^{90} = (605.65, 99.63, 2.68, 6.52)$. Charlie demands a utility $s_{a_c}(90) = 0.37$, and the offer with such utility that maximizes Equation 5.14 is

¹Calculated using *sqp* (non-linear constrained optimization) in Octave 3.2.4

```

t = 0;
while t ≤ TA do
  Send (Call For Proposals → A);
  XTt = ∅;
  foreach ai ∈ A do
    Receive (Xtai→A ← ai);
    XTt = XTt ∪ Xtai→A;
  end
  Send (Open Voting XTt → A);
  foreach ai ∈ A do Receive (Voteai ← ai);
  XtA→op = argmaxXj ∈ XTt, ai ∈ A ∑ Voteai(j);
  Send (XtA→op → op, A);
  Receive (Xtop→A ← op);
  if Xtop→A = Withdraw then
    Send (Opponent Withdraw → A);
    Return Failure;
  end
  else if Xtop→A = Accept then
    Send (Offer Accepted → A);
    Return Success;
  end
  else
    Send (Open Voting Xtop→A → A);
    foreach ai ∈ A do Receive (acai(Xtop→A) ← ai);
    if acA(Xtop→A) = accept then
      Send (Accept → op, A);
      Return Success;
    end
    else
      Send (Opponent Offer Rejected → A);
    end
  end
  end
  t = t + 1;
end
Send (Withdraw → op, A);
Return Failure;

```

Algorithm 4: Pseudo-code algorithm for the mediator in the Similarity Simple Voting intra-team strategy.

5. INTRA-TEAM STRATEGIES FOR NEGOTIATION TEAMS IN PREDICTABLE DOMAINS

$X_{a_c \rightarrow A}^{90} = (511.82, 119.80, 10.66, 9.10)$. Since $U_{a_a}(X_{a_b \rightarrow A}^{90}) = 0.14 < s_{a_a}(90)$ and $U_{a_a}(X_{a_c \rightarrow A}^{90}) = 0.36 < s_{a_a}(90)$, Alice only votes positively for her own offer. Similarly, since $U_{a_b}(X_{a_a \rightarrow A}^{90}) = 0.21 < s_{a_b}(90)$ and $U_{a_b}(X_{a_c \rightarrow A}^{90}) = 0.27$, Bob only votes positively for his own offer. Finally, since $U_{a_c}(X_{a_a \rightarrow A}^{90}) = 0.32 < s_{a_c}(90)$ and $U_{a_c}(X_{a_b \rightarrow A}^{90}) = 0.19 < s_{a_c}(90)$, Charlie's only positive vote is for his own offer. A tie is produced between all of the proposals, and so the mediator has to select one randomly.

If there is a tie in the voting process that determines whether or not the opponent's offer is accepted, it means that the number of team members is even and accepting the opponent's offer has received half of the votes. The tie-breaking rule assures that, in this case, at least half of the team members are satisfied with the agreement at round t , and half of the members plus one are satisfied with the agreement in case that the number of team members is odd. For instance, at round $t = 60$, the mediator receives $X_{op \rightarrow A}^{60} = (406, 150, 20, 20)$ and makes it public among team members. In the voting process, Alice supports the opponent's offer since $U_{a_a}(X_{op \rightarrow A}^{60}) = 0.63 > s_{a_a}(61) = 0.57$, Bob does not support the opponent's offer since $U_{a_b}(X_{op \rightarrow A}^{60}) = 0.27 < s_{a_b}(61) = 0.51$, and Charlie supports the opponent's offer since $U_{a_c}(X_{op \rightarrow C}^{60}) = 0.64 > s_{a_c}(61) = 0.57$. The opponent's offer would be accepted and the negotiation would end with a final agreement since 2 out of 3 members supported the opponent's decision at round 60.

However, it should be noted that even if a team member a_i does not support a team decision at round t and a final agreement is found, it does not necessarily mean that the final agreement does not satisfy a_i 's aspiration (utility of the final agreement lower than its reservation utility). That situation is only mandatory when $t = T_A$, a final agreement has been found, and a_i did not support such decision. In other scenarios, the final agreement, even if not supported by every

5.6 Similarity Borda Voting (SBV)

team member, may or may not achieve utility levels below team members' aspirations. For example, in the second example in this discussion, even if Bob does not support the opponent offer at round 60, it still yields a utility which is higher than his reservation utility and the final offer would be acceptable in the end. In the first scenario, it depends on which offer is selected by the tie-breaking rule: If Alice's or Charlie's offer are selected and accepted by the opponent, then every team members achieves values over their reservation utilities. If Bob's offer is selected and accepted by the opponent, then neither Alice and Charlie are satisfied with the final agreement.

Regarding security, against any kind of malicious agent that infiltrates the team, SSV is more robust than RE due to the fact that a larger number of malicious agents may be needed to manipulate the team. In the case where team members decide on whether or not to accept the opponents' offer, the set of malicious agents has to be equal to half the number of team members plus one if they want to assure that the team is manipulated. In the case of the plurality voting carried out to decide on which offer is sent to the opponent, in the best case only two malicious agents need to infiltrate the team but a large number of team members (i.e., a majority) may be needed if manipulation wants to be assured.

5.6 Similarity Borda Voting (SBV)

SSV is capable of assuring majority and plurality decisions within the team. However, some scenarios may need of intra-team strategies that ensure higher levels of consensus. SBV and FUM (described later) are designed to solve this problem. The basic structure of SBV remains the same than in SSV, but the voting rules employed are different. More specifically, when each team member votes team proposals, borda count is employed to determine the winner, and a unanimity rule

5. INTRA-TEAM STRATEGIES FOR NEGOTIATION TEAMS IN PREDICTABLE DOMAINS

is used to determine opponent's offer acceptance. Next, we briefly describe the aspects which make SBV different to SSV.

5.6.1 Offer proposal

As in SSV, when the team has to propose an offer to the opponent, the mediator opens a call for proposals where each team member can propose an offer to the mediator. Then, once every offer has been gathered, the mediator makes them public to the team members and a voting process starts. The main difference between both intra-team strategies resides in the fact that team members vote according to a Borda count rule (131). Basically, each team member a_i ranks the proposals XT^t in ascending order according to its own utility function $U_{a_i}(\cdot)$. We denote as $rank_{a_i}(XT^t)$ the ascending rank according to a_i 's utility function, and $Position(X, rank_{a_i}(XT^t))$ as the position (1 to $|XT^t|$) that the offer X occupies in $rank_{a_i}(XT^t)$. The vote emitted by a_i for offer j -th in XT^t is the position occupied by such offer in the ranked list minus one unit:

$$Vote_{a_i}(j) = Position(XT^t(j), rank_{a_i}(XT^t)) - 1 \quad (5.18)$$

Numerical votes for each offer are summed up by the mediator, who finally selects the offer that received the highest sum of scores from the team members (see Equation 5.11). It should be highlighted that the similarity heuristic employed by team members is the same than the one employed in SSV.

5.6.2 Offer acceptance

As for the offer acceptance, the only difference resides in the rule used by the mediator. The opponent's offer is accepted only if it is acceptable for all the team members.

5.6.3 Discussion

When describing the minimum unanimity level guaranteed by SBV, we mentioned the term semi-unanimity. It is clear that if an opponent offer is accepted by the team, it is acceptable for every team member due to the unanimity rule employed. Thus, Equation 5.3 is guaranteed if the final agreement is an offer accepted from the opponent. However, unanimity is not guaranteed regarding the team offer sent to the opponent. Borda count is generally referred as a method that selects broadly accepted options as winners instead of the majority/plurality option (e.g., avoid the tyranny of the majority). In this sense, Borda count entails some degree of unanimity. Nevertheless, the specific degree of unanimity that Borda assures is difficult to determine in our negotiation scenario.

Some problematic situations that arose in SSV are solved with this sort of voting. If we recall the extreme tie case in SSV's discussion, Alice proposed $X_{a_c}^{90} = (509.65, 133.04, 10.54, 5.68)$, Bob proposed $X_{a_b \rightarrow A}^{90} = (605.65, 99.63, 2.68, 6.52)$, and Charlie proposed $X_{a_c \rightarrow A}^{90} = (511.82, 119.80, 10.66, 9.10)$. The preference rankings of each team member would be $U_{a_a}(X_{a_b \rightarrow A}^{90}) < U_{a_a}(X_{a_c \rightarrow A}^{90}) < U_{a_a}(X_{a_a \rightarrow A}^{90})$ for Alice, $U_{a_b}(X_{a_a \rightarrow A}^{90}) < U_{a_b}(X_{a_c \rightarrow A}^{90}) < U_{a_b}(X_{a_b \rightarrow A}^{90})$ for Bob, and $U_{a_c}(X_{a_b \rightarrow A}^{90}) < U_{a_c}(X_{a_a \rightarrow A}^{90}) < U_{a_c}(X_{a_c \rightarrow A}^{90})$ for Charlie. Alice's offer gets 2 points from Alice, 0 from Bob, and 1 from Charlie, Bob's offer gets just 2 points from Bob, and Charlie's offer gets 2 points from Charlie, 1 point from Alice, and 1 point from Bob. Hence, Charlie's offer is selected to be sent to the opponent, and if accepted, it is one of the proposed offers that was over every team member's reservation utility, making it a unanimously acceptable deal.

We may find scenarios where Borda count, despite generally selecting broadly accepted options, selects winners that are not supported/acceptable by every team member. For example, if $t = 100$, and the last offer received from the opponent is

5. INTRA-TEAM STRATEGIES FOR NEGOTIATION TEAMS IN PREDICTABLE DOMAINS

$X_{op \rightarrow A}^{99} = (700, 120, 1.5, 8)$. Alice would propose $X_{a_a \rightarrow A}^{100} = (542.63, 133.84, 8.36, 5.53)$, Bob would propose $X_{a_b \rightarrow A}^{100} = (636.37, 114.79, 2.35, 6.52)$, and Charlie would propose $X_{a_c \rightarrow A}^{100} = (551.44, 123.68, 8.18, 8.09)$. In the Borda voting process, each team member ranking would be $U_{a_a}(X_{a_b \rightarrow A}^{100}) = 0.10 < U_{a_a}(X_{a_c \rightarrow A}^{100}) = 0.28 < U_{a_a}(X_{a_a \rightarrow A}^{100})$ for Alice, $U_{a_b}(X_{a_a \rightarrow A}^{100}) = 0.18 < U_{a_b}(X_{a_b \rightarrow A}^{100}) = 0.20 < U_{a_b}(X_{a_c \rightarrow A}^{100}) = 0.23$ for Bob, $U_{a_c}(X_{a_a \rightarrow A}^{100}) = 0.16 < U_{a_c}(X_{a_b \rightarrow A}^{100}) = 0.27 < U_{a_c}(X_{a_c \rightarrow A}^{100}) = 0.30$ for Charlie. According to these rankings, Alice's offer would receive 2 points from Alice and 1 point from Charlie, Bob's offer would only get 1 point from Bob, and Charlie's offer would get 1 point from Alice, 2 points from Bob, and 2 points from Charlie. Therefore, Charlie's offer would be selected to be sent to the opponent. If it is accepted, the final deal would report utilities higher than reservation utilities for Charlie and Bob, but not for Alice.

Since Borda count generally selects broadly accepted candidates, a larger number of malicious agents may be necessary to manipulate the team compared to the number of malicious agents necessary to manipulate SSV. However, given the unanimity rule employed in the opponent's offer acceptance mechanism, it is impossible for opponents posing as team members to manipulate the team regarding the acceptability of the offer sent by the opponent. Since as long as one of the team members does not agree with the opponent's offer, the offer will be rejected, it is not possible for opponent agents to manipulate the team even if they infiltrate the team in large numbers.

5.7 Full Unanimity Mediated (FUM)

The last intra-team strategy, Full Unanimity Mediated (FUM), seeks to reach unanimity regarding all team decisions. In fact, every team decision taken (i.e., offer acceptance, offer proposal) following this intra-team strategy entails unanimity at

5.7 Full Unanimity Mediated (FUM)

each round t of the negotiation process. However, the type of mediator required for FUM is more sophisticated than mediators in the rest of intra-team strategies presented in this chapter. It requires for the mediator to participate in a pre-negotiation process where team members hand decision rights over issues that are not interesting for them, to be able to infer issues' importance for the opponent, to coordinate unanimity voting processes, and to coordinate an iterated building process that constructs the offers sent to the opponent. A complete view of the pseudo-algorithm carried out by the mediator can be observed in Algorithm 5.

5.7.1 Pre-negotiation: information sharing

During the pre-negotiation, team members are allowed to hand over decision rights over issues that they do not consider interesting. The iterated offer building process relies on a mechanism which sets issue values one-per-one according to team members' will. When an agent hands over decision rights on an issue, it does not participate in the setting of such issue. All the communications in the pre-negotiation are private with the mediator, who asks each team member regarding the set of issues which it is willing to hand over. The rationale behind the idea of handing over decision rights is that conflict may be reduced, and, so, the chances to build a more likeable offer for the opponent may be increased while maintaining a good quality for one's own utility function. The fact that some issues may yield little or no importance at all for some team members is also feasible in a team setting, since some of these issues may have been introduced to satisfy the interests of a subgroup of team members.

The pre-negotiation protocol goes as follows. First, the mediator opens a call for decision rights, where each team member a_i is allowed to send (to the mediator) a set of negotiation issues NI_{a_i} , whose decision rights are handed over by a_i . Once all the responses have been gathered, the mediator keeps track of those issues that

5. INTRA-TEAM STRATEGIES FOR NEGOTIATION TEAMS IN PREDICTABLE DOMAINS

```

Send (Ask for  $NI_{a_i} \rightarrow A$ );
foreach  $a_i \in A$  do Receive ( $NI_{a_i} \leftarrow a_i$ );
 $t = 0$ ;
while  $t \leq T_A$  do
   $agenda = build\_agenda()$ ;
   $A' = A$ ;  $X'_{A \rightarrow op} = \emptyset$ ;
  foreach  $j \in \bigcap_{i=1}^M NI_{a_i}$  do
     $x_j = maximize\_for\_opponent(j)$ ;  $X'_{A \rightarrow op} = X'_{A \rightarrow op} \cup x_j$ ;
  end
  foreach  $j \in agenda \wedge issue\_not\_set(j)$  do
    Send (Needed value  $j$ , given  $X'_{A \rightarrow op} \rightarrow a_i | j \notin NI_{a_i} \wedge a_i \in A'$ );
    Receive ( $x_{a_i, j} \leftarrow a_i | j \notin NI_{a_i} \wedge a_i \in A'$ );
    if monotonically_increasing( $j$ ) then  $x_j = \max_i x_{a_i, j}$ ;
    else  $x_j = \min_i x_{a_i, j}$ ;
     $X'_{A \rightarrow op} = X'_{A \rightarrow op} \cup x_j$ ;
    Send (Acceptable  $X'_{A \rightarrow op} ? \rightarrow a_i | a_i \in A'$ );
    foreach  $a_i \in A'$  do
      Receive ( $ac'_{a_i}(X'_{A \rightarrow op}) \leftarrow a_i$ );
      if  $ac'_{a_i}(X'_{A \rightarrow op}) = true$  then  $A' = A' - a_i$ ;
    end
    if  $A' = \emptyset$  then break;
  end
  foreach  $j \in agenda \wedge issue\_not\_set(j)$  do
     $x_j = maximize\_for\_opponent(j)$ ;  $X'_{A \rightarrow op} = X'_{A \rightarrow op} \cup x_j$ ;
  end
   $X^t_{A \rightarrow op} = X'_{A \rightarrow op}$ ; Send ( $X^t_{A \rightarrow op} \rightarrow op, A$ );
  Receive ( $X^t_{op \rightarrow A} \leftarrow op$ );
  if  $X^t_{op \rightarrow A} = Withdraw$  then
    Send (Opponent Withdraw  $\rightarrow A$ ); Return Failure;
  else if  $X^t_{op \rightarrow A} = Accept$  then
    Send (Offer Accepted  $\rightarrow A$ ); Return Success;
  else
    Send (Open Voting  $X^t_{op \rightarrow A} \rightarrow A$ );
    foreach  $a_i \in A$  do Receive ( $ac_{a_i}(X^t_{op \rightarrow A}) \leftarrow a_i$ );
    if  $ac_A(X^t_{op \rightarrow A}) = accept$  then
      Send (Accept  $\rightarrow op, A$ ); Return Success;
    else
      Send (Opponent Offer Rejected  $\rightarrow A$ );
    end
  end
   $t = t + 1$ ;
end
Send (Withdraw  $\rightarrow op, A$ ); Return Failure;

```

Algorithm 5: Pseudo-code algorithm for the mediator in the FUM intra-team strategy.

5.7 Full Unanimity Mediated (FUM)

are not interesting for each agent NI_{a_i} , and those issues that are not interesting for all team members $\bigcap_{i=1}^M NI_{a_i}$. Once this process has finished, the team and the mediator are ready to start the negotiation process.

Of course, the set of issues handed over by each team member is not controllable by the mediator. It depends on the behavior of each agent. In our model, the set of issues handed over by each agent depends on a private parameter ϵ_{a_i} . The value of such parameter is related to the weight of the different negotiation issues in one's own utility function. More precisely, if $\epsilon_{a_i} = 0$, then the agent is only willing to hand over the decision rights over those issues that are not interesting for himself (i.e., weight equal to zero in the utility function). When $\epsilon_{a_i} = 1$, the agent is willing to hand over decision rights over every issue in the negotiation. In general, the agent is willing to hand over decision rights over issues whose sum of weight in the utility function is equal to or lower than ϵ_{a_i} :

$$\sum_{j \in NI_{a_i}} w_{a_i, j} \leq \epsilon_{a_i} \quad (5.19)$$

Given a certain ϵ_{a_i} , a reasonable heuristic is to assume that the agent is willing to concede as many decision rights as possible since this will enhance the possibility of finding an agreement with the opponent. Hence, each team member a_i chooses the largest possible set NI_{a_i} that fulfills Eq. 5.19. A simple algorithm that solves this problem is ordering the negotiation issues in ascending order by weight in the utility function. The set NI_{a_i} starts empty, and, then, the array of ordered issues is followed. If the issue weight plus the weights of those issues already in NI_{a_i} exceeds ϵ_{a_i} , then the search stops. Otherwise, the issue is added to NI_{a_i} and the algorithm continues with the next issue.

Let us imagine that $\epsilon_{a_a} = 0$, $\epsilon_{a_b} = 0.1$, and $\epsilon_{a_c} = 0.1$. Alice's ranking of issues (from less important to more important) would be *db, cf, pd, pp*. The list of issues whom decision rights are handed over by Alice, NI_{a_a} , would start as an empty set.

5. INTRA-TEAM STRATEGIES FOR NEGOTIATION TEAMS IN PREDICTABLE DOMAINS

Then, Alice would start by looking at *db*. Since its weight in Alice's utility function is 0, the accumulated weights in NI_{a_i} (equal to 0) and $0+w_{db}$ does not exceed ϵ_{a_a} , Alice would add *db* to the set. The next issue that would be looked at by Alice is *cf*. As before, *cf* is added to NI_{a_a} . Then, Alice would look at *pd*. Due to the fact that *pd*'s weight in Alice's utility function plus the accumulated weights in NI_{a_a} exceed ϵ_{a_a} , the process would stop and *cf* would not be added to NI_{a_a} . As for Bob, its ranking of issues would be *db, pd, pp, cf*. Bob would start looking at *db*. Since its weight in Bob's utility function does not exceed ϵ_{a_b} , *db* is added to NI_{a_b} . The next issue would be *pd*, and it would be added to NI_{a_b} since its weight plus *db*'s weight exactly match ϵ_{a_b} . No more issues would be added to NI_{a_b} . Similarly, Charlie would only hand over decision rights for *cf* ($NI_{a_c} = \{cf\}$).

5.7.2 Negotiation: Offer proposal

In order to determine which offer is sent to opponent, the mediator governs an iterated building process. The aim of this iterated process is building an offer issue per issue according to team members' needs so that the offer sent to the opponent is acceptable for every team member. The order in which the issues are adjusted is determined by an agenda built by the mediator. The details of this agenda are discussed in Section 5.7.3. Briefly, the iterated building process goes as follows:

1. The agenda of issues *agenda* is built by the mediator according to the available information.
2. When the iterated process starts, every team member is considered an active member ($a_i \in A'$) in the construction process.
3. The initial partial offer $X'_{A \rightarrow op}$ starts as an offer whose issues have not been set.

5.7 Full Unanimity Mediated (FUM)

4. The mediator checks those issues that are not interesting for every team member $\bigcap_{i=1}^M NI_{a_i}$. These issues are maximized according to the opponent's preferences. For instance, if the price is one of these issues and the opponent is a seller, it would be maximized for the opponent, thus, acquiring its maximum value. The partial offer $X_{A \rightarrow op}^t$ is updated with the new issue values.
5. The next issue j in the agenda is selected. Those team members active in the construction process ($a_i \in A'$) and interested in j ($j \notin NI_{a_i}$) are asked by the mediator to submit the value $x_{a_i,j}$ needed of issue j to get as close as possible to their aspiration levels.
6. The values $x_{a_i,j}$ gathered from team members are aggregated. The best value is selected according to the ranking of issue values (issues are *predictable and compatible*). For instance, if the assumed valuation function is monotonically increasing, then the *max* operator is used to aggregate the values and obtain the final value for the issue x_j .
7. x_j is set in $X_{A \rightarrow op}^t$ and the new partial offer is broadcasted among team members. Every team member that is active in the construction phase is asked if the current partial offer already satisfies its current demands.
8. Every response is gathered by the mediator. Those agents that answered positively are removed from the list of active agents. If there are still active agents, the mediator goes back to 5.
9. When every team member has been satisfied by the partial offer $X_{A \rightarrow op}^t$, if there are still issues that have not been set, those issues are maximized according to the opponent's preferences. Then, a final offer $X_{A \rightarrow op}^t$ is obtained, made public among team members, and sent to the opponent.

5. INTRA-TEAM STRATEGIES FOR NEGOTIATION TEAMS IN PREDICTABLE DOMAINS

In the protocol described above, team members are asked to submit a value for issues in which they are interested and to determine whether or not the partial offer satisfies their needs. In both cases, and as in previous strategies, we have assumed that team members follow time-based concession strategies similar to the one described in Equation 5.12. However, since team members may have handed over some issue decision rights, it is not possible for agents to demand the maximum utility. The value ϵ_{a_i} has to be subtracted from the maximum utility. Therefore, the concession strategy $s_{a_i}(t)$, which determines the level of demand at each negotiation round, can be formalized as:

$$s_{a_i}(t) = (1 - \epsilon_{a_i}) - (1 - \epsilon_{a_i} - RU_{a_i}) \left(\frac{t}{T_A} \right)^{\frac{1}{\beta_A}} \quad (5.20)$$

When team members are asked about a value for j , each team member communicates anonymously the value $x_{a_i,j}$. The value communicated is the one that gets as close as possible to its desired aspiration level $s_{a_i}(t)$ at round t . Taking the linear additive utility function formula, this can be calculated as:

$$x_{a_i,j} = \operatorname{argmin}_{x \in D_j} (s_{a_i}(t) - U_{a_i}(X_{A \rightarrow op}^t) - w_{a_i,j} V_{a_i,j}(x)) \quad (5.21)$$

where $s_{a_i}(t)$ is the utility demanded by the agent a_i at round t , $U_{a_i}(X_{A \rightarrow op}^t)$ is the utility reported by the current partial offer, and $w_{a_i,j} V_{a_i,j}(x)$ is the weighted utility reported by the value demanded by the agent. Since the value demanded looks to be as close as possible to the utility necessary to get to the current aspiration, the function is minimized. However, the following constrain is fulfilled by team members in order to avoid surpassing the utility demanded:

$$s_{a_i}(t) - U_{a_i}(X_{A \rightarrow op}^t) - w_{a_i,j} V_{a_i,j}(x_{a_i,j}) \geq 0 \quad (5.22)$$

As for determining when a partial offer is acceptable, team members follow a similar criterion to the method proposed in other intra-team strategies. Basically,

a partial offer is acceptable for an agent a_i if it reports a utility that is greater than or equal to the aspiration level marked by its concession strategy:

$$ac'_{a_i}(X'_{A \rightarrow op}) = \begin{cases} true & \text{if } U_{a_i}(X'_{A \rightarrow op}) \geq s_{a_i}(t) \\ false & \text{otherwise} \end{cases} \quad (5.23)$$

where *true* indicates that the partial offer is acceptable at its current state for agent a_i , and *false* indicates the opposite.

5.7.3 Negotiation: observing opponent's concessions and building an issue agenda

Once the negotiation starts, the mediator attempts to guess a ranking of issues according to the opponent's preferences. This ranking is used to build the agenda of issues used by the team in the iterated offer building process. The idea behind the agenda is attempting to satisfy team members as much as possible with those issues that are less important for the opponent. This way, team members may reach their desired aspiration level with those issues less interesting for the opponent, and use the rest of issues to make the offer as satisfactory as possible for the opponent. The only information available for the mediator regarding the opponent's preferences are the offers received. Thus, the mediator has to infer a ranking of issues according to that information. A possible heuristic is assuming that agents usually concede less in important issues and greater concessions are performed in lesser important issues at the first rounds of the negotiation.

Our proposed heuristic assumes that the mediator observes opponent's offers for the first k interactions. Then, it calculates the concession performed in each issue. Since our practical model assumed that the opponent's utility function employs the opposite type of valuation function than team members for each issue, it is feasible to calculate the amount of concession performed at each issue. For instance, if the opponent is a seller, it is reasonable to assume that its valuation

5. INTRA-TEAM STRATEGIES FOR NEGOTIATION TEAMS IN PREDICTABLE DOMAINS

functions is monotonically increasing (e.g., higher prices report higher utilities) and, thus, any value below the maximum price can be considered a concession with respect to the maximum price. Therefore, the amount of concession can be calculated in each issue. For each issue j , we may calculate the total amount of relative concession C_j in the first k offers:

$$C_j = \sum_{t=0}^{k-1} \frac{|X(j)_{op \rightarrow A}^t - best_value(j)|}{max_value(j) - min_value(j)} \quad (5.24)$$

where $X(j)_{op \rightarrow A}^t$ is the value of issue j in the offer $X_{op \rightarrow A}^t$, $best_value(j)$ is the best possible value for the opponent in issue j , and $max_value(j)$ and $min_value(j)$ are the maximum and minimum value of the issue in the negotiation domain. The inner part of the summatory determines the concession on issue j in the offer received at interaction/round t with respect to the best issue value for the opponent. So, the summatory counts the total concession for issue j in the first k offers. The heuristic is that issues that score lower in Equation 5.24 are usually those more important for the opponent, whilst those issues scoring higher in Equation 5.24 are those less important for the opponent. Based on the available information (i.e., number of rounds up to k), the mediator builds an agenda of issues according to the scores of C_j in descending order. This way, lesser important issues for the opponent are first in the agenda.

5.7.4 Negotiation: Offer acceptance

Since this intra-team strategy looks for unanimity regarding team decisions, we employed the same mechanism employed in SBV for determining whether or not an opponent offer is acceptable. When the mediator receives the opponent's offer $X_{op \rightarrow A}^t$, the offer is publicly announced to all of the team members. Then, the mediator opens a private voting process where each team member a_i should specify

5.7 Full Unanimity Mediated (FUM)

whether or not it supports acceptance of the opponent's offer $ac_{a_i}(X_{op \rightarrow A}^t)$. The mediator counts the number of positive votes. The offer is accepted if the number of positive votes is equal to the number of team members. Otherwise, the offer is rejected.

Similarly to SBV, an opponent offer is acceptable for a team member at round t if it reports a utility that is greater than or equal to the aspiration level marked by the concession strategy in the next round:

$$ac_{a_i}(X_{op \rightarrow A}^t) = \begin{cases} true & \text{if } s_{a_i}(t+1) \leq U_{a_i}(X_{op \rightarrow A}^t) \\ false & \text{otherwise} \end{cases} \quad (5.25)$$

where *true* means that the agent supports the opponent's offer, *false* has the opposite meaning, and $s_{a_i}(\cdot)$ is the concession strategy employed by agent a_i to calculate the aspiration level at each negotiation round t .

5.7.5 Discussion

As mentioned in this chapter, FUM allows team members to reach unanimity regarding team's decisions. These decisions include the offer that is sent to the opponent and the acceptance/rejection of the opponent's offers. In the latter, it is clear that according to Equations 5.25 and 5.3, and the proposed acceptance mechanism, unanimity is assured since an opponent offer is only accepted when it is equal or greater than each team members' demands. In the former process, how unanimity is achieved is not straightforward.

The type of unanimity that can be guaranteed by FUM regarding the offers sent to the opponent is a strict unanimity. We define that an offer sent to the opponent $X_{A \rightarrow op}^t$ is a **strict unanimous decision** for the team when, for any team member a_i , the offer reports a utility that is greater than or equal to its

5. INTRA-TEAM STRATEGIES FOR NEGOTIATION TEAMS IN PREDICTABLE DOMAINS

current aspiration level $s_{a_i}(t)$:

$$\forall_{a_i \in A} U_{a_i}(X_{A \rightarrow op}^t) \geq s_{a_i}(t) \quad (5.26)$$

Achieving this definition of unanimity within the team ensures that if a final agreement is found, it reports a utility that is greater than or equal to each agent's private reservation utility, thus fulfilling Equation 5.3. But the intra-team strategy goes beyond that, since it is capable of satisfying team member's demands even when they are above the reservation utility. In fact, the definition of team unanimity in Equation 5.3 is included in strict unanimity since any offer sent by the team is equal to or greater than each team members' reservation utility as long as team members do not ask for less than their reservation utilities. In order to achieve the proposed definition of unanimous decision, some assumptions have to be made regarding the behavior of team members. Basically, team members have to be truthful in their responses to the mediator, following the behavior specified in Eq. 5.21, 5.22 and 5.23. Next, we prove that, if team members follow these behaviors, unanimity is achieved in team's decisions according to Equation 5.26.

Proof. $\forall_{a_i \in A} U_{a_i}(X_{A \rightarrow op}^t) \geq s_{a_i}(t)$

subject to: Eq. 5.21, Eq. 5.22, Eq. 5.23, and *compatible and predictable* issues. For the sake of simplicity, we assume that team members' valuation functions are monotonically increasing for any negotiation issue. It should be pointed out that, in that case, the aggregation operation carried out by the trusted mediator is the *max* operator. In any case, for any issue j , its value will be determined as $x_j = \max(x_{a_1,j}, x_{a_2,j}, \dots, x_{a_M,j})$ and then it holds true that $\forall_{a_i \in A}, w_{a_i,j} V_{a_i,j}(x_j) \geq w_{a_i,j} V_{a_i,j}(x_{a_i,j})$. The proof is quite straightforward. When the mediator declares that an issue j must be set, three different situations may arise for an agent a_i :

- a_i has already reached its aspiration level with the partial offer $U_{a_i}(X_{A \rightarrow op}^t) \geq s_{a_i}(t)$. Therefore, the value determined for x_j will add utility to the partial offer and the utility reported to a_i will further exceed its aspirations $U_{a_i}(X_{A \rightarrow op}^t) + w_{a_i,j} V_{a_i,j}(x_j) \geq s_{a_i}(t)$.

5.7 Full Unanimity Mediated (FUM)

- a_i can reach its current aspiration level $s_{a_i}(t)$ if it asks for a value $x_{a_i,j}$. Thus, $U_{a_i}(X'_{A \rightarrow op}{}^t) + w_{a_i,j}V_{i,j}(x_{a_i,j}) = s_{a_i}(t)$. Since the aggregation operation is $x_j = \max(x_{a_1,j}, x_{a_2,j}, \dots, x_{a_M,j})$, the new partial offer will have a utility that is equal to or greater than its aspirations, $U_{a_i}(X'_{A \rightarrow op}{}^t) + w_{a_i,j}V_{a_i,j}(x_j) \geq s_{a_i}(t)$.
- a_i cannot reach its aspirations by just setting x_j . In this case, a_i will demand the maximum possible value for j and then $x_j = x_{a_i,j}$. a_i will have to reach its aspiration level by adjusting the next issues in the agenda. In the worst case scenario, the next issue to be set x_N is the last one in the agenda. This means that a_i has demanded the maximum value for the previous issues and succeeded in getting its desired value for them. Thus, before the last issue is set, the utility reported by the partial offer to a_i is $\sum_j^{N-1} w_{a_i,j}$. Since $\sum_j^N w_{a_i,j} = 1$ and $0 \leq s_{a_i}(t) \leq 1$, the agent will reach its aspiration level by demanding a value for x_N that fulfills $V_{a_i,N}(x_N) \geq \frac{s_{a_i}(t) - \sum_j^{N-1} w_{a_i,j}}{w_{a_i,N}}$, which is ensured thanks to the morphology of the valuation functions ($0 \leq V_{a_i,j}(x) \leq 1$).

□

One might wonder whether or not it is reasonable to think that agents are truthful in this process. Members are not tempted to demand lesser value for issues since the process would not ensure that the final agreement would achieve its current aspiration level and, thus, reservation utility. On the other hand, it is true that agents may be inclined to demand a greater value for issues since the process ensures that the offer will be more profitable for them. Nevertheless, it should be pointed out that, generally, if more value is demanded for issues the offer may be less profitable for the opponent and the probabilities of reaching an agreement may be greatly reduced. This issue is studied in Subsection 5.8.3, where

5. INTRA-TEAM STRATEGIES FOR NEGOTIATION TEAMS IN PREDICTABLE DOMAINS

we analyze whether or not team members have strong incentives to deviate from the proposed behavior.

Another issue that needs to be considered is security with respect to malicious agents. The proposed intra-team strategy is robust against manipulations carried out by opponents' agents. However, it is vulnerable to malicious agents from competitor agents. Next we describe both types of manipulations and analyze why FUM is robust or vulnerable against such types of manipulation.

By opponent manipulation we refer to agents that infiltrate the team in order to increase the quality of the final agreement from the point of view of the opponent party. In a negotiation team setting formed by buyers, we are concerned about the fact that some seller parties may attempt to introduce agents among team members. This way, opponents may be able to maximize their own preferences by manipulating the decisions taken by the team. However, our proposed negotiation model is robust to this kind of manipulation. Let us imagine a situation where a negotiation team wants to buy a product and a seller has been able to infiltrate agents in the team. Due to the mechanism employed to build the offer sent to the opponent, and the mechanism employed to decide upon whether or not to accept the opponent's offer, it is not possible for opponent agents to manipulate the decisions taken within the team. Regarding the iterated offer construction process, an opponent agent would try to demand values that are close to the preferences of the opponent. In a generic electronic commerce application, an opponent agent might demand high values for the price and the dispatch date and low values for the product quality. However, the aggregation rules employed by the trusted mediator will ensure that team preferences prevail independently of the number of infiltrated opponent agents (the best value is chosen). As for the unanimous voting process, opponent agents might try to engage team members in accepting the opponent's offer. However, this is not possible due to the fact that

5.7 Full Unanimity Mediated (FUM)

as long as one team member does not support the opponent's offer, it will not be accepted. This is the case even in situations where the group of opponent agents is larger than the number of real team members.

Another kind of possible manipulation is the one carried out by competitor agents. Competitors are buyer agents (in the case that the team is made up of buyer agents) that are interested in the same product as the team. Some competitors may be interested in sabotaging team deals if that assures that competitors get better deals from the opponent. This is especially true in environments where goods or services are limited (e.g., personal sellers on Ebay). Thus, competitor agents may attempt to prevent the team from reaching an agreement with the opponent.

Even though the proposed model is robust against opponent agents, robustness is not maintained when dealing with infiltrated competitor agents. In that case, the strengths shown by the model become its weaknesses. In the voting process carried out to decide upon whether or not to accept the opponent's offer, only a single agent is needed to manipulate the process and prevent the team from accepting the opponent's offers. On the other hand, competitor agents may manipulate the offer construction phase by being highly demanding. In a generic electronic commerce application, the competitor agent would demand very low values for the price, short dispatch dates and very high product quality. This way, competitor agents make offers extremely undesirable for opponent agents, preventing the team from reaching a final agreement with the opponent. Due to the aggregation operators employed by the trusted mediator, only one competitor agent is needed to manipulate the offer construction process. Thus, this model should be employed only when team members are extremely sure that no competitor agent has infiltrated the team.

5.8 Empirical Analysis of Full Unanimity Mediated Intra-Team Strategy

As stated, one of the main goals of our work is obtaining intra-team strategies that are able to guarantee unanimity regarding team decisions. For that reason, we decided to explore Full Unanimity Mediated in depth. More specifically, we study the importance of the agenda of issues imposed by the mediator on the negotiation process, the impact of the number of decision rights that are handed over during the pre-negotiation, and whether or not team members have incentives to deviate from the proposed strategy.

5.8.1 Studying the Impact of Intra-Team Agenda

The mediator uses an agenda to determine which issues are set first in the iterated building process. A reasonable heuristic is to try to satisfy team members with those issues that are less important for the opponent. Otherwise, the resultant offer may be too demanding and the negotiation process may end in failure. Thus, ideally, the agenda should order the issues in ascendant order of importance for the opponent.

In our first experiment, we decided to study the importance of the agenda on the negotiation process. While every team member gets a utility that is greater than or equal to its desired aspiration level, the offer may be more or less demanding for the opponent. If the offer is less demanding for the opponent, it is more probable that it will be accepted by him. Therefore, we decided to study the utility reported by the teams' offer to the opponent at each negotiation round. We simulated a negotiation process where offers are not accepted (i.e., it always reaches the negotiation deadline) just to observe the utility of the offers proposed by the team from the opponent's perspective. Two different environments were

5.8 Empirical Analysis of Full Unanimity Mediated Intra-Team Strategy

tested: one with a short deadline $T_{op} = T_A = 10$, and one with a long deadline $T_{op} = T_A = 50$. Other parameters were set to the standard values of our negotiation model: $\beta_{op} = \beta_A = 1$, $\epsilon_{a_i} = 0$, and $RU_{op} = RU_{a_i} = 0$. Three different types of agendas for the FUM model were compared: a perfect agenda where the mediator knows perfectly the order of importance given by the opponent (FUM-perfect); the agenda built with the learning method described in this chapter (FUM-simple); and a random agenda that is built at each negotiation round (FUM-random). For FUM-simple, the number of initial negotiation rounds to be taken into account was set to $k = \lfloor \frac{T_A}{4} \rfloor$, which we found to be a good heuristic in practice. Additionally, the proposed negotiation model is compared for illustrative purposes with RE and SSV. These two models are expected to be less demanding in terms of utility due to the fact that less conflict is introduced with the opponent (i.e., a fewer number of team members may reach their aspiration level). A total of 100 random teams with size $M = 4$ and random utility functions ($N=4$ issues) were confronted with 11 randomly generated opponents. The preference profiles created for all of the experiments carried out in this chapter correspond to our motivating scenario: group booking. In order to capture stochastic variations in the different models, each possible negotiation was repeated 4 times. Thus, a total of 4400 negotiations were carried out per model and environment (i.e., short/long deadline). The results for this first experiment can be observed in Fig. 5.1.

As can be observed in the short deadline scenario (Fig. 5.1), the offers proposed by the representative model are more attractive for the opponent. This is reasonable since, in this case, the representative only negotiates attending to its own utility function. Therefore, it results in less conflict with the opponent and more trade-off possibilities. The behavior observed for the perfect agenda model and the similarity simple voting model are more surprising. Even though, in the first rounds, SSV proposes offers that report more utility for the opponent than

5. INTRA-TEAM STRATEGIES FOR NEGOTIATION TEAMS IN PREDICTABLE DOMAINS

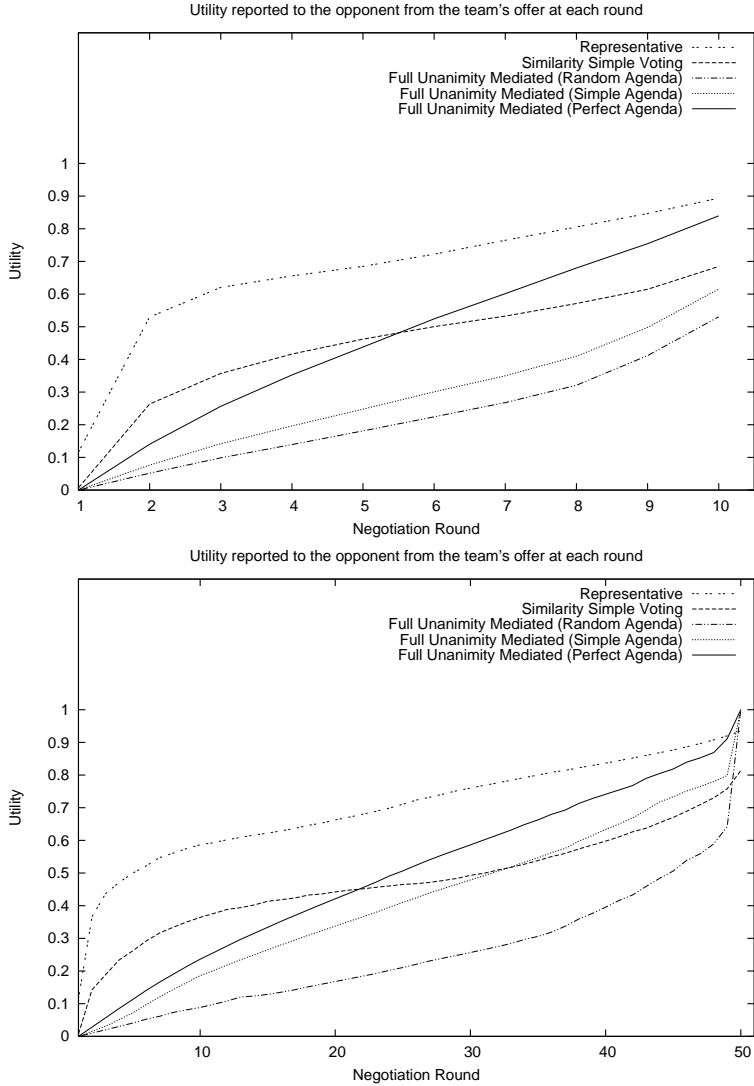


Figure 5.1: Average utility reported to the opponent by team's proposals per round in short (upper graphic) and long (lower graphic) deadline scenarios.

5.8 Empirical Analysis of Full Unanimity Mediated Intra-Team Strategy

those built by the perfect agenda model, as the negotiation advances, the perfect agenda model outperforms SSV. This happens at negotiation round 6. This may be explained by the fact that, at that point, more trade-off possibilities arise between all of the team members and the opponent, and the perfect agenda model is capable of exploiting them while assuring the desired aspiration level for each teammate. As for the simple agenda model, it performs slightly better than the random agenda model, but worse than the other methods in the experiment. This is explainable by the fact that, since the negotiation deadline is short, limited information can be used to learn the opponent's preferences. Consequently, the agenda built is closer to a random agenda than to the perfect agenda. In the case of the long deadline scenario, there are some differences that are worth highlighting. First, the representative model is still the one that is the most attractive for the opponent's interests. However, in this scenario, both the perfect agenda model and the simple agenda model are able to outperform SSV at some points of the negotiation process. Obviously, this happens earlier for the perfect agenda model since it represents perfect knowledge about the opponent's preferences. Hence, it is able to take advantage of possible trade-offs earlier in the negotiation. It happens approximately at round 22. Regarding the simple agenda model, it is able to outperform SSV around round 33. Differently to the first scenario, since the amount of information to learn from is greater, the simple agenda model is able to get closer to the perfect agenda and offer more attractive offers to the opponent.

In conclusion, results seem to point out that as the agenda gets closer to the ideal agenda, the offers are more likeable for the opponent. Obviously, the utility reported to the opponent by offers proposed by FUM is initially lower than other intra-team strategies that guarantee less degree of unanimity like RE and SSV, but as more information becomes available and the negotiation is longer, FUM with a proper learning mechanism is able to propose offers that are more likeable

5. INTRA-TEAM STRATEGIES FOR NEGOTIATION TEAMS IN PREDICTABLE DOMAINS

for the opponent. This suggests that FUM may benefit from negotiation processes where there is a long deadline.

5.8.2 Studying the Impact of ϵ_{a_i}

In this second experiment, we decided to study the impact of ϵ_{a_i} on the team's performance. It seems reasonable to think that low values of this parameter should help to construct offers that are more interesting for the opponent, but high values should impact negatively on the utility obtained by a_i . We devised an experiment where the value of ϵ_{a_i} was set in a uniform way for all of the team members. More specifically, we used the values 0, 0.02, 0.05, 0.07, 0.1, 0.12, 0.15, 0.17, 0.2 for ϵ_{a_i} . For the quality measures, we observed the minimum and the average utility of the team members. Two different environments were tested: short/long deadline, whose lengths are drawn from the uniform distributions $T_{op} = T_A = U[5, 10]$, $T_{op} = T_A = U[30, 60]$, respectively. The concession speed for both parties was set to be drawn from $\beta_{op} = \beta_A = U[0.4, 0.99]$. The reservation utility for the agents was drawn from a uniform distribution $RU_{op} = RU_{a_i} = U[0, 0.25]$. In this case, the learning method for the agenda was set to FUM-simple and the number of initial rounds to be taken into account was set to $k = \lfloor \frac{T_A}{4} \rfloor$. A total of 100 randomly generated teams with size $M = 4$ and random utility functions (4 issues) were confronted with 12 randomly generated opponents. Each possible negotiation was repeated 4 times. Thus, a total of 4800 negotiation were carried out per model and environment. The results for this experiment are shown in Table 5.1.

The results show a slight decrease in the utility (minimum utility and average utility) as ϵ_{a_i} gets larger. This behavior is found in almost every scenario tested. Those scenarios that do not show this pattern usually obtain very similar results for all of the configurations. Thus, the agents should choose $\epsilon_{a_i} = 0$ independently of the type of scenario where they negotiate. In the best case, the agent will get

5.8 Empirical Analysis of Full Unanimity Mediated Intra-Team Strategy

ϵ_{a_i}	Long deadline		Short deadline	
	Min.	Ave.	Min.	Ave.
0.00	0.49	0.72	0.35	0.60
0.02	0.50	0.71	0.37	0.61
0.05	0.49	0.68	0.37	0.58
0.07	0.48	0.67	0.37	0.57
0.10	0.49	0.66	0.37	0.56
0.12	0.48	0.65	0.37	0.56
0.15	0.48	0.64	0.37	0.55
0.17	0.47	0.63	0.38	0.54
0.20	0.46	0.61	0.38	0.55

Table 5.1: Average impact of ϵ_{a_i} on team performance. Min: Minimum utility of team members, Ave: Average utility of team members

a slightly better utility than other values of the parameter. In the worst case scenario, the agent will get a very similar utility to other values of the parameter ϵ_{a_i} . The value $\epsilon_{a_i} = 0$ corresponds to the agents only handing over those decision rights associated to issues that yield no interest at all for the agent.

It can also be observed that the average utility is impacted more negatively by increment of the ϵ_{a_i} parameter in the long deadline scenario than in the short deadline scenario. A thorough analysis of our results gave an answer to this phenomenon. The results suggest that higher values of ϵ_{a_i} reduce the average utility for the team members. However, the number of negotiations that ended with no agreement in the long deadline scenario when $\epsilon_{a_i} = 0$ was 151 (3.1% of the negotiation cases ended with an average utility equal to 0), whereas the number of failed negotiations was 404 (8.41%) when $\epsilon_{a_i} = 0$ and the deadline was short. As ϵ_{a_i} was increased to 0.2, the number of failed negotiations decreased to 35 (0.7%) in the long deadline scenario and 91 (1.8%) in the short deadline scenario. Thus, higher values for ϵ_{a_i} contribute to reaching an agreement in cases where no deal was found. This effect is more notorious in the short deadline scenario. Since the

5. INTRA-TEAM STRATEGIES FOR NEGOTIATION TEAMS IN PREDICTABLE DOMAINS

number of failed negotiations is greatly reduced in the short deadline scenario, the negative effect of higher ϵ_{a_i} is moderated since the new negotiations contribute with values for the average utility that are greater than or equal to 0. Despite this, the reduction in the number of failed negotiations is not enough to counter the negative impact of ϵ_{a_i} .

In general, ϵ_{a_i} can be considered as some sort of moderator for the initial demand. According to our results, in general, agents should not give up any decision right over an issue that yields interest for him. Only those decision rights associated to issues that yield no interest at all should be handed over. Hence, team members should always start demanding their highest aspiration level. This situation resembles results obtained in bilateral negotiation (18), where it was found that if the deadline is reasonably long, the agent should start demanding values close to their maximum utility.

5.8.3 Strategy Deviation

The proposed model assumes that team members state the truth when asked about which issue values they need to reach their desired utility level during the offer construction phase. When dealing with selfish agents, one risk faced is the fact that selfish agents may not tell the truth in order to maximize their own utility. In this case, it seems clear that team members have no incentives to ask for less issue value than they need since it may end up in an agreement with a utility inferior to the desired level of utility. However, team members may have incentives to demand more value if that maximizes their utilities (be more demanding). For a team member to play strategically, it would need to have some knowledge about team members' and opponent's utility functions, deadlines, reservation utilities, and other agents' strategies. We aim to propose negotiation models for open environments, where information is private. Therefore, agents usually have limited

5.8 Empirical Analysis of Full Unanimity Mediated Intra-Team Strategy

and uncertain information regarding the negotiation conditions. This leads to the question of whether or not team members would achieve higher utilities by deviating from the proposed strategy.

In this subsection we analyze whether or not team members have incentives to deviate from the proposed strategy in the offer construction phase. For this matter, we designed two types of *deviated* team members. The first type of deviated agent, which we will name *slightly deviated*, behaves exactly as the standard behavior proposed for team members in this chapter. However, during the iterated offer construction phase, the agent does not ask for the value it needs from j , but a value that reports higher utility than it needs. The amount of extra utility that it attempts to achieve is controlled by a parameter d_i . When $d_i > 1$, the team member demands more value than it needs, as it can be appreciated in the formula:

$$x_{a_i,j} = \operatorname{argmin}_{x \in [0,1]} (d_i \times (s_{a_i}(t) - U_{a_i}(X_{A \rightarrow op}^t)) - w_{a_i,j} V_{a_i,j}(x)) \quad (5.27)$$

When the utility of the partial offer exceeds or equals the desired utility level $s_{a_i}(t)$, the agent abandons the offer construction phase at that round. The effect of this behavior is that, when the agent is asked to set an issue which can report the desired utility, it demands more value for that issue and then leaves the iterated building process. For instance, if a seller agent needs 250\$ for the price issue in order to reach its desired utility level and $d_i = 1.25$, it will ask for $250 \times 1.25 = 312.5$ \$ instead. The second type of deviated team member, named *highly deviated*, behaves as the *slightly deviated* team member but when it has reached its desired utility level, it stays an additional issue in the iterated building process. When asked about the value of that extra issue, the *highly deviated* agent asks for a random value that reports between 10% and 50% of the issue's utility. For instance, assuming that the price is scaled between 0\$ and 1000\$, a *highly deviated* seller that

5. INTRA-TEAM STRATEGIES FOR NEGOTIATION TEAMS IN PREDICTABLE DOMAINS

has reached its desired utility level would ask for a price value between 100\$ and 500\$. After setting the extra negotiation issue, the *highly deviated* team member leaves the offer construction phase.

We set the parameters of our model to the same values used in the previous experiment: $T_A = T_{op} = U[30, 60]$ for long deadline scenarios, $T_A = T_{op} = U[5, 10]$ for short deadline scenarios, $RU_{a_i} = RU_{op} = U[0, 0.25]$, and $\beta_A = \beta_{op} = U[0.4, 0.99]$. A total of 100 randomly generated teams with size $M = 4$ and random utility functions (4 issues) were confronted with 12 randomly generated opponents. Each possible negotiation was repeated 4 times. Thus, a total of 4800 negotiations were carried out per model and environment. We studied the effect of the number of *slightly deviated* agents $|A|_{sd} = \{1, 2, 3, 4\}$ (the rest of team members having the standard behavior), the effect of the number of *highly deviated* agents $|A|_{hd} = \{1, 2, 3, 4\}$ (the rest of team members having the standard behavior), and different values for $d_i = \{1.25, 1.50, 1.75\}$ (all of the deviated agents were set to have the same d_i). The quality measure studied was the average utility since an increment in the utility of one of the team members will always have a positive effect on the average utility (same type of valuation functions). The results of the experiment are depicted in Table 5.2. Some of the combinations are empty since they do not make sense in practice (e.g., 0 *deviated* agents and $d_i > 1$). We only show the results for the long deadline scenario, but it should be noted that the same pattern was found for short deadline scenarios. It can be observed that all the combinations obtain similar results in terms of average utility. There is only a slight decrement in the average utility as we move to more demanding attitudes (e.g., $|A|_{hd} = 4, d_i = 1.75$). Even though, the differences between the most demanding behaviors and other behaviors are not large enough to be considered significant. A closer look at the negotiation traces explained the previous results. While being more demanding may obtain higher utilities in successful negotiations,

5.8 Empirical Analysis of Full Unanimity Mediated Intra-Team Strategy

	$d_i = 1$	$d_i = 1.25$	$d_i = 1.5$	$d_i = 1.75$
$ A = 4$	[0.71-0.72]	-	-	-
$ A _{sd} = 1$	-	[0.70-0.72]	[0.71-0.72]	[0.70-0.71]
$ A _{sd} = 2$	-	[0.71-0.72]	[0.71-0.72]	[0.70-0.72]
$ A _{sd} = 3$	-	[0.71-0.72]	[0.70-0.72]	[0.69-0.70]
$ A _{sd} = 4$	-	[0.70-0.72]	[0.70-0.72]	[0.69-0.71]
$ A _{hd} = 1$	-	[0.70-0.72]	[0.71-0.72]	[0.70-0.71]
$ A _{hd} = 2$	-	[0.70-0.71]	[0.70-0.72]	[0.69-0.71]
$ A _{hd} = 3$	-	[0.69-0.71]	[0.69-0.71]	[0.69-0.70]
$ A _{hd} = 4$	-	[0.69-0.70]	[0.69-0.70]	[0.68-0.69]
	$d_i = 1$	$d_i = 1.25$	$d_i = 1.5$	$d_i = 1.75$
$ A = 4$	206	-	-	-
$ A _{sd} = 1$	-	208	205	199
$ A _{sd} = 2$	-	202	230	236
$ A _{sd} = 3$	-	199	240	243
$ A _{sd} = 4$	-	248	267	287
$ A _{hd} = 1$	-	202	189	236
$ A _{hd} = 2$	-	241	246	274
$ A _{hd} = 3$	-	256	301	302
$ A _{hd} = 4$	-	299	292	324

Table 5.2: Confidence intervals (upper table) for the average utility of team members and number of failed negotiations (lower table) depending on the number of *deviated* agents and d_i .

it may also lead to a higher number of failed negotiations, thus leading to lower or equal average utilities. These results can be observed also at Table 5.2, where there is a clear tendency for the number of failed negotiations to increase as team members deviate further from the standard behavior. Thus, the experimental results suggest that team members may not have incentives to deviate much from the proposed strategy.

5.9 Studying the Impact of the Negotiation Environment on Intra-Team Strategies' Performance

Section 5.8 allowed us to explore Full Unanimity Mediated in detail, an intra-team strategy that is capable of guaranteeing unanimity regarding team decisions. In this section we explore how the four intra-team strategies presented in this chapter perform under different environmental conditions. First, the environmental conditions and performance measures studied in this chapter are introduced to the reader. Then, we describe the experiments carried out, and we analyze the results provided by each intra-team strategy.

5.9.1 Negotiation Environment Conditions & Team Performance

We consider that the negotiation environment plays a very important part in team dynamics. It may not be the same using a representative approach in a setting where all of the team members' preferences are very similar than using the same strategy in a setting where team members' preferences are exactly the opposite. Since which conditions of the negotiation environment are available depend on the application, we decided to focus on those general conditions that are present in almost every negotiation scenario involving negotiation teams: opponent deadline, team deadline, team members' preference similarity, opponent concession speed, and team size.

Regarding team performance, it is acknowledged that there are several well known social welfare measures to assess the quality of decisions in a society. A negotiation team can be considered a small society, and, thus, social welfare measures

5.9 Studying the Impact of the Negotiation Environment on Intra-Team Strategies' Performance

can also be considered appropriate metrics for measuring negotiation teams' performance. More specifically, we study the impact of the negotiation environment on the minimum utility of team members (i.e., egalitarian social welfare (160)), and the average utility of team members (i.e., a special case of ordered weighted averaging (160)). However, we do not exclusively restrain our analysis to social welfare measures. Computational measures like the number of negotiation rounds are also analyzed for all of the intra-team strategies.

5.9.1.1 Environment Condition: Opponent Deadline Length

One of the issues that can affect the negotiation process is the number of interactions that the opponent has until he decides that negotiating is no longer worthy, namely opponent deadline T_{op} . We partitioned the opponent negotiation deadline in three different classes: short deadline $T_{op} = U[5, 10] = S$, medium deadline $T_{op} = U[11, 29] = M$, and long deadline $T_{op} = U[30, 60] = L$.

5.9.1.2 Environment Condition: Team Deadline Length

Similarly, the maximum number of rounds that the team has to negotiate may also impact the performance of the different intra-team strategies. As in the case of the opponent deadline, we partitioned the team deadline in three different classes: short deadline $T_A = U[5, 10] = S$, medium deadline $T_A = U[11, 29] = M$, and long deadline $T_A = U[30, 60] = L$.

5.9.1.3 Environment Condition: Team Similarity

25 different linear utility functions were randomly generated. These utility functions represented the preferences of potential team members for $n=4$ negotiation issues, whose $V_i(\cdot)$ is the same type (i.e., monotonically increasing or monotonically decreasing) for all of the team members. 25 linear utility functions were

5. INTRA-TEAM STRATEGIES FOR NEGOTIATION TEAMS IN PREDICTABLE DOMAINS

generated to represent the preferences of opponents. These utility functions were generated by taking potential teammates' utility functions and reversing the type of $V_i(\cdot)$.

In order to determine the preference diversity in a team, we decided to compare team members' utility functions. We introduce a dissimilarity measure based on the utility difference between offers. The dissimilarity between two teammates can be measured as follows:

$$D(U_{a_i}(\cdot), U_{a_j}(\cdot)) = \frac{\sum_{\forall X | \text{valid_offer}(X)} |U_{a_i}(X) - U_{a_j}(X)|}{\# \text{ possible offers}} \quad (5.28)$$

If the dissimilarity between two team members is to be measured exactly, it needs to sample all of the possible offers. However, this is not feasible in the current domain where there is an infinite number of offers. Therefore, we limited the number of sampled offers to 1000 per dissimilarity measure. Due to the fact that a team is composed by more than two members, it is necessary to provide a team dissimilarity measure. We define the team dissimilarity measure as the average of the dissimilarity between all of the possible pairs of teammates.

For all of the teams, we measured their dissimilarity and calculated the dissimilarity mean \bar{dt} and standard deviation σ . We used this information to divide the spectrum of negotiation teams according to their diversity. Our design decision was to consider those teams whose dissimilarity was greater than, or equal to $\bar{dt} + 1.5\sigma$ as very dissimilar, and those teams whose dissimilarity was lower than, or equal to $\bar{dt} - 1.5\sigma$ as very similar. The rest of the cases are considered as scenarios where teams have an average similarity. In each case, 100 random negotiation teams were selected for the tests, that is, 100 teams were selected to represent the very similar team case, and 100 teams were selected to represent the very dissimilar team case. These teams participate in the different environmental scenarios, where they are confronted with one random half of all of the possible

5.9 Studying the Impact of the Negotiation Environment on Intra-Team Strategies' Performance

individual opponents. Therefore, each environmental scenario (complete instantiation of all the environmental conditions) consists of $100 \times 12 \times 4 = 4800$ different negotiations (each negotiation is repeated 4 times to capture stochastic variations in the different intra-team strategies).

5.9.1.4 Environment Condition: Opponent Concession Speed

The concession speed of the opponent during the negotiation process β_{op} may determine the final quality of the agreement for team members. For instance, if the opponent concedes very quickly towards its reservation utility, better agreements for the team may come earlier in the negotiation process. In those cases, even intra-team strategies that guarantee less degree of unanimity may achieve good results. We divided the family of concession speeds based on the classic classification of time-tactics: we considered that when $\beta_{op} = U[0.1, 0.49] = VB$ the concession speed is very bouldware, when $\beta_{op} = U[0.5, 0.99] = B$ the concession speed is bouldware, when $\beta_{op} = U[1, 10] = C$ the concession speed is conceder, when $\beta_{op} = U[11, 40] = VC$ the concession speed is very conceder. Similarly, when we refer to β_A (the team concession speed), we will also employ the same partition in bouldware (B), very bouldware (VB), conceder (C), and very conceder (VC).

5.9.1.5 Environment Condition: Number of Team Members

We think that the number of team members may also influence the performance of the different intra-team strategies. Some of the strategies may become too demanding when the number of team members increases and it may result in more negotiations ending in failure. Therefore, we decided to study the effect of the team size on the performance of the different intra-team strategies. The number of team members $|A|$ ranged from 4 to 8. This number of team members

5. INTRA-TEAM STRATEGIES FOR NEGOTIATION TEAMS IN PREDICTABLE DOMAINS

is motivated by the negotiation case employed in our experiments. We consider that groups of friends from 4 to 8 persons are a reasonable number in practice.

5.9.1.6 Team Performance: Number of Negotiation Rounds

The number of negotiation rounds considers the number of interactions between the team and the opponent. It is a measure employed to assess the negotiation time employed to reach a final agreement. In our study, every pair offer/counter-offer in the negotiation thread is considered as a negotiation round. In equal conditions of utility performance, those intra-team strategies that spend less negotiation rounds are preferred since they employ less negotiation time to reach a final agreement.

5.9.1.7 Team Performance: Minimum Utility of Team Members

The minimum utility of team members (Min.) in a negotiation represents the utility of the final agreement for the less benefited team member. If the final agreement is X and the team is composed of M different team members $A = \{a_1, a_2, \dots, a_M\}$, the minimum utility of team members can be calculated as:

$$Min.(X) = \min_{1 \leq i \leq M} U_{a_i}(X) \quad (5.29)$$

In applications where there is a strong bond among team members (i.e., the group of travelling friends), team members may attempt to maximize the minimum utility of team members in order to avoid extremely unsatisfied team members and a degradation of the relationship among team members. Even if a strong bond is not present among team members, an agent may desire to maximize the minimum utility of team members if it thinks that its own utility is going to be less favored utility by the final agreement.

5.9 Studying the Impact of the Negotiation Environment on Intra-Team Strategies' Performance

5.9.1.8 Team Performance: Average Utility of Team Members

The average utility of team members (Ave.) in a negotiation represents the average utility of the final agreement for all of the team members. If the final agreement is X and the team is composed of M different team members $A = \{a_1, a_2, \dots, a_M\}$, the average utility of team members can be calculated as:

$$Ave.(X) = \frac{1}{M} \sum_{1 \leq i \leq M} U_{a_i}(X) \quad (5.30)$$

A less conservative agent may desire to maximize the average utility of team members if it thinks that its own utility is not going to be less favored utility by the final agreement.

5.9.2 Results

It should be highlighted that the number of variables included in the study give a large combination of scenarios. Due to space limitations and for the the comfort of the reader, we only include those results which are the most interesting from our point of view. Next, we analyze the results of the experiments that were carried out in this chapter.

5.9.2.1 Number of Negotiation Rounds

Although we measured the number of negotiation rounds in each experiment, we found that a general pattern was found in almost every experiment. Thus, instead of commenting the results for the number of negotiation rounds in each experimental section, we decided to present the performance of the four intra-team strategies according to the number of negotiation rounds just once. As a sample for this behavior, we can observe the number of negotiation rounds spent by each intra-team strategy when team and opponent have a long deadline ($T_{op} = L$ and

5. INTRA-TEAM STRATEGIES FOR NEGOTIATION TEAMS IN PREDICTABLE DOMAINS

$T_A = L$), the number of team $|A|$ members is set to 4, and the opponent uses different concessions speeds β_{op} in Table 5.3.

As long as the concession speed of the four intra-team strategies is the same, *RE is usually the fastest intra-team strategy, followed by SSV, then SBV, and finally FUM*. Since less unanimity is guaranteed among team members, it is logical that there may be less conflict with the opponent and, thus, agreements are found faster with low unanimity strategies like RE and SSV. The main exception for this rule is when team members are very similar and the opponent uses either bouldware or very bouldware concession speeds. In those cases, FUM is able to finalize negotiations successfully in fewer rounds than SBV (and sometimes SSV). The learning heuristic employed by FUM benefits from the fact that the opponent usually concedes more in those issues that are less important and, thus, it is able to infer a proper agenda and propose better offers to the opponent (i.e., ending the negotiation faster). This pattern disappears as team members get more dissimilar. In that case, FUM also has to deal with more intra-team conflict, which in turn results in more demanding offers needed to guarantee unanimity.

Additionally, as expected, *as the concession strategy of team members becomes more conceder, the number of negotiation rounds spent is lower*. Thus, RE using $\beta_A = VB$ is slower than RE using $\beta_A = B$, which is slower than RE using $\beta_A = C$, which is slower than RE using $\beta_A = VC$.

The number of negotiation rounds spent by each intra-team strategy is interesting as a selection criterion when intra-team strategies perform equally in utility terms (minimum or average utility). For instance, if SBV and FUM tie in utility terms, a team is suggested to select SBV most of the times due to the fact that it usually requires less negotiation rounds, if SSV and SBV tie in utility terms, the team should select SSV since it usually requires less rounds than SBV, and so forth.

5.9 Studying the Impact of the Negotiation Environment on Intra-Team Strategies' Performance

Very Similar, $T_{op} = T_A = L, M = 4$					Very Dissimilar, $T_{op} = T_A = L, M = 4$				
	$\beta_{op} = VC$	C	B	VB		$\beta_{op} = VC$	C	B	VB
RE $\beta_A = VC$	2.01	2.28	7.25	19.57	RE $\beta_A = VC$	2.03	2.33	7.78	19.71
SSV $\beta_A = VC$	2.02	2.41	8.35	22.08	SSV $\beta_A = VC$	2.00	2.71	9.97	24.35
SBV $\beta_A = VC$	2.01	2.70	10.48	24.83	SBV $\beta_A = VC$	2.05	3.44	12.99	27.33
FUM $\beta_A = VC$	2.11	2.63	10.31	24.10	FUM $\beta_A = VC$	2.90	4.52	16.29	30.70
RE $\beta_A = C$	2.39	3.77	11.07	23.47	RE $\beta_A = C$	2.28	3.30	10.98	22.84
SSV $\beta_A = C$	2.73	5.17	13.17	25.33	SSV $\beta_A = C$	2.45	5.17	14.83	27.43
SBV $\beta_A = C$	3.02	6.18	15.55	27.32	SBV $\beta_A = C$	2.99	7.12	18.64	30.08
FUM $\beta_A = C$	4.09	6.23	14.01	26.45	FUM $\beta_A = C$	6.54	10.47	21.13	32.66
RE $\beta_A = B$	9.17	13.63	22.48	30.73	RE $\beta_A = B$	6.57	10.02	19.94	29.19
SSV $\beta_A = B$	15.53	19.99	26.97	32.95	SSV $\beta_A = B$	12.09	18.26	26.42	33.52
SBV $\beta_A = B$	17.96	22.40	28.88	34.21	SBV $\beta_A = B$	16.50	22.74	30.54	35.76
FUM $\beta_A = B$	20.31	23.25	25.59	33.09	FUM $\beta_A = B$	25.93	28.53	30.97	36.96
RE $\beta_A = VB$	22.50	25.47	31.59	35.51	RE $\beta_A = VB$	17.22	21.14	28.94	34.50
SSV $\beta_A = VB$	28.62	31.44	35.27	37.22	SSV $\beta_A = VB$	25.44	30.04	34.59	37.64
SBV $\beta_A = VB$	31.50	33.21	36.24	37.80	SBV $\beta_A = VB$	30.10	33.29	36.77	38.74
FUM $\beta_A = VB$	32.77	33.67	33.97	37.15	FUM $\beta_A = VB$	35.00	36.59	36.39	39.01

Table 5.3: Average number of rounds when both parties have a long deadline.

5.9.2.2 Same Type of Deadlines

The next set of experiments that we conducted consisted in assessing which intra-team strategies work better when both parties have the same type of deadline. We studied those scenarios where both parties have short deadlines or long deadlines. For each deadline scenario, we tested very dissimilar teams, average similarity teams, and very similar team. We gathered information about the minimum and average utility of team members regarding each possible strategy configuration (e.g., team concession speeds, intra-team strategies, opponent concession speeds, etc.). The number of team members remained static at $|A| = 4$. The reservation utilities are drawn from uniform distributions $RU_{op} = U[0, 0.25]$ and $RU_{ai} = U[0, 0.25]$.

The results for this batch of experiments can be found in Table 5.4. It shows the average of the minimum utility of team members (Min.) and the average of the

5. INTRA-TEAM STRATEGIES FOR NEGOTIATION TEAMS IN PREDICTABLE DOMAINS

Very Similar, $T_{op} = T_A = S, M = 4$								
	$\beta_{op} = VC$		$\beta_{op} = C$		$\beta_{op} = B$		$\beta_{op} = VB$	
	Min.	Ave.	Min.	Ave.	Min.	Ave.	Min.	Ave.
RE $\beta = B$	0.68	0.79	0.59	0.71	0.41	0.55	0.29	0.40
SSV $\beta = B$	0.70	0.79	0.61	0.70	0.45	0.56	0.33	0.42
SBV $\beta = B$	0.71	0.79	0.62	0.70	0.47	0.56	0.34	0.42
FUM $\beta = B$	0.69	0.77	0.62	0.72	0.48	0.60	0.35	0.47
Very Similar, $T_{op} = T_A = L, M = 4$								
	$\beta_{op} = VC$		$\beta_{op} = C$		$\beta_{op} = B$		$\beta_{op} = VB$	
	Min.	Ave.	Min.	Ave.	Min.	Ave.	Min.	Ave.
RE $\beta = B$	0.75	0.85	0.66	0.77	0.47	0.61	0.32	0.45
SSV $\beta = B$	0.75	0.83	0.67	0.76	0.51	0.62	0.37	0.48
SBV $\beta = B$	0.77	0.83	0.69	0.76	0.54	0.62	0.40	0.48
FUM $\beta = B$	0.75	0.81	0.69	0.77	0.60	0.72	0.44	0.56
Very Dissimilar, $T_{op} = T_A = S, M = 4$								
	$\beta_{op} = VC$		$\beta_{op} = C$		$\beta_{op} = B$		$\beta_{op} = VB$	
	Min.	Ave.	Min.	Ave.	Min.	Ave.	Min.	Ave.
RE $\beta = B$	0.26	0.62	0.19	0.54	0.11	0.42	0.07	0.30
SSV $\beta = B$	0.50	0.73	0.41	0.66	0.29	0.52	0.18	0.37
SBV $\beta = B$	0.55	0.73	0.46	0.65	0.34	0.52	0.21	0.36
FUM $\beta = B$	0.55	0.70	0.46	0.65	0.34	0.58	0.23	0.44
Very Dissimilar, $T_{op} = T_A = L, M = 4$								
	$\beta_{op} = VC$		$\beta_{op} = C$		$\beta_{op} = B$		$\beta_{op} = VB$	
	Min.	Ave.	Min.	Ave.	Min.	Ave.	Min.	Ave.
RE $\beta = B$	0.38	0.72	0.27	0.63	0.11	0.45	0.06	0.31
SSV $\beta = B$	0.57	0.78	0.45	0.70	0.28	0.57	0.18	0.43
SBV $\beta = B$	0.65	0.79	0.55	0.71	0.37	0.56	0.24	0.41
FUM $\beta = B$	0.62	0.75	0.56	0.71	0.48	0.72	0.31	0.54

Table 5.4: Average of the minimum utility of team members (Min.) and the average of the average utility of team members (Ave.).

average utility of team members (Ave.). It only shows the results for intra-team strategies using a Boulware concession speed since we found that the best results are found in this setting. Additionally, it should be highlighted that the results for average similarity teams were very similar to the very dissimilar case. Thus, we did not include this information in the tables to show a more compact and interpretable representation. The results in bold font indicate those configurations that are statistically better and different (t-test $\alpha = 0.05$) to the rest of configurations.

When both parties have a short deadline (first and third sub-table in Table 5.4), independently of team similarity, SBV $\beta = B$ and FUM $\beta = B$ are usually

5.9 Studying the Impact of the Negotiation Environment on Intra-Team Strategies' Performance

the best options for the minimum utility. The unanimity and semi-unanimity rules employed by this strategy make possible for the worst affected team member to ensure that its situation is better with respect to other rules. As for the average utility of team members, FUM $\beta = B$ is usually the best option. The only exception for this pattern is when the opponent uses conceder strategies ($\beta_{op} = VC$ or $\beta_{op} = C$). In that case, all of the strategies perform similarly, especially when team members are very similar. For instance, we can observe that RE, SSV, SBV $\beta = B$ are the best option for the average utility of team members when the deadline is short, team members are very similar, and the opponent uses a very conceder strategy. In the same setting, but with the opponent using a conceder strategy, FUM is statistically better but the differences are not very important (less than a 1.8%).

However, *when both parties have a long deadline to negotiate (subtables 2 and 4 in Table 5.8), FUM $\beta = B$ becomes the best choice for the minimum and average utility of team members in almost every scenario.* The only exceptions for this superiority are, again, scenarios where the opponent employs conceder strategies. For instance, when the deadline is long, team members are very dissimilar, and the opponent uses a very conceder strategy, SBV $\beta = B$ is the best intra-team strategy for the minimum and average utility of team members.

We can also observe that *RE and SSV are specially affected by very dissimilar preferences' scenarios.* When team members are very similar, both strategies are capable of being close to SBV and FUM in the minimum and average utility of team members as long as the opponent plays conceder strategies. However, both intra-team strategies' results get further from those of SBV and FUM when conflict is introduced inside the team (average similarity and very dissimilar scenario). These intra-team strategies are not able to tackle situations where team members have

5. INTRA-TEAM STRATEGIES FOR NEGOTIATION TEAMS IN PREDICTABLE DOMAINS

very dissimilar preferences due to the type of decision rule applied, and their use in such situations is discouraged.

The reason why several strategies perform similarly in utility terms when the opponent plays conceder strategies is simple: Since the opponent concedes very fast in the first rounds of the negotiation process, as long as the team does not concede very fast (i.e., bouldware strategy), all of the strategies are capable of finding a reasonable good agreement in the first rounds by letting the opponent concede and then accepting the opponent's offer. However, there is an additional reading that explains why strategies like FUM, which guarantees unanimity regarding team decisions, does not perform so well when the opponent uses conceder strategies. FUM relies on the assumption that the opponent concedes very little in those issues that are important for its interests at the first rounds. However, when the concession strategy carried out by the opponent is conceder or very conceder (a more acute effect) big concessions are usually carried out at the first rounds. Thus, FUM is not able to infer an appropriate agenda. In Section 5.8, it was shown that as the agenda gets further from the real ranking of opponent preferences, the more demanding becomes the strategy. This may have a negative effect in the negotiation, since more negotiations may end in failure due to the high demands of the team. In fact a slight effect is observed in the results: when the opponent uses a bouldware strategy, the percentage of successful negotiations is 94.6% which is greater than the 92.6% obtained when the opponent uses a conceder strategy and the 93.1% obtained when the opponent uses a very conceder strategy.

Another issue found in the results is the difference between FUM and other strategies when the deadline is long. FUM tends to obtain better results when the deadline is long for both parties and the differences with the other intra-team strategies become greater when compared with the short deadline scenario. The reason for this phenomenon is similar to the reason mentioned in the paragraph

5.9 Studying the Impact of the Negotiation Environment on Intra-Team Strategies' Performance

above. FUM is a strategy that relies on the information gathered in the negotiation process. Thus, when the number of interactions is lower, the agenda inferred by the trusted mediator is more different to the ideal agenda. When the agenda deviates from the ideal agenda, offers proposed by the team are more demanding and less probable to be accepted by the opponent. As a matter of fact, the reader can notice the difference on average between FUM $\beta = B$ in long deadline scenarios and the results obtained by FUM $\beta = B$ in short deadline scenarios counterpart is approximately 8.1%, whereas it is approximately 5.3% for SBV, 5.4% for SSV and 5.8% for RE. Logically, every intra-team strategy benefits from having a longer deadline, but the results suggest that FUM benefits still more than the rest of intra-team strategies due to its learning heuristic based on the amount of information.

5.9.2.3 Different Types of Deadlines

The next batch of experiments consisted in studying the behavior of the different intra-team strategies when both parties have strongly different types of deadline. Thus, in this case, one of the two parties has a deadline which is way lower than the deadline of the other party. Clearly, the party with a lower deadline is at disadvantage with respect to the other party since it has fewer offers to send before ending the negotiation in failure, and the pressure to accept the opponent's offers arises earlier.

Short Team Deadline and Long Opponent Deadline First, we start by analyzing the case where the deadline of the team is shorter than the deadline of the opponent party. Hence, $T_A = U[5, 10]$ and $T_{op} = U[30, 60]$. The reservation utilities are drawn from uniform distributions $RU_{op} = U[0, 0.25]$ and $RU_{a_i} = U[0, 0.25]$. The results of this experiment can be found in Table 5.5. The results

5. INTRA-TEAM STRATEGIES FOR NEGOTIATION TEAMS IN PREDICTABLE DOMAINS

in bold font indicate those configurations that are statistically better and different (t-test $\alpha = 0.05$) to the rest of configurations.

In this case, the team has a shorter deadline and, thus, it should be at disadvantage with respect to the opponent. However, we can observe that *when the opponent uses a conceder or very conceder strategy, the results are similar to the analogous case where both parties had a short deadline*. These results can be explained due to the fact that since the opponent concedes very quickly, a good deal can be found for the team in the first rounds of the negotiation process and the team is not affected by the fact that its deadline is shorter. Nevertheless, as the opponent starts to employ boulware strategies, there is a clear negative effect (i.e., a reduction) on the minimum and average utility of team members: *all of the strategies are affected by the fact that the team has a shorter deadline*. In the scenario where both parties have a short deadline, the average for the average utility of team members in conceder settings¹ is approximately 0.67, and the average for the average utility of team members in boulware settings² is approximately 0.45. Thus, the average utility for team members is reduced a 25%. In the present setting, the average of average utility of team members in conceder settings is approximately 0.63, whereas the average of the average utility of team members in boulware settings is approximately 0.10. Therefore, the average utility of team members is reduced a 53%, approximately doubling the difference found in the case where both parties had a short deadline.

When *team members are very similar* (upper sub-table in Table 5.5), it can be observed that, as in the scenario where both parties have a short deadline and team members are very similar, several strategies perform very similarly. The

¹This measure is calculated averaging the average for the average utility of team members for all of the intra-team strategies when $\beta_{op} = C$ and $\beta_{op} = VC$ in Table 5.4

²This measure is calculated averaging the average for the average utility of team members for all of the intra-team strategies when $\beta_{op} = B$ and $\beta_{op} = VB$ in Table 5.4

5.9 Studying the Impact of the Negotiation Environment on Intra-Team Strategies' Performance

main difference resides in the fact that the only strategy capable of reaching similar results to FUM $\beta = B$ in the minimum and average utility is RE $\beta = B$. Differently to the case when team members are very similar and the deadline for both parties is short, the RE $\beta_A = B$ strategy is capable of achieving similar results to the other intra-team strategies even in less conceding settings ($\beta_{op} = C$, $\beta_{op} = B$, and $\beta_{op} = VB$). These results suggest that, *despite not assuring any minimum level of unanimity, employing a representative with a reasonably slow concession (boulware) leads to good results compared with those obtained by other intra-team strategies*. A closer look at the experiments threw some light over these results. For instance, when $\beta_{op} = B$, the number of successful negotiations was 2695 for RE $\beta_A = B$, 1925 for FUM $\beta_A = B$, 1855 for SBV $\beta_A = B$, and 2394 for SSV $\beta_A = B$. The average utility for successful negotiations was 0.32 for RE $\beta_A = B$, 0.34 for SSV $\beta_A = B$, 0.39 for SBV $\beta_A = B$, and 0.42 for FUM $\beta_A = B$. Hence, despite obtaining less quality results in successful negotiations, the representative approach becomes a good option for these scenarios because it leads to a great number of negotiations ending in success where other intra-team strategies fail to succeed (utility=0). SSV, UBS, and FUM need more interactions to find a satisfactory deal, but when they find it, it is better in utility terms. However, in average, a representative approach may be more adequate for settings where the team has a shorter deadline than the opponent.

As for the scenario where *team members are very dissimilar* (lower sub-table in Table 5.5), we can observe that the negative effect produced by having a shorter deadline is especially acute when the opponent uses boulware or very boulware concessions. The dissimilarities between team members, and the fact that there are very few interactions to find a deal that satisfies both team and opponent, contribute to a strong reduction in the minimum and the average utility of team members. In terms of the minimum utility of team members, *FUM* and *SBV*

5. INTRA-TEAM STRATEGIES FOR NEGOTIATION TEAMS IN PREDICTABLE DOMAINS

$\beta_A = B$ work better when the opponent uses conceder or very conceder concessions. However, *almost every intra-team strategy performs equally bad in terms of the minimum utility of team members when the opponent moves to bouldware concessions* (especially in the very bouldware case). In this case, the representative approach can no longer compete with the rest of strategies in terms of utility in most scenarios. Nevertheless, despite team members being very dissimilar and RE not guaranteeing any unanimity regarding team decisions, RE performs slightly better than the rest in terms of the average utility of team members when the opponents concedes using bouldware. The explanation to this phenomenon is similar to the case where team members were very similar: a lesser number of negotiations end in failure (26% failures for RE, 33% for SSV, 48% for SBV, and 46% for FUM), which compensates for the dissimilarity between team members' preferences and the unanimity level guaranteed by RE. In any case, the utility obtained for team members is so low in the average and minimum utility of team members that, in some cases, *it may even be better not to negotiate with such kind of opponent and spend computational resources in looking for another alternative.*

Long Team Deadline and Short Opponent Deadline In this case, the team has an advantage over the opponent since its maximum deadline is way longer than the opponent's deadline. The goal of these experiments is to determine the combination of intra-team strategies and negotiation parameters that maximize the different social welfare measures employed. Thus, if the team has a maximum deadline equal to the uniform distribution $T_A = U[30, 60]$, the team may decide to play (prior to the negotiation) a different class of deadline like a medium deadline ($T_A = U[11, 29]$) or a short deadline ($T_A = U[5, 10]$) if the results of the simulation suggest that better results are obtained by not playing the maximum deadline. Thus, we also show the results for teams that play a medium deadline, and teams

5.9 Studying the Impact of the Negotiation Environment on Intra-Team Strategies' Performance

Very Similar, $T_{op} = L, T_A = S, M = 4$								
	$\hat{\beta}_{op} = VC$		$\hat{\beta}_{op} = C$		$\hat{\beta}_{op} = B$		$\hat{\beta}_{op} = VB$	
	Min.	Ave.	Min.	Ave.	Min.	Ave.	Min.	Ave.
RE $\beta = B$	0.64	0.75	0.48	0.60	0.15	0.24	0.02	0.04
SSV $\beta = B$	0.64	0.74	0.46	0.57	0.14	0.22	0.02	0.04
SBV $\beta = B$	0.65	0.74	0.47	0.56	0.13	0.19	0.01	0.02
FUM $\beta = B$	0.65	0.74	0.49	0.61	0.15	0.22	0.02	0.04
Very Dissimilar, $T_{op} = L, T_A = S, M = 4$								
	$\hat{\beta}_{op} = VC$		$\hat{\beta}_{op} = C$		$\hat{\beta}_{op} = B$		$\hat{\beta}_{op} = VB$	
	Min.	Ave.	Min.	Ave.	Min.	Ave.	Min.	Ave.
RE $\beta = B$	0.245	0.596	0.153	0.482	0.025	0.170	0.002	0.040
SSV $\beta = B$	0.459	0.703	0.280	0.540	0.035	0.156	0.002	0.017
SBV $\beta = B$	0.511	0.706	0.313	0.496	0.028	0.082	0.001	0.064
FUM $\beta = B$	0.520	0.704	0.336	0.545	0.026	0.084	0.001	0.060

Table 5.5: Average for the minimum utility of team members (Min.) and the average utility of team members (Ave.) when the team has a short deadline and the opponent has a long deadline.

that play a short deadline. In this experiment, the opponent plays a short deadline $T_{op} = U[5, 10]$. The reservation utilities are drawn from uniform distributions $RU_{op} = U[0, 0.25]$ and $RU_{a_i} = U[0, 0.25]$. The results of this experiment for the very similar scenario can be observed in Fig. 5.2 and 5.3, whereas the results for the very dissimilar scenario can be observed in Fig. 5.4 and 5.5. The dots indicate those configurations that perform statistically better than the rest (t-test, $\alpha = 0.05$).

We start by analyzing the results for scenarios where team members are very similar (Fig. 5.2 and 5.3). We can observe that for situations where the opponent is very concenter, the team benefits from playing strategies with the same deadline. Since the opponent concedes very fast in the first negotiation rounds, the best deals for the team may be proposed in the first negotiation rounds. Playing a longer deadline may be risky since the team may have extremely high aspirations during the whole negotiation, which results in most offers being rejected and ending the negotiation in failure. As a matter of fact, the number of successful negotiations for

5. INTRA-TEAM STRATEGIES FOR NEGOTIATION TEAMS IN PREDICTABLE DOMAINS

intra-team strategies playing a short deadline and bouldware concession was 95.1%, 68% for medium deadline and bouldware concession, and 45% for long deadline and bouldware concession, 29% for medium deadline and very bouldware concession, and 14% for long deadline and very bouldware concession. Other configurations may have a higher number of successful negotiations, but they are not able to retain as much utility as the bouldware configuration. As the opponent starts to move towards strategies that concede more slowly, the best intra-team strategies for the team are those played with a medium deadline and bouldware strategy (RE, SBV and SSV $\beta = B$). In those cases, the opponent may not propose the best deals for the team until its last negotiation rounds. Thus, playing a slightly longer deadline with a bouldware concession comes at an advantage for the team since the team does not fully concede in the whole negotiation and still accepts last opponent's offers. Some strategies played with a medium deadline like FUM $\beta = B$ are still too demanding and end in more negotiation failures (they have very little information to learn the opponents' preferences).

The very dissimilar scenario (Fig. 5.4 and 5.5) is a little bit different. In this scenario, the team needs to deal with strong divergences in their preferences too. Thus, teams are prone to be more demanding in order to accommodate the preferences of as many team members as possible. We can observe that for cases where the opponent uses conceder strategies, the team should play bouldware strategies with the same deadline. Similarly to the very similar scenario, playing a longer deadline is risky since it results in extremely high aspirations and most offers being rejected. However, in the very dissimilar scenario, the transition from selecting short deadline strategies to selecting medium deadline strategies does not appear until the opponent uses bouldware strategies. This may be explained precisely due to the dissimilarity among team members, which requires stronger demands that are not met when playing medium deadline. As the opponent starts

5.9 Studying the Impact of the Negotiation Environment on Intra-Team Strategies' Performance

to concede using boultware strategies, the best intra-team strategies are usually found in the medium deadline, as in the very similar scenario case.

In conclusion, in these experiments we have observed that, generally, *even though the team is able to play a long deadline and the opponent plays a short deadline, the team would benefit more from playing the same type of deadline than the opponent or a slightly longer deadline.*

5.9.2.4 Team size effect on intra-team strategies

We decided to analyze the effect of the team size in the performance of intra-team strategies. Thus, we repeated the conditions in 5.9.2.2 increasing the number of team members. We excluded the RE strategy from the analysis. Since team members do not interact in RE and no unanimity level is guaranteed, the inclusion of additional team members should not affect the way in which the strategy works. The results of this experiment can be found in Figure 5.6 and 5.7. It shows the average and minimum utility of team members for teams of size $|A| = \{4, 5, 6, 7, 8\}$.

Generally, it can be observed in all of the graphics that, as the number of team members increases, *the quality of the results in terms of the minimum and the average utility is reduced.* This behavior was expected since as the number of agents increases, the set of possible agreements is reduced and the conflict inside the team and with the opponent is increased. However, the reduction in utility terms can be appreciated more easily in the minimum utility of team members. The average for the average utility of team members when $|A| = 4$ is 0.70 and 0.67 for $|A| = 8$, whereas the average for the minimum utility of team members when $|A| = 4$ is 0.48 and 0.41 for $|A| = 8$. As the number of team members increases, the contribution of each team member to the average utility is lesser, and that is the reason why the negative effect of team size on utility measures can be observed

5. INTRA-TEAM STRATEGIES FOR NEGOTIATION TEAMS IN PREDICTABLE DOMAINS

more easily in the minimum utility of team members than in the average utility of team members.

We expected that as the number of team members increased, the performance of unanimity intra-team strategies like FUM would greatly decrease compared to the performance of SSV since more team members would increase the demands of the team and make offers less interesting for the opponent. However, the difference in performance between the three strategies is approximately maintained in almost every graphic as the number of team members increases. Therefore, *team size did not have a different effect on the performance of the three intra-team strategies, affecting all of intra-team strategies equally.* The decision on which intra-team strategy should be chosen seems to be unaltered by team size.

The only exceptions to this rule are scenarios where the opponent uses conceder strategies ($\beta_{op} = C$ and $\beta_{op} = VC$) and team members' preferences are very dissimilar (rows 1 and 3, Figure 5.7). In these scenarios, we can observe that there is a special negative effect of team size on the performance of FUM with respect to the other intra-team strategies, which results in FUM being one of the worst choices when the number of team members is large. As a numeric example of the reduction in the performance of FUM, the difference in the average utility between SBV and FUM goes from approximately a 2% ($|A| = 4$) to 10% ($|A| = 8$) when $\beta_{op} = VC$ and the deadline is short, from approximately a 0% ($|A| = 4$) to 5% ($|A| = 8$) when the deadline is short and $\beta_{op} = C$, and from 3% ($|A| = 4$) to 8% ($|A| = 8$) when the deadline is long and $\beta_{op} = VC$. This phenomenon has a reasonable explanation. When the opponent uses conceder strategies FUM has greater difficulties to learn a proper issue agenda. If the number of team members increases and they are very dissimilar, the demands of team members increase, which summed up to the fact that the agenda does not properly reflect the preferences of the opponent, results in very demanding team proposals.

5.10 Conclusions

In this chapter we have focused on studying intra-team strategies for negotiation teams that negotiate with a single opponent by means of the alternating offers bilateral protocol (12). More specifically, the focus of our analysis has been intra-team strategies for negotiation domains where negotiation issues are *predictable and compatible* among team members. In this setting, there is potential for cooperation among team members since they share the same type of valuation functions (e.g., monotonically increasing or monotonically decreasing). Therefore, if one of the team members demands more from the negotiation issues and increases its welfare, it will result in the rest of team members staying at the same level of utility or increasing their respective utilities.

We have proposed four different intra-team strategies that are able to guarantee four different levels of unanimity regarding team decisions: representative (RE; no unanimity), similarity simple voting (SSV; majority/plurality), similarity borda voting (SBV; semi-unanimity) and full unanimity mediated (FUM; unanimity). Among these intra-team strategies, we have put a special emphasis on full unanimity mediated since it is able to guarantee unanimity for domains with *predictable and compatible* issues among team members. Results have shown that team members, in practice, do not have much incentive to deviate from the proposed team member behavior due to the fact that offers become very demanding and negotiations end in failure. Additionally, we have found that full unanimity mediated is robust against infiltrated agents from the opponent that attempt to manipulate the team into accepting/proposing offers that are very close to the opponent's preferences. However, the intra-team strategy is prone to be manipulated by agents from the competition whose aim is to prevent the team from reaching a deal with the opponent.

5. INTRA-TEAM STRATEGIES FOR NEGOTIATION TEAMS IN PREDICTABLE DOMAINS

Another of the goals of this chapter was studying how environmental conditions affected the performance of the different intra-team strategies. We studied how the deadline of both parties, the concession speed of the opponent, similarity among team members' preferences and team size affected the performance in terms of the minimum utility of team members, the average utility of team members and the number of negotiation rounds. *The results suggest that depending on the environmental conditions and the team performance metric, team members should select different intra-team strategies, which confirms our initial hypothesis in this thesis.* Next, we summarize some of the most important results found in this chapter:

- Generally, when the concession speed is the same for the different intra-team strategies, RE takes less numbers of negotiation rounds than SSV, which takes less number of rounds than SBV, which takes less number of rounds than FUM. The exception for this rule is when team members are very similar and the opponent uses boulware or very boulware strategies, which makes FUM usually faster than SBV.
- FUM tends to outperform the rest of intra-team strategies studied in utility terms (minimum and average utility of team members) when there is a long time to negotiate and the opponent uses either boulware or very boulware concession strategies. When the opponent uses conceder or very conceder strategies, different intra-team strategies tie in terms of the minimum and average utility of team members.
- When the team deadline is way shorter than the opponent's deadline, all of the intra-team strategies are negatively affected in the results obtained in the minimum and average utility of team members. When team members are very similar, RE becomes one of the best choices for the average utility

of team members since it is capable of ending more negotiations successfully where other intra-team strategies fail. When team members are very dissimilar, FUM and SBV tend to work better in terms of utility. However, if the opponent uses *boulware* or very *boulware* concession strategies every intra-team strategy performs equally bad and team members are encouraged to look for other negotiation alternatives.

- In situations where the team's maximum deadline is way larger than the opponent's deadline, the team should not play intra-team strategies with the maximum deadline but intra-team strategies with the same type of deadline than the opponent or a slightly longer type of deadline. Otherwise, the team performance in utility terms is not maximized due to more negotiations ending in failure.
- As the number of team members increases, the performance in utility terms of all of the intra-team strategies is negatively affected. However, in general, all of the intra-team strategies studied are equally affected by the increment in the number of team members. Thus, team size did not have an effect on the intra-team strategy that should be selected by team members to maximize the minimum or the average utility of team members.

With this chapter we have proposed intra-team strategies for negotiation domains exclusively composed by *predictable* and *compatible* issues. While these types of domains represent an important number of possible scenarios in electronic commerce, other scenarios may also exist where *unpredictable* issues are present. An issue can be considered as *unpredictable* among team members if no common order of issue values can be inferred for team members. For instance, if a group of travelers has to negotiate with the hotel on the orientation of their room,

5. INTRA-TEAM STRATEGIES FOR NEGOTIATION TEAMS IN PREDICTABLE DOMAINS

where do team members stand between rooms oriented towards the sea with respect to rooms oriented towards the pool? It may be the case that all of the team members prefer the same order of issue values, but nothing can be guaranteed since it is also possible that some team members prefer different orders of issue values. And more importantly, how can unanimity be guaranteed regarding team decisions when issues are unpredictable? This type of scenario is covered in the following chapter.




Figure 5.2: Results for very similar team members when the team has a long deadline, the opponent has a short deadline and the opponent uses concenter or very concenter tactics.

5. INTRA-TEAM STRATEGIES FOR NEGOTIATION TEAMS IN PREDICTABLE DOMAINS




Figure 5.3: Results for very similar team members when the team has a long deadline, the opponent has a short deadline and the opponent uses Boulware or very Boulware tactics.




Figure 5.4: Results for very dissimilar team members when the team has a long deadline, the opponent has a short deadline and the opponent uses concedes or very concedes tactics.

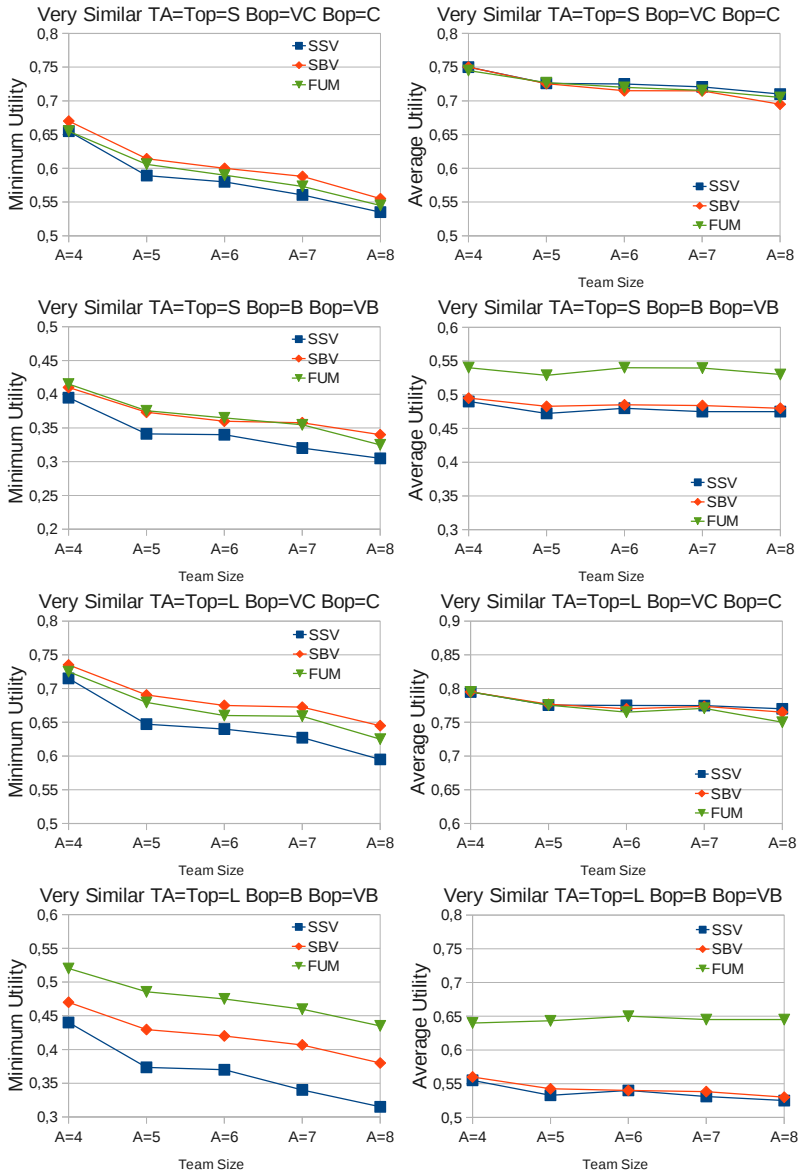


Figure 5.6: Effect of the size of the team on team performance when both parties have the same type of deadline and team members are very similar.

5. INTRA-TEAM STRATEGIES FOR NEGOTIATION TEAMS IN PREDICTABLE DOMAINS

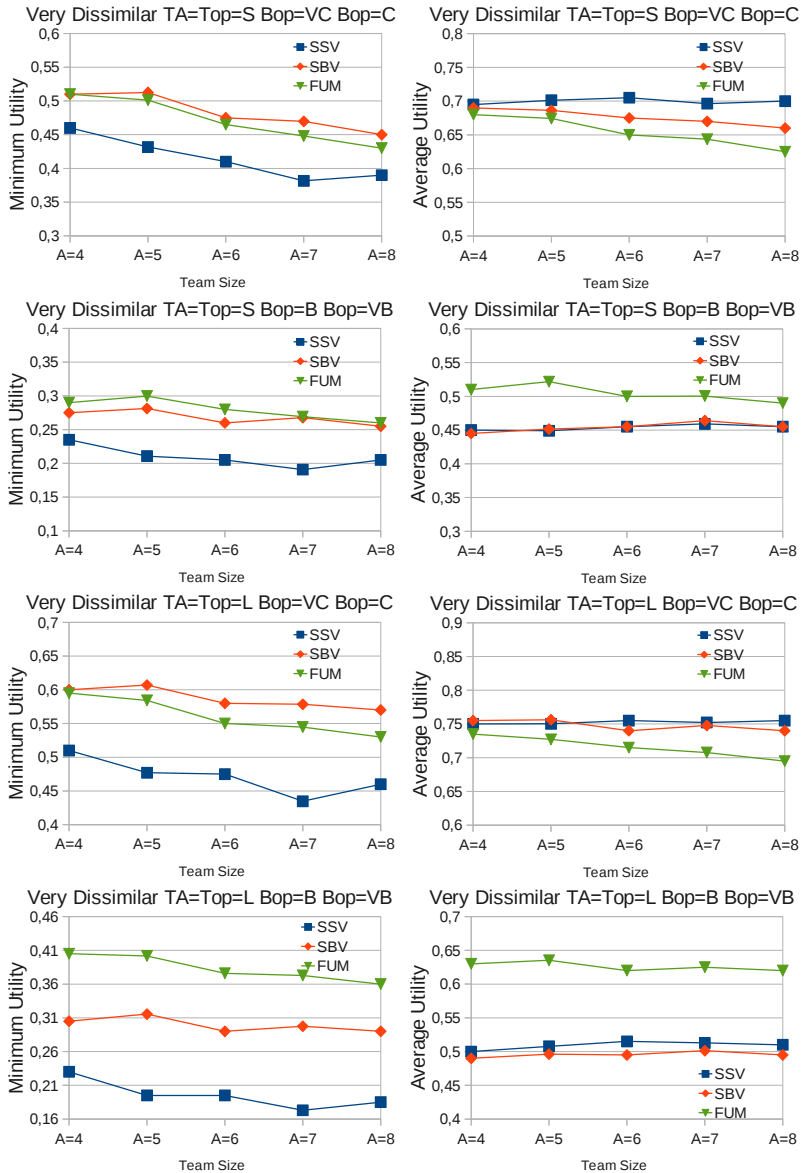


Figure 5.7: Effect of the size of the team on team performance when both parties have the same type of deadline and team members are very dissimilar.

6

Negotiation Teams in Unpredictable Domains

6.1 Introduction

The model presented in this chapter extends FUM to make it capable of reaching unanimously acceptable agreements in domains that contain both *predictable* and *unpredictable* issues. As stated, *unpredictable* issues among team members are those whose issue ranking cannot be inferred prior to the negotiation process. Therefore, the preferences of team members may or may not generate intra-team conflict. For instance, in a team of travelers that negotiates travel accommodation with a hotel, it is difficult to tell whether or not team members would prefer free internet, free drinks, gym service, or a guided tour as a complimentary activity. This extension allows us to cover a wider range of negotiation scenarios that were not initially supported by strategies proposed in Chapter 5. From the strategies presented in Chapter 6, RE, SSV and SBV support any kind of negotiation domain as long as little modifications are introduced (i.e., changes in similarity heuristics).

6. NEGOTIATION TEAMS IN UNPREDICTABLE DOMAINS

However, FUM, which is the strategy that is capable of guaranteeing unanimity regarding team decisions, cannot work in domains with *unpredictables* issues. Our interest in designing intra-team strategies capable of guaranteeing unanimity drove us to extend FUM for domains with *unpredictable* issues too.

We propose two negotiation strategies for team members: a basic negotiation strategy, and a negotiation strategy based on Bayesian learning to model its teammates' and opponent's preferences for *unpredictable* issues. The performance of the model proposed in this Chapter is evaluated in a set of environmental conditions.

We describe our general framework in Section 6.2 and the motivating scenario in Section 6.3. The intra-team protocol that allows team members to reach unanimity is detailed in Section 6.4. After that, we propose two negotiation strategies for team members in Section 6.5 and we explain why unanimity is guaranteed among team members in Section 6.6. After analyzing the experimental results in Section 6.7. Finally, we briefly highlight the conclusions of this chapter in Section 6.8.

6.2 Negotiation Setting

Most of the characteristics of the negotiation setting coincide with those of Chapter 5. To avoid redundancy, in this chapter we only highlight those features which are different from the previous chapter.

- A team mediator is present in the negotiation. The team mediator plays a key role during the negotiation since it does not only broadcast the messages between team members and opponent party but also coordinates the team members and helps to reach unanimously acceptable deals.
- Among the n different negotiation issues that compose the negotiation domain, we consider that there are issues that are *predictable and compatible*

among team members and issues that are *unpredictable* among team members. An issue j with domain D_j is compatible among team members if for each pair of team members $a, b \in A$, for each pair of issue values $v_1, v_2 \in D_j$, the following expression is true:

$$V_{a,j}(v_2) \geq V_{a,j}(v_1) \longleftrightarrow V_{b,j}(v_2) \geq V_{b,j}(v_1). \quad (6.1)$$

Hence, an issue is compatible among team members if changing its value (v_1) with another value (v_2) increases/decreases a team member's utility, then v_2 would also increase/decrease the utility for other members. An issue is *predictable* for an agent if the preference ordering of issue values is known due to the negotiation domain. Therefore, an issue is *compatible and predictable* among team members if the preferences regarding issue values are known due to the negotiation domain and increasing the utility of one of the team members by selecting one specific issue value results in other team members staying at the same utility or also increasing their respective utilities. Thus, there is potential for cooperation among teammates in *compatible and predictable* issues. Our proposed intra-team strategy takes advantage of these issues to satisfy team members as much as possible and guarantee unanimously acceptable agreements. On the other hand, an issue is *unpredictable* among team members if the preference ordering of the issue values cannot be accurately predicted and Equation 6.1 may not hold for the issue. In this framework, PR is the set of *predictable* and *compatible* issues, while UN is the set of *unpredictable* issues inside the team.

- An offer is unanimously acceptable for a team A if it is acceptable for all of the team members inside the negotiation team:

$$\forall a_i \in A, U_{a_i}(X) \geq RU_{a_i}. \quad (6.2)$$

6. NEGOTIATION TEAMS IN UNPREDICTABLE DOMAINS

The intra-team strategy proposed in this chapter assures that team members only accept those offers that are unanimously acceptable and that offers proposed to the opponent are over each team members' reservation utilities, thus, making it unanimously acceptable.

6.3 Motivating Example: Advanced Group Booking

The case of study of this chapter is similar to the one presented in the previous chapter. A group of travelers (i.e., Alice, Bob, Carol, and Dave) wants to go on a holiday together and arrange their accommodation accordingly. To do this, they may need to negotiate together with a hotel on the following issues.

- Price (p): It represents the price per night that each traveler pays to the hotel for the booking service. The value goes from 200\$, which is the minimum rate applicable by the hotel, to 400\$, which is the maximum rate found in the hotel. This negotiation issue is considered to be *predictable and compatible* among team members since all of the travelers obviously prefer low prices to high prices. Contrarily, the hotel prefers high prices to low prices.
- Cancellation fee (cf): This issue represents the amount of the final price that each friend pays if the reservation is canceled. Possible values for this negotiation issue go from 0% to 50%. This is a *predictable and compatible* issue among team members since all of the travelers prefer low cancellation fees to high cancellation fees. On the contrary, the opponent prefers high cancellation fees to low cancellation fees.
- Arranged Foods Included (af): The hotel may also offer some diets included in the deal with the travelers. The type of dietary plans included are *none*,

6.3 Motivating Example: Advanced Group Booking

breakfast, *breakfast+lunch*, *breakfast+dinner*, *lunch+dinner*, and *all*. In our negotiation scenario, we have considered that this negotiation issue is *unpredictable* among team members since preferences of team members on this issue may vary, and it cannot be assumed to be the same for each member.

- Type of room (*tr*): The four travelers can be accommodated in different types of room depending on their preferences. More specifically, the hotel offers *4 individual rooms*, *2 twin rooms*, *1 triple and 1 individual room*, or *1 apartment*. As the previous issue, it is considered that the type of room is an *unpredictable* negotiation issue among team members.
- Payment method (*pm*): The amount of money paid by the travelers may be paid by different methods. The hotel allows for the payment to be made in *cash*, via *credit card*, by *bank transfer*, in a *3 months deferred payment* through the bank, and in a *6 months deferred payment*. This negotiation issue is *unpredictable* since team members may prefer to choose different payment methods.
- Room orientation (*ro*): If possible, the team members can decide upon an orientation for the balcony of their rooms. The different options are *inner garden*, *main street*, *pool*, *sea*, and *outer garden*. This issue is also considered an *unpredictable* issue among team members.
- Free amenity (*fa*): As a token of generosity for booking as a group, the hotel offers one free service to all of the team members. More specifically, the team members can choose between *gym service*, *free wi-fi*, *1 free drink per day*, *1 free spa session*, *pool service*, *cable tv service*, and one *free guided tour*. Since the preferences of team members vary for this issue and no assumption about their preferences can be made, this issue is also considered as *unpredictable*.

6. NEGOTIATION TEAMS IN UNPREDICTABLE DOMAINS

To sum up, for this case study we have that $PR = \{p, cf\}$ and $UN = \{af, tr, pm, ro, fa\}$, with a total of 4200 different combinations of discrete issue values (af, tr, pm, ro, fa) , and two real issues (p, cf) . We assume that the team mediator knows which issues are predictable and can apply an operator such as min/max (monotonically decreasing/increasing valuation functions). For *unpredictable* issues, team members can have different types of valuation functions and the mediator does not know which issue values are better for team members. Each team member may assign different weights (i.e., priorities) to negotiation issues and the opponent's valuation functions and issue weights may be different from those of team members.

6.4 Intra-Team Protocol

We propose an intra-team protocol that is governed by a team mediator. Basically, the team mediator regulates the interactions that can be carried out among team members and, accordingly, helps team members to reach unanimous acceptable decisions inside the team. The proposed protocol is clearly differentiated into two different phases: *Pre-negotiation* and *Negotiation*. On the one hand, during the pre-negotiation, the mediator helps team members to identify potential offers that are not unanimously acceptable for every teammate. On the other hand, during the negotiation the mediator coordinates the offer proposal mechanism (composed of a voting process for *unpredictable* issues and an iterated building process for predictable issues), and the evaluation of opponent's offers.

6.4.1 Pre-negotiation Phase

The assignment of *unpredictable* issues during the negotiation is more complicated than assigning predictable issues since preferences on *unpredictable* issues may

be incompatible within the team members. The fact that team members and the team mediator do not have any prior knowledge about the incompatibility of preferences within the team, makes it more difficult for the team mediator to detect which offers are unanimously acceptable for the team. In our proposal, the mediator attempts to find the combinations of *unpredictable* issue values that will not result in an unanimously acceptable agreement under any circumstance. We say that a combination of *unpredictable* issue values will not result in an unanimously acceptable agreement when setting the most desired value for the *compatible and predictable* issues, there is at least one team member that cannot reach its reservation value. The rationale behind identifying these combinations of *unpredictable* issue values is pruning the negotiation space inside the team. Hence, team members exclusively work with combinations of *unpredictable* issue values that can result in unanimously acceptable offers.

We define an *unpredictable* partial offer X' as a partial offer that has a concrete instantiation of all the *unpredictable* issues in UN . Similarly, we consider that a complete offer X is the offer that has all of the issues in $UN \cup PR$ instantiated. The utility of an *unpredictable* partial offer is calculated as $U_{a_i}(X') = \sum_{j \in UN} w_{i,j} V_{i,j}(x'_j)$. For a team member a_i , an *unpredictable* partial offer X' is not acceptable when the sum of the utility of $U_{a_i}(X')$ and the maximum utility that can be taken from predictable issues $maxPR_{a_i} = \sum_{j \in PR} w_{i,j} \times \max_{v \in D_j} V_{i,j}(v)$ is less than its reservation value RU_{a_i} . For a team member a_i , we define the set of *unpredictable* partial offers that under any circumstance will result in an unacceptable offer as forbidden *unpredictable* partial offers, F_{a_i} (see Equation 6.3).

$$F_{a_i} = \{X' | U_{a_i}(X') + maxPR_{a_i} < RU_{a_i}\} \quad (6.3)$$

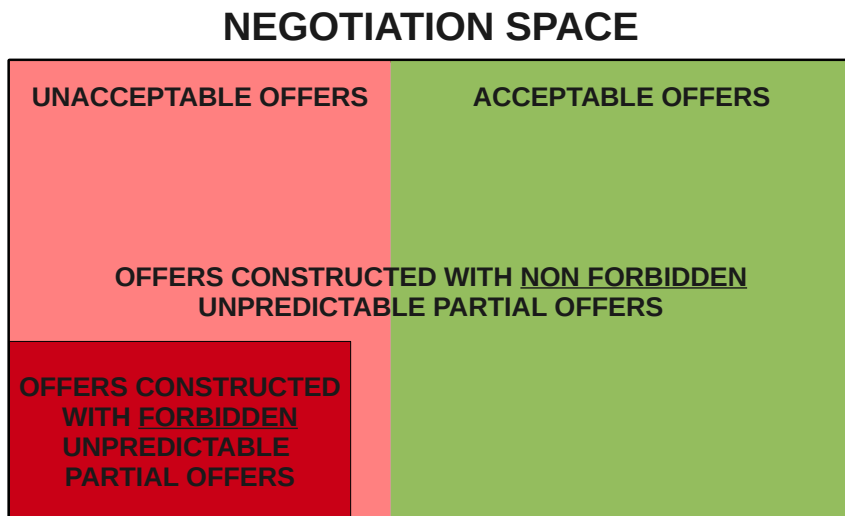


Figure 6.1: Representation of the negotiation space of an agent

It is worth noting that F_{a_i} is not the whole negotiation space that is unacceptable for a_i , but just a portion of it. In fact, some *unpredictable* partial offers that are not contained in F_{a_i} , can become unacceptable when the agent does not get the value needed from predictable issues. An intuitive idea of the negotiation space of an agent can be observed in Figure 6.1. As expected, the offers generated with forbidden *unpredictable* partial offers only end up in unacceptable offers for the agent, whereas offers generated by non forbidden *unpredictable* partial offers include acceptable and unacceptable offers for the agent. That is, an acceptable offer can be reached only by using non forbidden *unpredictable* partial offers. However, it does not mean that all of the offers generated by non forbidden *unpredictable* partial offers will be acceptable. The size of F_{a_i} may grow as the reservation utility of the agent increases. Thus, agents with high reservation utilities are expected to have larger sets of F_{a_i} than agents with low reservation utilities.

In the pre-negotiation phase, the mediator coordinates the following intra-team

6.4 Intra-Team Protocol

protocol to discover the set of forbidden *unpredictable* partial offers for the team F_A . Figure 6.2 shows an overview of the proposed intra-team protocol. Speech balloons represent public communications inside the team, while directional arrows represent private communications among two different agents. According to the proposed protocol, the team mediator initiates the pre-negotiation phase by asking each team member a_i to calculate its own set of forbidden *unpredictable* partial offers F_{a_i} (See the first frame in Figure 6.2). Each team member builds its own (forbidden) set as requested, and it is communicated to the mediator privately as depicted in the second frame (see Figure 6.2). When the mediator receives the sets from team members, it aggregates them in order to construct the set of forbidden *unpredictable* partial offers for team A, $F_A = \bigcup_{a_i \in A} F_{a_i}$. Then, as observed in the third frame (see Figure 6.2), the team mediator makes public the list of forbidden *unpredictable* partial offers of the team F_A . It should be stated that, since any *unpredictable* partial offer in this set will prevent one of the team members from reaching its reservation utility, the team is not allowed to generate an offer involving any of the partial offers in F_A . After the team mediator has shared F_A with team members, the negotiation phase starts.

The set of forbidden *unpredictable* partial offers F_A may be a useful tool. The team mediator may be able to detect prior to the negotiation if no unanimous acceptable agreement is possible among team members (i.e., when F_A covers all of the possible partial offers). Additionally, the information gathered from each agent may facilitate team formation under the rationale that those F_{a_i} that are the most similar may reduce intra-team conflict. The use of this information in team formation algorithms is considered as a future line of work.

6. NEGOTIATION TEAMS IN UNPREDICTABLE DOMAINS

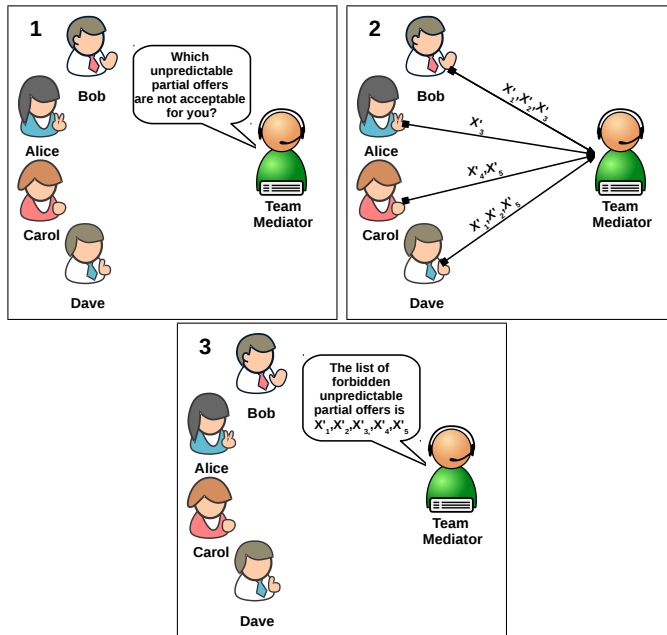


Figure 6.2: Overview of the intra-team protocol carried out during the pre-negotiation

6.4.2 Negotiation Phase

In the negotiation, two mechanisms are carried out at each round: a mechanism for deciding to accept/reject the opponent's offer (Evaluation of Opponent's Offer), and a mechanism for proposing an offer to the opponent (Offer Proposal). For the former, a unanimity voting process is employed, while for the latter an offer building process is governed by the team mediator. Next, we describe both processes in detail.

6.4.2.1 Evaluation of Opponent's Offer

This mechanism is carried out each time the team mediator receives an offer from the opponent. Since the main goal of the proposed intra-team strategy is achieving unanimously acceptable agreements for the team, a unanimity voting is carried out to decide whether or not the opponent's offer is acceptable for the team. With this mechanism, as long as one of the team members is not satisfied with the opponent's offer, the offer is not accepted by the team, precluding the team from reaching agreements that are not unanimously acceptable. The intra-team protocol used for this mechanism goes as follows. First, the team mediator receives the offer X^t from the opponent at time t . If X^t involves any forbidden *unpredictable* partial offer in F_A , the opponent offer is automatically rejected. However, the opponent's offer is also informed to team members in order to allow each team member to process the new information leaked by the opponent if they see it necessary. Otherwise, if the combination of *unpredictable* issue values is not in F_A , the mediator makes the opponent's offer public among team members and starts an anonymous voting process (i.e., votes are communicated privately to the team mediator and only the final result of the voting is publicly communicated to team members). Each team member a_i states to the mediator whether he is willing to accept X^t (positive vote) or to reject it (negative vote) at that specific instant. Since we desire to guarantee unanimity, the offer is only accepted if all of the team members emit a positive vote. Otherwise, the offer is rejected and a counter-offer is proposed as explained in Section 6.4.2.2.

6.4.2.2 Offer Proposal

Proposing an offer to the opponent is a complex task, since the space of offers may be huge and the preferences of the team members should be reflected in the offer sent to the opponent. Moreover, the offer sent should be unanimously acceptable

6. NEGOTIATION TEAMS IN UNPREDICTABLE DOMAINS

for team members. The process is divided into two sub-phases: constructing an *unpredictable* partial offer, and setting up predictable issues. In both phases, the team mediator acts according to Algorithm 6. We explain both processes in detail below.

- **Constructing an unpredictable partial offer:** The first step is proposing an *unpredictable* partial offer, a partial offer which has all of the *unpredictable* issues instantiated. Since team members know from F_A the list of *unpredictable* partial offers that will not result in unanimously acceptable offers under any circumstance, any offer proposed by the team should avoid being constructed from *unpredictable* partial offers found in F_A . The method used to propose offers to the opponent relies on the fact that *unpredictable* issues are those issues where more intra-team conflict may be present, whereas full potential for cooperation is present in *predictable and compatible* issues. Hence, in order to build an offer to be sent to the opponent, it seems more appropriate to set *unpredictable* issue values first and then, depending on the remaining needs of team members, allow team members to set *compatible and predictable* issues as they require to reach their demands. The proposed mechanism for the first part, proposing an *unpredictable* partial offer, is based on voting and social choice. The voting process goes as follows.

1. The mediator asks each team member to anonymously propose one *unpredictable* partial offer $X'_{a_i}{}^t$.
2. Each team member privately sends its proposal to the team mediator, who gathers all of the proposals in a list that will be later sent to team members. If any *unpredictable* partial offer proposed by a_i is contained in F_A , the mediator automatically ignores this proposal.

3. Once all of the proposals have been gathered, the team mediator makes public the list of proposal UPO'^t among team members and opens a Borda scoring process (131) on proposed candidates.
 4. Each team member anonymously scores the *unpredictable* partial offers and anonymously sends the scores to the team mediator. The team mediator sums up scores and selects the candidate with the highest score $X'_A{}^t$, making it public among team members. This candidate, an *unpredictable* partial offer, is the base for the full offer that is to the opponent.
- **Setting up predictable and compatible issue values:** Once *unpredictable* issues have been set, it is necessary to set *predictable and compatible* issues to construct a complete offer. As it has been stated along this thesis, there is full potential for cooperation among team members in these issues, since increasing the utility of one of the team members by selecting one issue value will result in the other team members staying at the same utility or increasing their utility. The selected *unpredictable* partial offer does not satisfy equally the needs of all the team members. Nevertheless, team members can make use of *predictable and compatible* issues to satisfy their remaining needs while not generating conflict inside the team. To complete the partial offer $X'_A{}^t$, the iterative mechanism proposed in FUM (see Chapter 5) is used to build the final offer issue per issue.

6.5 Team Members' Strategies

The team mediator defines the coordination mechanisms and the rules of the game inside the team. However, each team member's internal strategy has a great effect

6. NEGOTIATION TEAMS IN UNPREDICTABLE DOMAINS

```

/*Proposing an unpredictable offer*/;
Send (Ask for  $X'_{a_i} \longrightarrow \forall a_i \in A$ );
Receive( $X'_{a_i} \longleftarrow \forall a_i \in A$ );
 $UPO'^t = (\bigcup_{a_i \in A} X'_{a_i}) - F_A$ ;
Send (Score borda  $UPO'^t \longrightarrow \forall a_i \in A$ );
Receive ( $score_{a_i} \longleftarrow \forall a_i \in A$ );
 $X'_A{}^t = \operatorname{argmax}_{X' \in UPO'^t} \sum_{a_i \in A} score(a_i, X')$ ;
agenda=build_predictable_agenda();  $A' = A$ ;
/*Setting predictable issues*/;
foreach  $j \in agenda$  do
    Send (Needed value  $x_{a_i,j}$ , for  $X'_A{}^t \longrightarrow \forall a_i | a_i \in A'$ );
    Receive ( $x_{a_i,j} \longleftarrow \forall a_i | a_i \in A'$ );
    if monotonically_increasing( $j$ ) then
        |  $x_j = \max_i x_{a_i,j}$ ;
    end
    else
        |  $x_j = \min_i x_{a_i,j}$ ;
    end
     $X'_A{}^t = X'_A{}^t \cup \{x_j\}$ ;
    Send (Acceptable  $X'_A{}^t? \longrightarrow \forall a_i | a_i \in A'$ );
    foreach  $a_i \in A'$  do
        | Receive ( $ac'_{a_i}(X'_A{}^t) \longleftarrow a_i$ );
        | if  $ac'_{a_i}(X'_A{}^t) = true$  then  $A' = A' - \{a_i\}$ ;
    end
    if  $A' = \emptyset$  then break;
end
foreach  $j \in agenda \wedge issue\_not\_set(j)$  do
    |  $x_j = maximize\_for\_opponent(j)$ ;
    |  $X'_A{}^t = X'_A{}^t \cup \{x_j\}$ ;
    |  $X_A^t = X'_A{}^t$ ;
end

```

Algorithm 6: Pseudo-algorithm for the offer construction from the point of view of the mediator. Send (*message* \longrightarrow *condition*) means that *message* is sent to every agent that fulfills *condition*

on team dynamics. In this chapter, we propose two types of negotiation strategy for team members. According to the first strategy (i.e., our basic team member), the team member only considers its own utility function to take decisions. In the second strategy, team members also take into account its teammates' and opponent's preferences by employing Bayesian learning (i.e., Bayesian team member).

6.5.1 Basic Strategy for Team Members

Since negotiations are time-bounded, we consider that team members have to perform some kind of concession if an agreement is to be found. For this purpose we have designed basic team members as agents whose demands are controlled by an individual and private concession tactic. More specifically, the concession strategy for a team member a_i is based on time-based tactics (18, 57). This concession strategy estimates the utility demanded by a_i at time t by using the formula in Equation 6.4, where RU_{a_i} is its reservation utility, T is the negotiation deadline, and β_{a_i} is the concession speed, which determines how fast the agent's demands are lowered towards the reservation utility.

$$s_{a_i}(t) = 1 - (1 - RU_{a_i}) \times \left(\frac{t}{T}\right)^{\frac{1}{\beta_{a_i}}} \quad (6.4)$$

6.5.1.1 Evaluation of Opponent's Offer

Given an offer X^t proposed by the opponent at round t , the team member emits a positive vote for this offer in the unanimity voting process if it reports a utility which is greater than or equal to its current demands $s_{a_i}(t)$. Otherwise, a negative vote is emitted.

$$ac_{a_i}(X^t) = \begin{cases} true & \text{if } s_{a_i}(t) \leq U_{a_i}(X^t) \\ false & \text{otherwise} \end{cases} \quad (6.5)$$

6. NEGOTIATION TEAMS IN UNPREDICTABLE DOMAINS

6.5.1.2 Offer Proposal

As documented in Section 6.4.2, team members interact at three points during the offer proposal. First, they propose an *unpredictable* partial offer to the team mediator. Since each team member a_i has its demands regulated by a time-based tactic, when proposing an *unpredictable* partial offer to the mediator at instant t , the proposed *unpredictable* partial offer $X'_{a_i}{}^t$ fulfills:

$$X'_{a_i}{}^t \notin F_A \wedge (U_{a_i}(X'_{a_i}{}^t) + \max PR_{a_i} \geq s_{a_i}(t)) \quad (6.6)$$

Hence, a_i selects an *unpredictable* partial offer which is not forbidden inside the team (since it will be ignored by the team mediator) and whose utility allows him to achieve or surpass its current demands at time t . This way, the team member assures that if the proposed *unpredictable* partial offer is the winner of the Borda voting process, it can reach its current demands. However, one should be aware that many *unpredictable* partial offers may fulfill Equation 6.6. Therefore, it is necessary to select one of them as the proposed candidate. Being our basic team member, from the set of partial offers that fulfill Equation 6.6, a team member selects one of the candidates randomly.

The second time that a team member interacts with the team mediator is for scoring *unpredictable* partial offers that have been proposed by team members. For scoring candidate partial offers in the Borda voting process, a basic team member orders the candidates according to the partial utility reported by each of the candidates. That is, the team member assigns the highest score to the partial offer whose utility is the highest for itself, and the second highest score to the partial offer whose utility is the second best one, and so forth.

Finally, team members also interact with the mediator during the iterative mechanism used to set *predictable and compatible* issues. In this part of the intra-

team negotiation, team members employ the same strategy describe in FUM (see Chapter 5).

6.5.2 Bayesian-based Strategy for Team Members

The Bayesian negotiation strategy for a team member is based on modeling the team's (as a whole) and its opponent's preferences on *unpredictable* issues. For this purpose, two Bayesian models are employed to predict whether *unpredictable* partial offers are acceptable for both teammates and the opponent. One of the Bayesian models is employed to capture the preferences of the team on *unpredictable issues*, whereas the other is used for capturing the preferences of the opponent on *unpredictable* issues. The strategy used to evaluate opponent's offer is the same than the one described in the basic strategy.

6.5.2.1 Bayesian Learning

Bayesian learning is a probabilistic learning method based on Bayes' theorem (58). Given a certain set of hypothesis H and some observation e , Bayesian learning attempts to compute the probability $p(h|e)$ that a certain hypothesis h is true after observing e . In our case, we want to determine whether or not the proposed offer will be acceptable for the opponent (or the team) ($H=\{acc, \neg acc\}$) given a certain *unpredictable* partial offer ($e = X'^t$).

Since we assume that there is no interdependence among negotiation issues, we can consider that each negotiation issue contributes individually to the acceptability of an offer/*unpredictable* partial offer. Thus, applying Bayes' theorem under independence assumption we have:

$$p(acc|X'^t) = \frac{p(acc) \prod_{j \in UN} p(x_j|acc)}{p(acc) \prod_{j \in UN} p(x_j|acc) + p(\neg acc) \prod_{j \in UN} p(x_j|\neg acc)} \quad (6.7)$$

6. NEGOTIATION TEAMS IN UNPREDICTABLE DOMAINS

where $p(acc)$ is the prior probability for an *unpredictable* partial offer to be acceptable, $p(\neg acc)$ is the prior probability for an *unpredictable* partial offer to be non-acceptable, and $p(x_j|acc)$ is the conditional probability for the value of the j -th issue to be part of an acceptable offer.

We consider positive examples S_{acc} as those examples that correspond to the acceptable hypothesis (acc) and negative examples $S_{\neg acc}$ as those examples that correspond to the not acceptable hypothesis ($\neg acc$). For the opponent's model, we employ *unpredictable* partial offers that have appeared in opponent's offers as positive examples, and *unpredictable* partial offers that appear in offers rejected by the opponent as negative samples. For the team's model, we use F_A and those opponent's offers rejected by team members as the set of negative examples. Winners in the Borda votings (i.e., *unpredictable* partial offers contained in offers sent to the opponent) are considered as positive examples.

When computing $p(x_j|h)$, we calculate the proportion between the number of times that x_j appeared in hypothesis h (acc or $\neg acc$), and the total number of examples for h :

$$p(x_j|h) = \frac{\#\{x_j \in S_h\}}{|S_h|} \quad (6.8)$$

6.5.2.2 Offer Proposal

Up to this point, we have explained how the team members model other team members' and the opponent's preferences by means of Bayesian models. However, we have not explained yet where these models come into play. Basically, Bayesian models are employed to help in the selection of the *unpredictable* partial offer that is proposed to the other team members. If we remember from the basic team member formalization, team members propose at t *unpredictable* partial offers in the set defined in Equation 6.6. Bayesian models help to the select a candidate from such set.

However, it is reasonable to think that in the first interactions Bayesian models do not accurately represent other agents' preferences. For that purpose, a team member invests part of the negotiation time t_{exp} in exploring the negotiation space and collecting information regarding the opponent's and the team's preferences. As long as the negotiation process has not surpassed t_{exp} , the team member just selects randomly one of candidate *unpredictable* partial offers as basic team members do. Meanwhile, the Bayesian models are continuously updated with the new information that becomes available during the negotiation. After reaching the time threshold, the team member starts to use Bayesian models in order to select the *unpredictable* partial offer. The heuristic used in the selection of the candidate is proposing an *unpredictable* partial offer that is both acceptable for the team and the opponent. The model has an additional parameter named p_{esc} . It represents the probability of avoiding the Bayesian models (using the random proposal model) when the negotiation time has gone beyond t_{exp} . This parameter is included in the model in order to: (i) explore further the negotiation space; (ii) escape from local optima induced by inaccurate Bayesian models (e.g., wrong samples, limited number of samples, etc.). We can formalize the selection as follows:

$$X'_{a_i} = \begin{cases} \underset{X \in B}{\operatorname{argmax}} (w_A p_A(\operatorname{acc}|X) + w_{op} p_{op}(\operatorname{acc}|X)) & \text{if } \operatorname{rand} \leq p_{esc} \wedge t \geq t_{exp} \\ \operatorname{select_random_partial_offer}(B) & \text{otherwise} \end{cases} \quad (6.9)$$

where B is the set of candidate *unpredictable* partial offers that fulfill Equation 6.6, rand is a random number in $[0,1]$, $p_A(\operatorname{acc}|X)$ is the probability for a candidate *unpredictable* partial offer to be acceptable for the team, $p_{op}(\operatorname{acc}|X)$ is the probability for the candidate *unpredictable* partial offer to be acceptable for the opponent, and w_A and w_{op} ¹ represent the weights or importances given to the

¹ $w_A + w_{op} = 1$

6. NEGOTIATION TEAMS IN UNPREDICTABLE DOMAINS

acceptability of the *unpredictable* partial offer for the team and the opponent, respectively. Varying these weights allow team members to show different behaviors depending on their inclination to satisfy either the team or the opponent with the *unpredictable* partial offer.

6.6 Unanimously Acceptable Proof

As stated, one of our research goals is proposing negotiation team models that are able to guarantee unanimously acceptable team decisions. Next, we show that under the assumption of rationality¹, team members are able to achieve unanimously acceptable final agreements, if an agreement is found. For that matter, let us employ *reductio ad absurdum* (reduction to absurdity).

If X is the final agreement, let us suppose that Equation 6.2 (unanimously acceptable) is violated in a negotiation: unanimity is not reached because a_i obtained a utility below its reservation utility.

$$\exists a_i \in A, U_{a_i}(X'^t) + \sum_{j \in PR} w_{i,j} V_{i,j}(x_j) < RU_{a_i} \quad (6.10)$$

The final agreement is found when (1) team members accept an opponent's offer or (2) the opponent accepts a team's offer. Next, we show that in both cases, Equation 6.10 is never true.

1. When the team members accept an opponent's offer, a unanimity voting process has been carried to decide whether or not to accept the final offer. The offer is only accepted if all of the team members have emitted a positive vote. Since a rational agent a_i would never have incentive to emit a positive vote if the offer reports a utility below its reservation utility, this scenario is never true.

¹Rational agents seek to improve their current welfare. Thus, they would not take actions that lead to utilities below their reservation utilities

6.6 Unanimously Acceptable Proof

2. When the opponent accepts a team's offer X , the offer has been proposed by the intra-team mechanism mentioned in Section 6.4.2.2. The offer can be decomposed into an *unpredictable* partial offer X'^t and an instantiation of predictable issues. The team member a_i is not able to get over its reservation utility if and only if $X'^t \in F_{a_i}$ or when $X'^t \notin F_{a_i}$ and a_i could not get what it demanded in predictable issues. A rational agent has no incentive to exclude a forbidden *unpredictable* partial offer X'^t when declaring F_{a_i} . Since $F_A = \bigcup_{a_i \in A} F_{a_i}$ and the mediator ignores *unpredictable* partial offers in F_A , an *unpredictable* partial offer X'^t that forms a team offer is never in F_{a_i} . If $X'^t \notin F_{a_i}$ then the agent can accomplish to satisfy the following expression $U_{a_i}(X'^t) + \max PR_{a_i} \geq RU_{a_i}$. a_i could not get over its reservation value because he could not demand the most of predictable issues. However, when the team mediator aggregates predictable issues inside the team, the team mediator selects the highest value for team members in the the list of values proposed by them. Thus, a_i can obtain the maximum utility from predictable issues. If an agreement is found, Equation 6.10 is never true.

Since both possible scenarios are never true under our initial assumption, we have shown by *reduction ad absurdum* that, if a final agreement is found, it is unanimously acceptable. Another research issue is the presence of exaggerating team members (i.e., agents that exaggerate their preferences to get the most from the negotiation). In our setting, even if team members exaggerate and decide to include in F_{a_i} *unpredictable* partial offers that are acceptable but report low utility or they demand more than they need from predictable issues, if a final agreement is found it will be unanimous among team members. However, by doing so, they may be pruning negotiation space and lowering the probability of finding agreements. This is an interesting situation that we plan to study in the future.

6. NEGOTIATION TEAMS IN UNPREDICTABLE DOMAINS

6.7 Experiments

In this section, we explore the behavior of the proposed negotiation model in different environments. In order to assess the performance of the proposed negotiation approach, we have performed three different experiments. All of the experiments carried out have been done in the negotiation domain introduced in Section 6.3. The first experiment aims to compare the performance of the proposed model with basic and Bayesian team members against other negotiation team models in this thesis. The comparison is carried out in scenarios with different degrees of team's preference dissimilarity. In the second experiment we study how the weights w_A and w_{op} , which control the importance given to the preferences of the team and the opponent in the *unpredictable* partial offer proposed to teammates, impact the performance of the proposed model when team members employ the Bayesian strategy. Finally, we conduct an experiment to study the effect of team members' reservation utility on the performance of the proposed negotiation model.

The implementations of this chapter have been carried out in GENIUS (150), a well-known simulation framework for negotiations. It supports simulation of sessions and tournaments based on bilateral negotiations. Users are able to design their own agents and test them against a wide variety of different agents designed by the community. The framework provides information critical for analysis (e.g., utility, Pareto optimality, etc.) which is extremely useful for research tasks. Moreover, the use of GENIUS as a testbed for bilateral negotiations is testified by its use in the annual automated negotiating agent competition (ANAC) (161). The ANAC competition provided GENIUS with a large repository of agents. The integration of ABNT in GENIUS additionally facilitates the following objectives: (i) the use of GENIUS in ANAC has provided with wide variety of opponent agents; (ii) GENIUS is a consolidated testbed among the agent community. Thus, the in-

clusion of ABNT inside GENIUS can facilitate research on agent-based negotiation teams by other scholars, and even give room to a future negotiating competition involving teams.

6.7.1 First Experiment: Measuring Model Performance

As stated above, in this first set of experiments we study the performance of the proposed model. The study is carried out with an emphasis on observing the performance difference in settings having different degrees of preference dissimilarity among team members. In our experiment, we consider two configurations for the proposed model: when all the team members use the basic strategy (i.e., basic), and when all the team members employ the Bayesian strategy proposed in this chapter (i.e., Bayesian). We also included the Similarity Borda Voting model (i.e., SBV, see Chapter 5) in our experiment. The reason to include this intra-team strategy in our study is due to the fact that it is capable of achieving similar results to FUM under certain circumstances. In order to adapt this approach for domains with *unpredictable* issues, we use a similarity heuristic that uses Euclidean distance for real/integer issues and string matching for other types of issues.

In our framework, we are also interested in studying how team members' preferences impact on the performance of team negotiation models. The team dissimilarity measure is calculated as shown in Chapter 5. For this experiment, we decided to explore teams whose preferences are very dissimilar, teams whose preferences are very similar, and teams with an average degree of similarity/dissimilarity (i.e., average similarity). For the scenario of very dissimilar preferences, 9 negotiation cases were randomly generated (i.e., a combination of 3 different negotiation teams consisting of four team members with 3 different opponents), while 9 negotiation cases were randomly generated for the very similar preferences scenario (i.e., a combination of 3 different negotiation teams consisting of four team members

6. NEGOTIATION TEAMS IN UNPREDICTABLE DOMAINS

with 3 different opponents) and 12 negotiation cases were randomly generated for the average similarity scenario (i.e., a combination of 4 different negotiation teams consisting of four team members with 3 different opponents).

For the models proposed in this chapter (i.e., basic and Bayesian members), there are several parameters that need to be fixed. Initially, the reservation utility of each team member was set to $RU_{a_i} = 0.5$ to simulate negotiation scenarios where team members have outside options besides the current negotiation. Additionally, for each team member (i.e., basic, Bayesian and SBV) the concession speed was randomly selected from a uniform distribution of bouldware strategies $\beta_{a_i} = U[0.5, 1]$. In the case of Bayesian team members, the time of exploration was set to $t_{exp} = 70\%$ and the probability of escape after the exploration phase was set to $p_{esc} = 30\%$ ¹. Initially, we set Bayesian team members to care equally about the probability for *unpredictable* partial offers to be accepted by the team and the opponent $w_A = w_{op} = 0.5$.

Since the model presented in this chapter has been implemented in GENIUS, we are able to study team performance against state-of-the-art opponents. We decided to test the negotiation team models against different families of opponents. More specifically, we followed the categorization of negotiation strategies proposed by Baarslag et al. (151), which divides negotiation strategies into four categories: competitors, conceders, matchers, and inverters. On the one hand, competitors hardly concede (independently of opponent behavior), whereas conceders yield independently of the opponent behavior. On the other hand, matchers concede when they perceive that the opponent concedes, and they do not concede if they perceive that the other party does concede. Inverters respond by implementing the opposite behavior shown by the other party. According to the practical categorization of Baarslag et al., we selected Agent K (162), winner of negotiating agent

¹These values were found to work well in practice for almost every scenario tested.

competition in 2010, as competitor agent for our tests, Nice Tit-for-Tat (163), participant in 2011’s negotiating competition, as a matcher strategy, a time-based tactic (18, 57) with $\beta_{op} = 0.2$ (i.e., very bouldware) as an inverter strategy, and a time-based tactic (18, 57) with $\beta_{op} = 2$ (i.e., concenter) as a concenter strategy.

Following the type of setting used in the annual agent competition, the negotiation time is set to $T = 180$ seconds. Each opponent strategy is faced against each negotiation team model in every possible negotiation case. A total of 20 repetitions are done per negotiation case in order to capture stochastic variations in negotiation strategies. Therefore, $3 \times 3 \times 3 \times 4 \times 20 = 2160$ (team preference profiles \times opponent preference profiles \times team negotiation models \times opponent strategies \times repetitions) negotiations were simulated in the very similar scenario, 2160 negotiations were simulated in the very dissimilar scenario, and 2880 negotiations were simulated in the average similarity scenario. Information was gathered regarding the joint utility of the team¹ in the final agreement, and the opponent utility (included to see the effect of Bayesian models) in the final agreement. The results of the experiment can be found in Table 6.1. An ANOVA test ($\alpha = 0.05$) with a Boniferroni post-hoc analysis was carried out to assess statistical differences among the different measures gathered. Those measures that are statistically the best configurations for each column are highlighted in bold style.

It can be observed that when *team members’ preferences are very similar, both basic and Bayesian models are statistically equivalent to each other and they are statistically better than SBV with respect to the average team joint utility*. Basic and Bayesian models outperform SBV with respect to the average team joint utility since they are able to guarantee unanimously acceptable agreements, while SBV does not guarantee such condition. The reason why Bayesian models do not give

¹We consider the joint utility of the team to be the product of the utilities of the team members.

6. NEGOTIATION TEAMS IN UNPREDICTABLE DOMAINS

Very Similar								
	Agent K		Nice Tit-for-Tat		Very Boulware		Conceder	
	T. Joint	Op.	T. Joint	Op.	T. Joint	Op.	T. Joint	Op.
SBV	0.181	0.743	0.150	0.694	0.184	0.755	0.552	0.482
Basic	0.259	0.683	0.173	0.760	0.223	0.696	0.561	0.468
Bayesian	0.263	0.690	0.164	0.746	0.224	0.695	0.557	0.472
Average Similarity								
	Agent K		Nice Tit-for-Tat		Very Boulware		Conceder	
	T. Joint	Op.	T. Joint	Op.	T. Joint	Op.	T. Joint	Op.
SBV	0.168	0.629	0.137	0.562	0.170	0.598	0.324	0.428
Basic	0.211	0.574	0.141	0.691	0.210	0.585	0.386	0.414
Bayesian	0.248	0.583	0.158	0.669	0.224	0.574	0.390	0.414
Very Dissimilar								
	Agent K		Nice Tit-for-Tat		Very Boulware		Conceder	
	T. Joint	Op.	T. Joint	Op.	T. Joint	Op.	T. Joint	Op.
SBV	0.07	0.522	0.160	0.457	0.128	0.547	0.257	0.430
Basic	0.174	0.397	0.184	0.572	0.254	0.505	0.472	0.367
Bayesian	0.209	0.457	0.196	0.559	0.271	0.489	0.475	0.367

Table 6.1: Average joint Utility (T. Joint) and average opponent Utility (Op) for the first set of experiments.

an advantage over the basic model in the very similar scenario can be explained due to the fact that, since team members are very similar, there is no necessity to carry out team modeling. If an offer is good for one of the team members, it will probably be good for other teammates. Additionally, since team members' preferences are similar, it is also easier for opponents to learn the team preferences, which helps to find better agreement even if only the opponent uses opponent modeling. The basic model, the Bayesian model and SBV perform statistically equal in terms of the average team joint utility only when the opponent is a conceder. Since the opponent concedes rapidly, the three team negotiation models are able to get similar results.

If we observe the opponent utility when *team members' preferences are very similar*, one can realize that the opponent may be benefited if the team members

employ the SBV model. This is observable in the case of the competitor agent (i.e., Agent K) and the inverter agent (i.e., very boultware), which exploit the team using the SBV model and that is reflected in a statistically lower team joint utility and a statistically higher opponent utility. We could not find this pattern in the case of matchers and conceders since they should not tend to exploit team members.

As conflict is introduced inside the team by making team members' preferences more dissimilar (i.e., very dissimilar and average similarity scenarios), it can be observed that usually team members get the statistically highest average team joint utility by employing the Bayesian strategy for modeling the team's preferences. In this case, the teammates' preferences are no longer very similar and some sort of modeling mechanism is needed in order to guide the intra-team negotiation towards agreements that are good for all of the team members. The only exception for this pattern is found in the concenter case, where the performance in terms of the team joint utility was found to be statistically equivalent among basic and Bayesian models. Taking a closer look at the negotiation traces, we observed that, in all of the negotiations, the exploration time t_{exp} was never surpassed. Therefore, Bayesian models do not get to be used. In fact, the average negotiation time against concenter agents was 62 seconds in the very similar scenario, 76 seconds in the average similarity scenario, and 88 seconds in the very dissimilar scenario. All of them are below the threshold of 126 seconds delimited by t_{exp} . As a result of reaching an agreement early, the team members have not used their Bayesian model while generating their proposals.

In the case of the average opponent's utility, we found a similar pattern to the one found in the very similar scenario, where the opponent exploit teams using the SBV model if a competitor or matcher strategy is played. Additionally, we also found that when teammates' preferences are very dissimilar and the opponent uses a concenter strategy, the opponent also gets a statistically higher utility if team

6. NEGOTIATION TEAMS IN UNPREDICTABLE DOMAINS

members employ the Similarity Borda Voting model. Since no full unanimity is guaranteed by SBV, the negotiation space considered by team members is larger than that considered by the basic and the Bayesian model. The opponent may benefit from those new portions of negotiation space since the agreement space may get larger.

As for the comparison between basic team members and Bayesian team members in terms of the opponent utility, the Bayesian model's performance is slightly better than the performance of the basic model when facing Agent K. Nevertheless, we could not observe this behavior against Nice Tit-for-Tat and Very Boulware. We analyzed the trace of different negotiations against Nice Tit-for-Tat and Very Boulware opponents. In the former case, we could observe that close to the end of the negotiation the Nice Tit-for-Tat opponent had only sent 5 different *unpredictable* partial offers in a domain that has 4200 different *unpredictable* partial offers. This behavior results in scarce information for any learning mechanism. In the case of negotiations against Very Boulware agents, one should consider that the Very Boulware strategies concede only towards the end of the negotiation and, most of the time, the aspirations are very high. On top of that, the Very Boulware implementation in GENIUS sends any offer with the demanded utility without considering any other information. Thus, most of the samples gathered by the Bayesian classifier when facing Very Boulware agents correspond to offers with high demands where usually only the best issue values appear. Other issue values do not appear in the samples or they have their frequency misrepresented with respect to the utility that they actually report. Therefore, any learning mechanism based on frequencies (i.e., Bayesian) would have difficulties in learning these opponents' preferences. This fact explains in part the reason why Bayesian classifiers improve the joint utility of the team over the basic model, but they do not improve the utility obtained by the opponent.

In conclusion, in this first experiment we have found that team members benefit in terms of the team joint utility by using the proposed model in this chapter with respect to other approaches like SBV. Additionally, we have shown that as long as some intra-team conflict is present in teammates' preferences, team members benefit from playing the Bayesian strategy proposed in this chapter.

6.7.2 Second Experiment: Analyzing the impact of parameters for the proposal of unpredictable partial offers

If we recall from Section 6.5.2 there are two weight parameters that control how important the opponent's and team's preferences are while generating the *unpredictable* partial offer (respectively w_{op}, w_A). w_A represents how important it is for us to make an *unpredictable* partial offer acceptable for the team, whereas w_{op} represents how important it is for us to make an *unpredictable* partial offer acceptable for the opponent. The use of these weights is not trivial, since one should consider that, it only refers to the acceptability of the *unpredictable* partial offer by one of the two parties. A complete offer is composed by the predictable and *unpredictable* issues. Therefore, using a high value of w_{op} may not have the desired effect on the opponent unless *unpredictable* issues are important for the opponent. In this second experiment we explore the impact of these weights in a wide variety of situations. More specifically, we consider situations where the team gives more importance to *unpredictable* partial issues than the opponent, situations where the opponent gives more importance to *unpredictable* partial issues than the team, and situations where both team and the opponent give the same importance to *unpredictable* partial issues.

To assess the importance given by an agent to *unpredictable* partial issues, we

6. NEGOTIATION TEAMS IN UNPREDICTABLE DOMAINS

consider the sum of *unpredictable* issue weights in utility functions.

$$I_{a_i} = \sum_{j \in UN} w_{i,j} \quad (6.11)$$

We consider that when $I_{a_i} \in [0.0, 0.33]$ the agent a_i gives low importance to *unpredictable* issues, when $I_{a_i} \in [0.33, 0.66]$ it gives average importance to *unpredictable* issues, and when $I_{a_i} \in [0.66, 1.0]$ the agent gives high importance to *unpredictable* issues. We generated 8 random negotiation cases where team members give a high importance to *unpredictable* issues and the opponent gives low (4 cases) and average (4 cases) importance to *unpredictable* issues, 8 different randomly generated negotiation cases where team members give a low importance to *unpredictable* issues and the opponent gives average (4 cases) and high (4 cases) importance to *unpredictable* issues, and 12 negotiation cases where the opponent and the team give the same importance to *unpredictable* issues (4 cases where both give low importance, 4 cases where both give average importance, and 4 cases where both give high importance).

We tested three different configurations for Bayesian team members: standard Bayesian team members that give the same importance to the acceptability of the *unpredictable* partial offer by the opponent and the team $w_A = w_{op} = 0.5$ (Normal), Bayesian team members that give more importance to the acceptability of the *unpredictable* partial offer by the opponent $w_A = 0.25$ $w_{op} = 0.75$ (Opponent Oriented), and Bayesian team members that give more importance to the acceptability of the *unpredictable* partial offer by the team $w_A = 0.75$ $w_{op} = 0.25$ (Team Oriented). As for the opponent's strategy, we selected Agent K since we observed that Bayesian classifiers are able to learn good models from offers sent by Agent K.

For each negotiation case, we repeated the negotiation 20 times in order to capture stochastic variations in strategies. Therefore, a total of 1680 negotiations

were carried out in this experiment. The results of this experiment can be observed in Table 6.2. It shows the average of the joint team utility and the average of the opponent utility. An ANOVA test ($\alpha = 0.05$) with Bonferroni post-hoc analysis was carried out to detect statistically different averages. The statistically better configurations for each of the three scenarios are highlighted in bold font style. It can be appreciated that *when unpredictable issues are more important for the opponent, the best results in average are obtained by taking an opponent oriented approach*: proposing *unpredictable* partial offers that are very likely to be acceptable for the opponent and satisfy remaining members' aspirations by demanding on *predictable and compatible issues*, which are less important for the opponent. This approach gets better results in terms of the opponent utility and the team joint utility¹ As for the scenario where *unpredictable issues are more important for the team, it is clearly observed that the best choice for team joint utility is to give a high weight to w_A , thus employing a team oriented approach*. Since *unpredictable* issues are more important for the team, they should satisfy their needs as much as possible with proposed *unpredictable* partial offers and demand less on predictable issues, which are more important for the opponent. The opponent utility is maximized when taking an opponent oriented approach, but it results in a considerable reduction in the team joint utility. Thus, in normal conditions, team members do not have any incentive to use an opponent oriented approach over a team oriented approach in this scenario. Finally, the last scenario corresponds to the case where *unpredictable issues have the same importance for both parties*. In this case, there may be more conflict between the team and its opponent since the parties do not have a clear trade-off opportunity such as increasing the demand

¹The p-value when comparing the team joint utility obtained by the opponent oriented approach and the normal approach was 0.07, which is very close to 0.05. Therefore, we decided to consider it as statistically different.

6. NEGOTIATION TEAMS IN UNPREDICTABLE DOMAINS

Equal importance on unpredictable issues		
	T. Joint	Op
Normal	0.168	0.480
Opponent Oriented	0.155	0.521
Team Oriented	0.116	0.323
Unpredictable issues more important for the team		
	T. Joint	Op
Normal	0.213	0.576
Opponent Oriented	0.200	0.595
Team Oriented	0.248	0.561
Unpredictable issues more important for the opponent		
	T. Joint	Op
Normal	0.280	0.627
Opponent Oriented	0.296	0.664
Team Oriented	0.271	0.559

Table 6.2: Impact of w_A and w_{op} on the team joint utility and the opponent utility in different scenarios

on *unpredictable* issues while decreasing the demand on predictable issues as appeared in two previous cases. One can observe that *the best team joint utility is obtained when using the standard team members*¹ ($w_A = w_{op} = 0.5$). Since both parties give the same importance to *unpredictable* issues, it seems natural to give the same importance to the acceptability of the *unpredictable* partial offer by the team and the opponent. The team oriented approach is clearly worse than the rest of approaches since most of the negotiations (42 %) ended in failure due to the team being too demanding and not satisfying the opponent's preferences. As for the opponent utility, the best option seems again the opponent oriented approach, but the cost is obtaining worse team joint utility than the best approach. Thus, normally, team members do not have incentive to use an opponent oriented approach over a normal approach in this scenario.

¹When comparing the team joint utility obtained by the normal approach and the opponent oriented approach, the p-value was 0.06, which is very close to 0.05. Thus, we considered both results as statistically different in practice.

However, for an effective adjustment of w_A and w_{op} , some information regarding which party gives more importance to *unpredictable* issues may be needed. It can be observed that in the three scenarios analyzed in this chapter, standard team members are the best option when both parties give the same importance to *unpredictable* issues and the second best option in the rest of the cases. The opponent oriented approach is the best option in one scenario (when *unpredictable* issues are more important for the opponent), the second best option in one scenario (when *unpredictable* issues are equally important for both parties), and the worse option in another scenario (when *unpredictable* issues are more important for the team). As for the team oriented approach, it seems to be the best option in one of the scenarios (when *unpredictable* issues are more important for the team), and the worse option in the other two scenarios. Hence, in absence of any prior information regarding this matter, a conservative approach suggests using standard team members and assuming that both parties give the same importance to *unpredictable* issues.

6.7.3 Third Experiment: Analyzing the Impact of the Reservation Utility

In the third experiment, we investigate the impact of team members' reservation utility on the team performance. As explained in Section 6.4.1, team members jointly prune a part of the negotiation space (i.e., a set of *unpredictable* partial offers) which does not contain, with absolute certainty, any unanimously acceptable offer. This pruning is directly related with the reservation utility of team members, which represents the minimum acceptable utility by team members. Therefore, any offer with a utility lower than the reservation utility is not acceptable for the agent.

6. NEGOTIATION TEAMS IN UNPREDICTABLE DOMAINS

Lower reservation utilities make it easier to obtain the needed utility by just setting *compatible and predictable* issues. Thus, each team member prunes less negotiation space with the *unpredictable* partial offers sent to the team mediator. Presumably, a joint list of forbidden *unpredictable* partial offers (i.e., the negotiation space that is pruned) with lower reservation utilities is smaller than lists constructed with higher reservation utilities. This leaves more room for finding an agreement with the opponent. However, if team members have very low reservation utilities, despite having more room for finding an agreement, they may end up with low utility agreements in the end. On the contrary, with higher reservation utilities, it is harder to obtain the needed utility with *compatible and predictable* issues. Therefore, each team member may need to prune more negotiation space and the joint list of forbidden *unpredictable* partial offers will be larger than the list constructed with lower reservation utilities. In fact, if team members set very high aspirations with their reservation utility, it may end up with all the negotiation space being pruned. However, if an agreement is found under these conditions, it may lead to team members achieving high levels of utility.

In this experiment, we test the impact of having different levels of reservation utility on team performance. More specifically, as we did in Section 6.7.1, we tested teams employing the Bayesian model against different families of strategies: competitor (i.e., Agent K), matcher (i.e., Nice Tit-for-Tat), inverter (i.e., very bouldware), and conceiver (i.e., conceiver). As an additional dimension to our study, we also introduced preference similarity as in our first experiment. Therefore, teams are tested in the scenario where team members' preferences are very dissimilar, the scenario where team members' preferences have an average degree of similarity, and the scenario where team members' preferences are very similar. We configured three different types of Bayesian teams (i.e., teams composed by Bayesian team members) with different levels of reservation utilities. First,

	Very Similar	Average Similarity	Very Dissimilar
$RU_{a_i} = 0.35$	0.4%	11.6%	35.3%
$RU_{a_i} = 0.50$	23.8%	34.2%	72.6%
$RU_{a_i} = 0.65$	73.7%	81.8%	90.8%

Table 6.3: Average percentage of unpredictable partial offers pruned in the pre-negotiation.

a Bayesian team with a relatively low reservation utility $RU_{a_i} = 0.35$. Second, a Bayesian team with a moderate reservation utility $RU_{a_i} = 0.5$, and, finally, a Bayesian team with a high reservation utility $RU_{a_i} = 0.65$. These three types of teams were faced in every scenario and negotiation case against every type of opponent for 20 repetitions. We gathered information on the team joint utility and the utility obtained by the opponent, and an ANOVA ($\alpha = 0.05$) with Bonferroni post-hoc analysis was carried out to determine results that are statistically better than the rest.

Table 6.3 shows the average percentage of *unpredictable* partial offers that were pruned in the pre-negotiation depending on the team configuration and team preference similarity. According to this table, it can be observed that as team dissimilarity increases, the number of *unpredictable* partial offers that are pruned in the pre-negotiation also increases. Since team members' preferences are gradually more dissimilar, the list of *unpredictable* partial offers shared with the team mediator by each team member tends to be more different from the rest of lists shared by other teammates. Thus, when joining all of the lists it is just natural that more *unpredictable* partial offers are pruned. The experiment shows that as reservation utilities for team members increase, the tendency is to prune more negotiation space. If team members play excessively high reservation utilities, this may effectively result in leaving no room at all for negotiation by pruning all the negotiation space.

6. NEGOTIATION TEAMS IN UNPREDICTABLE DOMAINS

Very Similar								
	Agent K		Nice Tit-for-Tat		Very Boulware		Conceder	
	T. Joint	Op.	T. Joint	Op.	T. Joint	Op.	T. Joint	Op.
$RU_{a_i} = 0.35$	0.195	0.760	0.117	0.799	0.160	0.761	0.526	0.493
$RU_{a_i} = 0.50$	0.263	0.690	0.164	0.746	0.224	0.574	0.557	0.472
$RU_{a_i} = 0.65$	0.350	0.619	0.286	0.667	0.354	0.634	0.635	0.431
Average Similarity								
	Agent K		Nice Tit-for-Tat		Very Boulware		Conceder	
	T. Joint	Op.	T. Joint	Op.	T. Joint	Op.	T. Joint	Op.
$RU_{a_i} = 0.35$	0.167	0.651	0.090	0.758	0.136	0.667	0.342	0.440
$RU_{a_i} = 0.50$	0.248	0.583	0.158	0.669	0.224	0.574	0.390	0.414
$RU_{a_i} = 0.65$	0.242	0.402	0.268	0.535	0.313	0.493	0.470	0.378
Very Dissimilar								
	Agent K		Nice Tit-for-Tat		Very Boulware		Conceder	
	T. Joint	Op.	T. Joint	Op.	T. Joint	Op.	T. Joint	Op.
$RU_{a_i} = 0.35$	0.193	0.577	0.115	0.661	0.173	0.568	0.373	0.408
$RU_{a_i} = 0.50$	0.209	0.457	0.196	0.559	0.271	0.489	0.475	0.367
$RU_{a_i} = 0.65$	0.068	0.091	0.346	0.459	0.409	0.408	0.580	0.331

Table 6.4: Average joint Utility (T. Joint) and average opponent Utility (Op) for teams composed by Bayesian team members with different reservation utilities.

Table 6.4 shows the results of this experiment in terms of the average joint utility and the average opponent utility. A bold font style is used to highlight those Bayesian team configurations that are statistically the best option against each opponent. It can be observed that despite the degree of dissimilarity among team members' preferences, in this experimental setting *team members obtained statistically better team joint utility by setting high reservation utilities*. This pattern arose against Nice Tit-for-Tat, Very Boulware and Conceder opponents. Nevertheless, when facing Agent K, this pattern could only be observed when team members' preferences are very similar or they have an average similarity. As preference dissimilarity increased, we can observe how setting a high reservation utility (i.e., $RU_{a_i} = 0.65$) gradually becomes the worst possible course of action (from the options studied) when facing Agent K. The reason for this behavior is mainly

explained due to the increase in the number of failed negotiations. If we observe the very similar scenario, the number of failed negotiations when facing Agent K and $RU_{a_i} = 0.65$ is 1.6%. If we change to average similarity scenarios, the number of failed negotiations is 34.2% when facing Agent K with $RU_{a_i} = 0.65$. The same measure is increased to 81.8% in the very dissimilar scenario. Recalling what we observed in Table 6.3, as team dissimilarity increases, the number of *unpredictable* partial offers to be pruned is larger. This leaves less negotiation space to be played with Agent K. Differently from Very Boulware, Conceder, and Nice Tit-for-Tat, Agent K is a competitor agent that attempts to concede as less as possible by estimating the maximum utility that can be obtained from the opponent and employing a limit of compromise when the opponent takes a hard stance. First of all, if reservation utilities are high, it can be considered that team members play a hard stance. Second, if too much negotiation space is pruned, it may be feasible that the set of remaining *unpredictable* partial offers precludes Agent K from reaching its limit of compromise. Thus, employing such a configuration against a competitor agent like Agent K may result, as we have observed in practice, in an increase in the number of failed negotiations. In the case of the opponent utility, it is always maximized when the team sets low reservation values. This is natural since, in the end, team members will concede more and any type of opponent can take advantage from this situation.

In conclusion, we have observed that generally team members may benefit from playing high reservation utilities against conceders, inverters, and matchers. However, if faced against competitors like Agent K, setting high reservation utilities may prune too much negotiation space, especially when team members are very dissimilar. This results in negotiation spaces that may not contain the minimum limits established by competitor agents, thus, ending negotiations with failure. In

6. NEGOTIATION TEAMS IN UNPREDICTABLE DOMAINS

general, team members should be cautious when setting the reservation utility since it may end up in more negotiation failures.

6.8 Conclusions

In this chapter we have presented an extension to FUM which tackles domains with *unpredictable* issues. This allows us to tackle a wide variety of negotiation scenarios in electronic commerce. The extension is capable of assuring unanimously acceptable agreements for all of the team members. It takes advantage of the categorization of negotiation issues as *predictable and compatible*, and *unpredictable*. We have proposed two different types of team members for the current model: a basic team member that proposes *unpredictable* partial offers during the negotiation solely guided by its own utility function, and a Bayesian team member that suggests *unpredictable* partial offers based on the preferences of the team and the preferences of the opponent.

Results have shown that, as long as preferential conflict is present in the team, team members have a incentive to employ the Bayesian strategy over the basic strategy. Additionally, we have shown that in absence of information regarding which party gives more importance to *unpredictable* issues, Bayesian team members should give the same importance to the team's preferences and the opponent's preferences over giving more importance to the team's preferences. Finally, we have also shown that team members may benefit from playing higher reservation utilities against several types of agents like conceders, matchers and inverters. Nevertheless, setting high reservation utilities may become the worst option as team members' preferences are more dissimilar and the opponent plays a competitor strategy.

7

Conclusions

As mentioned in the introduction of this thesis, the main goal of this work is providing negotiation models for complex negotiations in multi-agent systems. More specifically, we set as goals of this thesis the design of a negotiation model for Ambient Intelligence domains and the design of negotiation models for the novel topic of agent-based negotiation teams. The latter refers to groups of persons/agents that join together as a single negotiation party because they share a common interest at the negotiation at hand. In the following section, we outline how this thesis has contributed to the resolution of the aforementioned goals.

7.1 Contributions

The main contributions of this thesis are the following:

- We reviewed the state-of-the-art in automated negotiation from the point of view of Artificial Intelligence. A special emphasis was put to analyze which current negotiation models are more convenient for Ambient Intelligence due to the specificities of the domain: scarce computational resources,

7. CONCLUSIONS

the dynamicity of the system, and limited computation time. The analysis allowed us to identify that from the point of view of negotiation in Ambient Intelligence, the use of time-based concession tactics, the use of complex and interdependent utility functions for representing agents' preferences, and the use of similarity heuristics and genetic algorithms as learning mechanisms may be more adequate for Ambient Intelligence.

- We proposed a negotiation model for Ambient Intelligence domains that is applicable to bilateral negotiations which employ a k-alternating offer protocol. Each agent employs genetic algorithms during the pre-negotiation to sample good and significantly different offers for oneself. During the negotiation, agents propose offers from the iso-utility curve that are the most similar ones to the last offers received from the opponent. Additionally, each agent uses genetic operators over one's good own offers and offers received from the opponent in order to sample new offers that are interesting for both parties. The results show that the proposed negotiation model is capable of obtaining statistically equivalent results to negotiation models that sample the whole negotiation domain while sampling a significantly lower number of negotiation offers. This is significantly important for Ambient Intelligence domains, since it certainly reduces the amount of computational resources employed in the negotiation.
- We also contributed with a general workflow of negotiation tasks for agent-based negotiation teams. The workflow aims to identify the tasks that may help agent-based negotiation teams to reach success in negotiations. Each of the tasks is analyzed and related with current research that is being carried out in multi-agent systems and Artificial Intelligence. On top of that, the analysis also goes further and it points out those unsolved problems that may

appear in each of the phases due to the specific characteristics of agent-based negotiation teams.

- From the proposed workflow tasks, we have focused on proposing intra-team strategies, which govern team dynamics during the negotiation. More specifically, we explored the space of intra-team strategies for negotiation teams whose team members have different preferences. We analyzed intra-team strategies for team members in domains exclusively composed by *compatible and predicatable* negotiation issues: the representative intra-team strategy, the similarity simple voting intra-team strategy, the similarity borda voting intra-team strategy, and the full unanimity mediated intra-team strategy. Each of the aforementioned intra-team strategies is capable of guaranteeing different levels of unanimity regarding team decisions: no unanimity, majority/plurality, semi-unanimity, and strict unanimity. The different intra-team strategies are analyzed under different environmental conditions (e.g., deadline lengths, intra-team conflict, team size, etc.) and those intra-team strategies that perform better are identified. We have been able to identify that depending on the environmental conditions and the team's goal (e.g., average utility of team members, minimum utility of team members, number of rounds, etc.), some intra-team strategies work better than other intra-team strategies. Thus, environmental conditions play a key role in the selection of the strategy to be carried out during the negotiation.
- Additionally, we extended the full unanimity mediated intra-team strategy to tackle domains with *compatible and predictable* issues, and *unpredictable* issues. The extension is capable of guaranteeing that if an agreement is found, it is unanimously acceptable among team members. The implementation of this framework has been carried out in GENIUS, which allows us to (i) use a

7. CONCLUSIONS

wide variety of opponent agents; (ii) facilitate research on agent-based negotiation teams by other scholars, and even give room to a future negotiating competition involving teams. The performance of the model has been tested in some environmental conditions, and we have been able to observe how such conditions affect the performance of the different model configurations. Additionally, we have been able to show that the proposed extension performs statistically better than other intra-team strategies in domains with both types of issues.

7.2 Future Work

Due to the novelty of the topic, agent-based negotiation teams is the area of research that, as of today, possibly remains more unexplored and with vast opportunities for new research. With respect to negotiation in the Ambient Intelligence domain, we consider that there are still interesting areas that should be researched in the future. Next, we describe some of the future lines of work that we consider as potentially interesting for future research:

- As introduced in the general workflow of tasks presented in Chapter 4, several tasks should be carried out by negotiation teams during the pre-negotiation. These pre-negotiation tasks have been largely unexplored by multi-agent literature and compose a critical problem since factors like team formation, understanding the negotiation domain, and selecting a proper negotiation protocol may have an important impact on team performance during the negotiation. Thus, we consider that future work in negotiation teams should also explore team tasks carried out in the pre-negotiation.
- In Chapter 1 we commented that one of the reasons to employ a negotiation team is gathering together individuals with different expertise and

complementary skills. When the negotiation domain is complex, it requires of different knowledge areas and skills in order to be tackled successfully. Thus, if agent-based negotiation teams are to participate in complex negotiations, different expertise, complementary skills, and proper coordination mechanisms are needed. This is a scenario that we have not studied in this thesis, where we have focused on negotiation teams which are composed by team members with different preferences.

- In this thesis, it was proved that environmental conditions may affect the performance of different intra-team strategies. In fact, the optimal intra-team strategy from a given intra-team strategy may vary according to the environmental conditions. Thus, negotiation teams should select the current intra-team strategy according to such information. We argue that a counselor agent may help team members to select the optimal intra-team strategy. The counselor agent would observe for changes in the environment and also attempt to learn the strategies employed by opponents given those environmental conditions. Based on this information, team necessities, and the results of simulations like the ones carried out in this thesis, the counselor agent may be able to advise teams on which intra-team strategies would work better.
- As stated in Chapter 6, the team mediator may employ the information on forbidden unpredictable partial offers to detect whether or not unanimously acceptable agreements are possible among team members and to form negotiation teams with lower degrees of intra-team conflict. We are interested on designing mechanisms that allow us to employ such information for successfully forming negotiation teams.

7. CONCLUSIONS

- At some points of this thesis, we have discussed about the possibility for some opponent agents and some competitor agents to infiltrate the negotiation team. These agents aim to sabotage the negotiation teams with different purposes and methods. The consequences of such manipulations may be very negative for teams (e.g., very expensive deals). Therefore, one of our concerns and future lines of work is the design of mechanisms (e.g., trust and reputation) or intra-team strategies (i.e., full unanimity mediated in the case of manipulations carried out by opponent agents) that are robust against these types of manipulations.
- As we have observed in this thesis, intra-team conflict with regards to team members' preferences is one of the factors that affects team performance the most. As we also discussed, the formation of negotiation teams has special considerations that should be addressed like social power over sellers, the social network of team members, the skill distribution among team members, and so forth. Another important issue to be considered is the dynamicity of the multi-agent system. For instance, in Ambient Intelligence domains, users may enter and leave the system very quickly, resulting in a more challenging scenario for negotiation team formation. Therefore, we think that, due to the aforementioned reasons, negotiation team formation is a topic worthy of being researched. As a case of study we could think of a ubiquitous mall where users with similar needs are grouped together in order to take advantage of group discounts. In this scenario, we would be able to contribute in the state-of-the-art of agent-based negotiation teams and negotiation in Ambient Intelligence.

Along the aforementioned research issues that may be worthy of being researched in the future, we have identified some potential scenarios that may be

supported by some of the technologies developed in this thesis:

- Even though there are some works in the area of human negotiation teams, the topic remains uncovered with respect to dyadic negotiations. One of the reasons for this issue is the fact that team dynamics are hard to study by social scientists, and experiments with human negotiation teams require more economic efforts. In that sense, computational models that mimic human negotiation teams may help social scientists. Firstly, simulations with computational models are cheap and it may allow social scientists to initially explore several scenarios before expending vast economic resources. Secondly, computational models can be employed as substitutes for human participants, which may also save economic resources for social scientists. And lastly, but not the least important, computational models can be used to train real negotiators in several scenarios. With respect to computational models for human negotiation teams, we are currently working in computational models that take into account cultural factors in negotiation team dynamics (164).
- The Smart Grid is the next generation electricity distribution grid. The grid abases its functioning on information networks that allow customers and sellers to purchase/sell energy intelligently. Those decisions may be based on consumption peaks, user requirements, energy saving policies, and so forth. We think that agent-based negotiation teams may be employed in these networks to gather together groups of consumers or groups of sellers. In the first case, groups of buyers with similar characteristics (e.g., similar energetic needs, physical proximity, etc.) may form groups which allow them to take advantage from buying at bulk. In the second case, groups of small green energy producers may form a virtual power plant, which may allow

7. CONCLUSIONS

to provide a notable amount of energy. This may result in green energy producers being able to compete more fairly with big producers. Therefore, we consider that applying agent-based negotiation teams in the smart grid is a potentially interesting scenario.

7.3 Scientific Publications

Next, all of the publications describing the results of this thesis work are listed.

7.3.1 Publications in SCI Journals

- V. Sanchez-Anguix, V. Julian, V. Botti and A. Garcia-Fornes. *Reaching Unanimous Agreements within Agent-based Negotiation Teams with Linear and Monotonic Utility Functions*. IEEE Transactions on Systems, Man and Cybernetics - Part B. Volume 42(3), pages 778-792, 2012. **Impact Factor 2.699 (Q1 Computer Science, Artificial Intelligence)**.
- V. Sanchez-Anguix, S. Valero, V. Julian, V. Botti and A. Garcia-Fornes. *Evolutionary-aided negotiation model for bilateral bargaining in Ambient Intelligence domains with complex utility functions*. Information Sciences, In Press, 2011. **Impact Factor 3.291 (Q1 Computer Science, Information Systems)**.
- V. Sanchez-Anguix, V. Julian, V. Botti and A. Garcia-Fornes. *Studying the Impact of Negotiation Environments on Negotiation Teams' Performance*. Information Sciences, In Press, 2012. **Impact Factor 2.833 (Q1 Computer Science, Information Systems)**.
- V. Sanchez-Anguix, V. Julian, V. Botti and A. Garcia-Fornes. *A Workflow of Tasks for Agent-Based Negotiation Teams: Analysis and Challenges*.

Submitted to IEEE Transactions on Systems, Man and Cybernetics - Part C.

- V. Sanchez-Anguix, R. Aydogan, V. Julian, A. Garcia-Fornes and C.M. Jonker. *Unanimously Acceptable Agreements for Negotiation Teams in Unpredictable Domains*. Submitted to Knowledge and Information Systems.

7.3.2 Publications in Conferences

- V. Sanchez-Anguix, R. Aydogan, V. Julian and C. Jonker. *Analysis of Intra-Team Strategies for Teams Negotiating Against Competitor, Matchers, and Conceders*. In the 5th International Workshop on Agent-based Complex Automated Negotiations (ACAN2012@AAMAS2012), pages 1-8, 2012.
- V. Sanchez-Anguix, T. Dai, Z. Semnani-Azad, K. Sycara and V. Botti. *Modeling power distance and individualism/collectivism in negotiation team dynamics*. In the 45 Hawaii International Conference on System Sciences (HICSS-45), pages 628-637, 2012. **ERA Conference Ranking A.**
- V. Sanchez-Anguix, V. Julian and A. Garcia-Fornes. *Agent-based Negotiation Teams*. In the 22nd International Joint Conference on Artificial Intelligence (IJCAI 2011), pages 2844-2845, 2011. **ERA Conference Ranking A.**
- V. Sanchez-Anguix. *Negotiation Teams in Multiagent Systems (Extended Abstract)*. In the 10th International International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2011), pages 1357-1358, 2011. **ERA Conference Ranking A.**

7. CONCLUSIONS

- V. Sanchez-Anguix, V. Julian, V. Botti and A. Garcia-Fornes. *Analyzing Intra-Team Strategies for Agent-Based Negotiation Teams*. In the 10th International International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2011), pages 929-936, 2011. **ERA Conference Ranking A.**
- V. Sanchez-Anguix, V. Julian and A. Garcia-Fornes. *From an Individual Perspective to a Team Perspective in Agent-Based Negotiation*. In the 9th International Conference on Practical Applications of Agents and Multi-Agent Systems (PAAMS 2011), pages 217-223, 2011. **ERA Conference Ranking C.**
- V. Sanchez-Anguix, V. Julian, V. Botti and A. Garcia-Fornes. *Towards agent-based negotiation teams*. In Group Decision and Negotiation 2010 (GDN 2010), pages 328-331, 2010. **ERA Conference Ranking B.**
- V. Sanchez-Anguix, S. Valero, V. Julian, V. Botti and A. Garcia-Fornes. *Genetic-Aided Multi-Issue Bilateral Bargaining for Complex Utility Functions*. In the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2010), pages 1601-1602, 2010. **ERA Conference Ranking A.**

7.4 Scientific Research Stays

- 02-12-2011 to 11-12-2011. *Delft University of Technology, Delft, The Netherlands*, COST Action IC0801 Short Term Scientific Mission supervised by Dr. Reyhan Aydogan and Prof. Catholijn Jonker on Agent-Based Negotiation Teams in Genius.

7.4 Scientific Research Stays

- 01-09-2010 to 30-11-2011. *Carnegie Mellon University, Pittsburgh, Pennsylvania, USA*, research stay supervised by Prof. Katia Sycara on Cultural Factors in Negotiation Teams.

7. CONCLUSIONS

References

- [1] M. WOOLDRIDGE. *Multi-agent systems : an introduction*. Wiley, 2001. 1, 17
- [2] R. H. GUTTMAN, A. G. MOUKAS, AND P. MAES. **Agent-mediated electronic commerce: a survey**. *Knowledge Engineering Review*, **13**(2):147–159, 1998. 1, 16
- [3] J. OTA. **Multi-agent robot systems as distributed autonomous systems**. *Advanced engineering informatics*, **20**(1):59–70, 2006. 1
- [4] J. PITA, M. JAIN, J. MARECKI, F. ORDÓÑEZ, C. PORTWAY, M. TAMBE, C. WESTERN, P. PARUCHURI, AND S. KRAUS. **Deployed ARMOR protection: the application of a game theoretic model for security at the Los Angeles International Airport**. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems: industrial track*, pages 125–132. International Foundation for Autonomous Agents and Multiagent Systems, 2008. 1
- [5] E. ARGENTE, V. BOTTI, C. CARRASCOSA, A. GIRET, V. JULIAN, AND M. REBOLLO. **An Abstract Architecture for Virtual Organizations: The THOMAS approach**. *Knowledge and Information Systems*, **29**(2):379–403, 2011. 1, 13
- [6] E. ARGENTE, V. BOTTI, AND V. JULIAN. **GORMAS: An Organizational-Oriented Methodological Guideline for Open MAS**. In *Agent-Oriented Software Engineering X- 10th Int. Workshop AOSE 2009 - Revised Selected Papers*, **6038**, pages 32–47. Springer, 2011. 1
- [7] M. LUCK AND P. MCBURNEY. **Computing as Interaction: Agent and Agreement Technologies**. In *Proc. IEEE Conference on Distributed Human-Machine Systems*, pages 1–6, 2008. 1, 12
- [8] C. SIERRA, V. BOTTI, AND S. OSSOWSKI. **Agreement Computing**. *KI-Kunstliche Intelligenz*, pages 1–5, 2011. 1, 12
- [9] D. G. PRUITT. *Negotiation Behavior*. Academic Press, 1981. 1, 15
- [10] L. THOMPSON. *The Mind and heart of the negotiator*. Prentice Hall Press, Upper Saddle River, NJ, USA, 2005. 2, 37
- [11] J.F. NASH. **The bargaining problem**. *Econometrica*, **18**:155–162, 1950. 2, 15
- [12] A. RUBINSTEIN. **Perfect equilibrium in a bargaining model**. *Econometrica*, **50**:155–162, 1982. 2, 15, 105, 106, 107, 120, 122, 185
- [13] K. BINMORE, A. RUBINSTEIN, AND A. WOLINSKY. **The Nash bargaining solution in economic modelling**. *Rand Journal of Economics*, **17**:176–188, 1986. 2, 15
- [14] K.P. SYCARA. **Multiagent compromise via negotiation**. In *Distributed Artificial Intelligence*, **2**, pages 119–137. Morgan Kaufmann Publishers Inc., 1989. 2, 20, 41, 42
- [15] K. SYCARA. *Negotiation Planning: An AI Approach*, **46**(1):216–234, May 1990. 2, 20, 41, 42
- [16] K. SYCARA. **Problem restructuring in negotiation**. *Management Science*, **37**(10):1248–1268, 1991. 2, 20, 41, 42
- [17] M. TAWFIK JELASSI AND A. FOROUGHI. **Negotiation support systems: An overview of design issues and existing software**. *Decision Support Systems*, **5**(2):167–181, 1989. 2
- [18] P. FARATIN, C. SIERRA, AND N. R. JENNINGS. **Negotiation Decision Functions for Autonomous Agents**. *International Journal of Robotics and Autonomous Systems*, **24**(3-4):159–182, 1998. 3, 18, 26, 41, 59, 121, 130, 162, 209, 219
- [19] H. H. BUI, S. VENKATESH, AND D. KIERONSKA. **Learning Other Agents’ Preferences in Multi-Agent Negotiation Using the Bayesian Classifier**. *International Journal of Cooperative Information Systems*, **8**(4):275–293, 1999. 3, 22, 41, 42
- [20] P. FARATIN. *Automated service negotiation between autonomous computational agents*. PhD thesis, University of London, 2000. 3, 104
- [21] C.M. JONKER, V. ROBU, AND J. TREUR. **An agent architecture for multi-attribute negotiation using incomplete preference information**. *Autonomous Agents and Multi-Agent Systems*, **15**(2):221–252, 2007. 3, 20, 41
- [22] J. M. CORCHADO, J. BAJO, Y. DE PAZ, AND D. I. TAPIA. **Intelligent environment for monitoring Alzheimer patients, agent technology for health care**. *Decision Support Systems*, **44**(2):382–396, 2008. 3, 17
- [23] W. WEBER, J. M. RABAAY, AND E. AARTS. *Ambient Intelligence*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2005. 3, 17
- [24] L. THOMPSON, E. PETERSON, AND S.E. BRODT. **Team negotiation: An examination of integrative and distributive bargaining**. *Journal of Personality and Social Psychology*, **70**(1):66–78, 1996. 4, 35, 37

REFERENCES

- [25] E. PETERSON AND L. THOMPSON. **Negotiation Teamwork: The Impact of Information Distribution and Accountability on Performance Depends on the Relationship among Team Members.** *Organizational Behavior and Human Decision Processes*, **72**(3):364–383, 1997. 4, 35, 36, 37
- [26] SUSAN BRODT AND LEIGH THOMPSON. **Negotiating Teams: A levels of analysis.** *Group Dynamics*, **5**(3):208–219, 2001. 4, 35
- [27] C. HEWITT. **Open information systems semantics for distributed artificial intelligence.** *Artificial intelligence*, **47**(1-3):79–106, 1991. 11
- [28] I. FOSTER, C. KESSELMAN, AND S. TUECKE. **The anatomy of the grid: Enabling scalable virtual organizations.** *International Journal of High Performance Computing Applications*, **15**(3):200–222, 2001. 12
- [29] H. KITANO AND S. TADOKORO. **RoboCup Rescue: A Grand Challenge for Multiagent and Intelligent Systems.** *AI Magazine*, **22**(1):39–52, 2001. 12
- [30] M.P. PAPAIOGLOU AND D. GEORGAKOPOULOS. **Service-oriented computing.** *Communications of the ACM*, **46**(10):25–28, 2003. 13
- [31] Y. KALFOGLOU AND M. SCHORLEMMER. **IF-Map: An ontology-mapping method based on information-flow theory.** *Journal on data semantics I*, pages 98–127, 2003. 13
- [32] N.F. NOY. **Semantic integration: a survey of ontology-based approaches.** *ACM Sigmod Record*, **33**(4):65–70, 2004. 13
- [33] N. CHOI, I.Y. SONG, AND H. HAN. **A survey on ontology mapping.** *ACM Sigmod Record*, **35**(3):34–41, 2006. 13
- [34] G. BOELLA, L. VAN DER TORRE, AND H. VERHAGEN. **Introduction to normative multiagent systems.** *Computational & Mathematical Organization Theory*, **12**(2):71–79, 2006. 13
- [35] N. CRIADO, E. ARGENTE, AND V. BOTTI. **THOMAS: An Agent Platform For Supporting Normative Multi-Agent Systems.** 13
- [36] B. HORLING AND V. LESSER. **Using quantitative models to search for appropriate organizational designs.** *Autonomous Agents and Multi-Agent Systems*, **16**:95–149, 2008. 13, 108
- [37] S. ESPARCIA AND E. ARGENTE. **Defining Virtual Organizations Following a Formal Approach.** 13
- [38] J. SABATER AND C. SIERRA. **Review on computational trust and reputation models.** *Artificial Intelligence Review*, **24**(1):33–60, 2005. 14, 98, 132
- [39] J. M. SUCH, A. ESPINOSA, A. GARCIA-FORNES, AND V. BOTTI. **Partial identities as a foundation for Trust and Reputation.** *Engineering Applications of Artificial Intelligence*, **24**(7):1128–1136, 2011. 14, 98, 99, 132
- [40] S. S. FATIMA, M. J. WOOLDRIDGE, AND N. R. JENNINGS. **An agenda-based framework for multi-issue negotiation.** *Artificial Intelligence*, **152**(1):1–45, 2004. 15
- [41] Y. IN AND R. SERRANO. **Agenda restrictions in multi-issue bargaining (II): unrestricted agendas.** *Economics Letters*, **79**(3):325–331, 2003. 15
- [42] C. BEAM AND A. SEGEV. **Automated negotiations: A survey of the state of the art.** *Wirtschaftsinformatik*, **39**(3):263–268, 1997. 16
- [43] S. KRAUS. **Negotiation and cooperation in multi-agent environments.** *Artificial Intelligence*, **94**(1-2):79–97, 1997. 16
- [44] N. R. JENNINGS, P. FARATIN, A. R. LOMUSCIO, S. PARSONS, M. J. WOOLDRIDGE, AND C. SIERRA. **Automated Negotiation: Prospects, Methods and Challenges.** *Group Decision and Negotiation*, **10**:199–215, Mar 2001. 16
- [45] ALESSIO R. LOMUSCIO, MICHAEL WOOLDRIDGE, AND NICHOLAS R. JENNINGS. **A Classification Scheme for Negotiation in Electronic Commerce.** *Group Decision and Negotiation*, **12**:31–56, 2003. 16
- [46] F. LOPES, M. WOOLDRIDGE, AND A.Q. NOVAIS. **Negotiation among autonomous computational agents: principles, analysis and challenges.** *Artificial Intelligence Review*, **29**(1):1–44, 2008. 16, 91, 92, 93
- [47] J. A. FRAILE, J. BAJO, A. ABRAHAM, AND J. M. CORCHADO. **HoCaMA: Home Care Hybrid Multiagent Architecture.** In *Pervasive Computing*, Computer Communications and Networks, pages 259–285. Springer London, 2009. 17
- [48] P. REMAGNINO AND G.L. FORESTI. **Ambient Intelligence: A New Multidisciplinary Paradigm.** *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans*, **35**(1):1–6, Jan. 2005. 17
- [49] S. KEEGAN, G.M.P. O’HARE, AND M.J. O’GRADY. **Easishop: Ambient intelligence assists everyday shopping.** *Information Sciences*, **178**(3):588–611, 2008. 17
- [50] J. BAJO, J. M. CORCHADO, Y. DE PAZ, J. FRANCISCO DE PAZ, S. RODRÍGUEZ, Q. MARTIN, AND A. ABRAHAM. **SHOMAS: Intelligent guidance and suggestions in shopping centres.** *Applied Soft Computing*, **9**(2):851–862, 2009. 17
- [51] M. J. OSBORNE AND A. RUBINSTEIN. *A Course in Game Theory*. MIT Press, 1999. 18, 75

REFERENCES

- [52] K. V. HINDRIKS AND D. TYKHONOV. **Opponent modelling in automated multi-issue negotiation using Bayesian learning.** In *The 7th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'08)*, pages 331–338, Richland, SC, 2008. International Foundation for Autonomous Agents and Multiagent Systems. 18, 24, 41, 42
- [53] C. M. JONKER AND J. TREUR. **An Agent Architecture for Multi-Attribute Negotiation.** In *The 17th International Joint Conference on Artificial Intelligence (IJCAI'01)*, pages 1195–1201. Morgan Kaufman, 2001. 20, 41
- [54] D.B. LEAKE. **Case-based reasoning.** *The Knowledge Engineering Review*, **9**(01):61–64, 1994. 20, 94
- [55] R. LOPEZ DE MANTARAS. **Case-Based Reasoning.** In *Machine Learning and Its Applications, 2049 of Lecture Notes in Computer Science*, pages 127–145. Springer Berlin / Heidelberg, 2001. 20, 94
- [56] P. FARATIN, C. SIERRA, AND N. R. JENNINGS. **Using Similarity Criteria to Make Issue Trade-Offs in Automated Negotiations.** *Artificial Intelligence*, **142**:205–237, 2002. 20, 21, 41, 42, 52, 55, 130
- [57] G. LAI, K. SYCARA, AND C. LI. **A decentralized model for automated multi-attribute negotiations with incomplete information and general utility functions.** *Multiagent and Grid Systems*, **4**(1):45–65, 2008. 20, 21, 41, 42, 47, 52, 53, 55, 61, 73, 76, 80, 84, 85, 105, 106, 130, 209, 219
- [58] S.J. RUSSELL AND P. NORVIG. *Artificial Intelligence: A Modern Approach*. Pearson Education, 2003. 22, 211
- [59] D. ZENG AND K. SYCARA. **Bayesian learning in negotiation.** *International Journal of Human-Computer Studies*, **48**(1):125–141, 1998. 22, 40, 41
- [60] V. NARAYANAN AND N. R. JENNINGS. **Learning to Negotiate Optimally in Non-stationary Environments.** In *Cooperative Information Agents X, 10th International Workshop, CIA 2006, Edinburgh, UK, September 11-13, 2006, Proceedings*, **4149 of Lecture Notes in Computer Science**, pages 288–300. Springer, 2006. 23, 40, 41, 112
- [61] S. BUFFETT AND B. SPENCER. **A bayesian classifier for learning opponents' preferences in multi-object automated negotiation.** *Electronic Commerce Research and Applications*, **6**(3):274–284, 2007. 23, 29, 41, 42
- [62] D. E. GOLDBERG. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1989. 25, 57
- [63] J. R. OLIVER. **On Artificial Agents for Negotiation in Electronic Commerce.** In *The 29 Hawaii International Conference on System Sciences (HICSS'96)*, **4**, pages 337–346, Washington, DC, USA, 1996. IEEE Computer Society. 25, 27, 41, 42, 109
- [64] N. MATOS, C. SIERRA, AND N. R. JENNINGS. **Determining Successful Negotiation Strategies: An Evolutionary Approach.** In *The 3rd International Conference on Multi Agent Systems (ICMAS '98)*, pages 182–189, Washington, DC, USA, 1998. IEEE Computer Society. 26, 27, 41, 42, 109
- [65] T. TU, E. WOLFF, AND W. LAMERSDORF. **Genetic Algorithms for Automated Negotiations: A FSM-based Application Approach.** In *Proceedings of the 11th International Conference on Database and Expert Systems (DEXA 2000)*, pages 1029–1033. IEEE, 9 2000. 27, 41, 42, 109
- [66] E. GERDING, D. VAN BRAGT, AND H. LA POUTRÉ. **Multi-Issue Negotiation Processes by Evolutionary Simulation, Validation and Social Extensions.** *Computational Economics*, **22**(1):39–63, 2003. 27, 41, 42, 109
- [67] R. KROVI, A. C. GRAESSER, AND W. E. PRACHT. **Agent behaviors in virtual negotiation environments.** *IEEE Transactions on Systems, Man, and Cybernetics, Part C*, **29**(1):15–25, 1999. 28, 41, 42
- [68] S. P. M. CHOI, J. LIU, AND S. CHAN. **A genetic agent-based negotiation system.** *Computer Networks*, **37**(2):195–204, 2001. 29, 41, 42
- [69] H.J. KUSHNER AND G. YIN. *Stochastic approximation and recursive algorithms and applications*, **35**. Springer Verlag, 2003. 29
- [70] R.Y.K. LAU. **Towards Genetically Optimised Multi-Agent Multi-Issue Negotiations.** In *The 38 Hawaii International Conference on System Sciences (HICSS'05)*, pages 35–45, 2005. 29, 41, 42
- [71] ROBERT M. COEHOORN AND NICHOLAS R. JENNINGS. **Learning on opponent's preferences to make effective multi-issue negotiation trade-offs.** In *The 6th International Conference on Electronic Commerce (ICEC'04)*, pages 59–68, 2004. 29, 40, 41
- [72] S.J. SHEATHER AND M.C. JONES. **A reliable data-based bandwidth selection method for kernel density estimation.** *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 683–690, 1991. 29
- [73] R. CARBONNEAU, G.E. KERSTEN, AND R. VAHDOV. **Predicting opponent's moves in electronic negotiations using neural networks.** *Expert Systems with Applications*, **34**(2):1266–1273, 2008. 30, 40, 41
- [74] RÉAL A. CARBONNEAU, GREGORY E. KERSTEN, AND RUSTAM M. VAHDOV. **Pairwise issue modeling for negotiation counteroffer prediction using neural networks.** *Decision Support Systems*, **50**(2):449 – 459, 2011. 30, 40, 41
- [75] J. WALLENIUS, J. S. DYER, P. C. FISHBURN, R. E. STEUER, S. ZIONTS, AND K. DEB. **Multiple Criteria Decision Making, Multiattribute Utility Theory: Recent Accomplishments and What Lies Ahead.** *Management Science*, **54**(7):1336–1349, 2008. 31

REFERENCES

- [76] R. L. KEENEY AND H. RAIFFA. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. Cambridge University Press, 1993. 31
- [77] M. KLEIN, P. FARATIN, H. SAYAMA, AND Y. BAR-YAM. **Negotiating Complex Contracts**. *Group Decision and Negotiation*, **12**:111–125, March 2003. 31, 32, 40, 41
- [78] T. ITO, M. KLEIN, AND H. HATTORI. **A multi-issue negotiation protocol among agents with nonlinear utility functions**. *Multiagent and Grid Systems*, **4**(1):67–83, 2008. 32, 34, 35, 40, 41, 66, 68, 73
- [79] V. ROBU, D. J. A. SOMEFUN, AND J. A. LA POUTRÉ. **Modeling complex multi-issue negotiations using utility graphs**. In *Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems (AAMAS'05)*, pages 280–287, New York, NY, USA, 2005. ACM. 33, 40, 41
- [80] V. ROBU AND J. A. LA POUTRÉ. **Retrieving the Structure of Utility Graphs Used in Multi-Item Negotiation through Collaborative Filtering of Aggregate Buyer Preferences**. In *Rational, Robust and Secure Negotiations*, **89** of *Computational Intelligence*. Springer, 2008. 33, 40, 41
- [81] H. HATTORI, M. KLEIN, AND T. ITO. **Using iterative narrowing to enable multi-party negotiations with multiple interdependent issues**. In *Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems (AAMAS'07)*, pages 1–3, New York, NY, USA, 2007. ACM. 34, 35, 40, 41
- [82] I. MARSA-MAESTRE, M. A. LOPEZ-CARMONA, J. R. VELASCO, AND E. DE LA HOZ. **Effective bidding and deal identification for negotiations in highly nonlinear scenarios**. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'09)*, pages 1057–1064, Richland, SC, 2009. International Foundation for Autonomous Agents and Multiagent Systems. 35, 40, 41, 73
- [83] I. MARSA-MAESTRE, M. A. LÓPEZ-CARMONA, J. R. VELASCO, T. ITO, M. KLEIN, AND K. FUJITA. **Balancing Utility and Deal Probability for Auction-Based Negotiations in Highly Nonlinear Utility Spaces**. In *International Joint Conference on Artificial Intelligence*, pages 214–219, 2009. 35, 40, 41, 73
- [84] G. BROWDER. **An analysis of the negotiations for the 1995 Mekong Agreement**. *International Negotiation*, **5**(2):237–261, 2000. 36
- [85] K.M. O'CONNOR. **Groups and solos in context: The effects of accountability on team negotiation**. *Organizational Behavior and Human Decision Processes*, **72**(3):384–407, 1997. 37
- [86] M. J. GELFAND, J. M. BRETT, L. IMAI, H. TSAI, AND D. HUANG. **Team Negotiation across Cultures: When and Where Are Two Heads Better Than One?** *Proceedings of 18th Annual Conference for the International Association for Conflict Management (IACM-2005)*, 2005. 37
- [87] K. BEHFAR, R. A. FRIEDMAN, AND J. M. BRETT. **The Team Negotiation Challenge: Defining and Managing the Internal Challenges of Negotiating Teams**. In *Proceedings of the 21st Annual Conference for the International Association for Conflict Management (IACM-2008)*, 2008. 37, 45
- [88] E. MANNIX. **Strength in Numbers: Negotiating as a Team**. *Negotiation*, **8**(5), 2005. 37, 38
- [89] S. E. BRODT AND M. TUCHINSKY. **Working Together but in Opposition: An Examination of the Good-Cop/Bad-Cop Negotiating Team Tactic**. *Organizational Behavior and Human Decision Processes*, **81**(2):155 – 177, 2000. 37
- [90] NIR HALEVY. **Team Negotiation: Social, Epistemic, Economic, and Psychological Consequences of Subgroup Conflict**. *Personality and Social Psychology Bulletin*, **34**:1687–1702, 2008. 37, 38
- [91] H. TAYLOR. **Most people are privacy pragmatists who, while concerned about privacy, will sometimes trade it off for other benefits**. *The Harris Poll*, **17**:19, 2003. 44
- [92] J. HOLLAND. *Adaptation in natural and artificial systems*. University of Michigan Press, 1975. 48, 57
- [93] S. W. MAHFOUD. *Niching methods for genetic algorithms*. PhD thesis, University of Illinois at Urbana-Champaign, Champaign, IL, USA, 1995. 48, 57
- [94] M. OSBORNE AND A. RUBINSTEIN. *Bargaining and Markets*. The Academic Press, 1990. 51, 52
- [95] O. J. MENGSHOEL AND D. E. GOLDBERG. **The crowding approach to niching in genetic algorithms**. *Evolutionary Computation*, **16**(3):315–354, 2008. 57
- [96] C. TOMLIN, G.J. PAPPAS, AND S. SASTRY. **Conflict resolution for air traffic management: A study in multiagent hybrid systems**. *IEEE Transactions on Automatic Control*, **43**(4):509–521, 1998. 93
- [97] J. CHU-CARROLL AND S. CARBERRY. **Conflict resolution in collaborative planning dialogs**. *International Journal of Human-Computer Studies*, **53**(6):969 – 1015, 2000. 93
- [98] K.S. BARBER, T.H. LIU, AND S. RAMASWAMY. **Conflict detection during plan integration for multi-agent systems**. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, **31**(4):616–628, 2001. 93
- [99] S. WOLLKIND, J. VALASEK, AND T.R. IOERGER. **Automated conflict resolution for air traffic management using cooperative multiagent negotiation**. In *AIAA guidance, navigation, and control conference*, pages 1–11, 2004. 93

REFERENCES

- [100] F. LOPES, N. MAMEDE, A.Q. NOVAIS, AND H. COELHO. **A negotiation model for autonomous computational agents: Formal description and empirical evaluation.** *Journal of Intelligent & Fuzzy Systems*, **12**(3):195–212, 2002. 93
- [101] KWANG MONG SIM AND CHUNG YU CHOI. **Agents that react to changing market situations.** *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, **33**(2):188 – 201, 2003. 95, 112
- [102] C. LI, J. GIAMPAPA, AND K. SYCARA. **Bilateral negotiation decisions with uncertain dynamic outside options.** *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews*, **36**(1):31–44, 2006. 95, 112
- [103] B. AN, K. M. SIM, C. Y. MIAO, AND Z. Q. SHEN. **Decision making of negotiation agents using markov chains.** *Multiagent and Grid Systems*, **4**(1):5–23, 2008. 95, 112
- [104] K. MANSOUR AND R. KOWALCZYK. **A Meta-Strategy for Coordinating of One-to-Many Negotiation over Multiple Issues.** In *Foundations of Intelligent Systems*, **122**, pages 343–353. Springer Berlin / Heidelberg, 2012. 95
- [105] D. SARNE, E. MANISTERSKI, AND S. KRAUS. **Multi-goal economic search using dynamic search structures.** *Autonomous Agents and Multi-Agent Systems*, **21**:204–236, 2010. 95
- [106] E. DEL VAL, M. REBOLLO, AND V. BOTTI. **Self-Organized Service Management in Social Systems.** In *Proceedings of the 45 Hawaii International Conference on System Sciences (HICSS 2012)*, pages 810–817, 2012. 95
- [107] H. RAIFFA. *The Art and Science of Negotiation.* Cambridge, USA, 1982. 96
- [108] M. TSVETOVAT, K. SYCARA, Y. CHEN, AND J. YING. **Customer Coalitions in Electronic Markets.** In *Agent-Mediated Electronic Commerce III, 2003 of Lecture Notes in Computer Science*, pages 121–138. Springer Berlin / Heidelberg, 2001. 96, 97
- [109] J. YAMAMOTO AND K. SYCARA. **A stable and efficient buyer coalition formation scheme for e-marketplaces.** In *Proceedings of the fifth international conference on Autonomous agents (AGENTS'01)*, pages 576–583, New York, NY, USA, 2001. ACM. 96, 97
- [110] T. ITO, H. OCHI, AND T. SHINTANI. **A group-buy protocol based on coalition formation for agent-mediated e-commerce.** *International Journal of Co-operative Information Systems*, **3**(1):11–20, 2002. 96, 97
- [111] C. LI, U. RAJAN, S. CHAWLA, AND K. SYCARA. **Mechanisms for coalition formation and cost sharing in an electronic marketplace.** In *Proceedings of the 5th international conference on Electronic commerce (ICEC'03)*, pages 68–77, New York, NY, USA, 2003. ACM. 96, 97
- [112] M. BELMONTE, R. CONEJO, J. PÉREZ-DE-LA CRUZ, AND F. TRIGUERO. **A Stable and Feasible Payoff Division for Coalition Formation in a Class of Task Oriented Domains.** In *Intelligent Agents VIII, 2333 of Lecture Notes in Computer Science*, pages 324–334. Springer Berlin / Heidelberg, 2002. 96
- [113] V. CONITZER AND T. SANDHOLM. **Computing shapley values, manipulating value division schemes, and checking core membership in multi-issue domains.** In *Proceedings of the 19th national conference on Artificial intelligence (AAAI 2004)*, pages 219–225. AAAI Press, 2004. 96
- [114] B. BLANKENBURG AND M. KLUSCH. **On Safe Kernel Stable Coalition Forming among Agents.** In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'04)*, pages 580–587. IEEE Computer Society, 2004. 96
- [115] N. OHTA, V. CONITZER, Y. SATOH, A. IWASAKI, AND M. YOKOO. **Anonymity-proof shapley value: extending shapley value for coalitional games in open environments.** In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems*, pages 927–934. International Foundation for Autonomous Agents and Multiagent Systems, 2008. 96
- [116] Y. ZICK AND E. ELKIND. **Arbitrators in overlapping coalition formation games.** In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'11)*, pages 55–62. International Foundation for Autonomous Agents and Multiagent Systems, 2011. 96
- [117] F. DIGNUM, B. DUNIN-KEPLICZ, AND R. VERBRUGGE. **Agent Theory for Team Formation by Dialogue.** In *Proceedings of the 7th International Workshop on Intelligent Agents VII. Agent Theories Architectures and Languages (ATAL'00)*, pages 150–166, London, UK, 2001. Springer-Verlag. 97
- [118] P. COHEN, H. LEVESQUE, AND I. SMITH. **On Team Formation.** In *Contemporary Action Theory. Synthese*, pages 87–114. Kluwer Academic Publishers, 1999. 97
- [119] M. E. GASTON AND M. DESJARDINS. **Agent-organized networks for dynamic team formation.** In *Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems (AAMAS'05)*, pages 230–237, New York, NY, USA, 2005. ACM. 97
- [120] J. FRENCH AND B. RAVEN. **The bases of social power.** In *Studies in social power*, pages 150–167. Ann Arbor, MI: Institute for Social Research, 1959. 98
- [121] B. REZABAKHSH, D. BORNEMANN, U. HANSEN, AND U. SCHRADER. **Consumer Power: A Comparison of the Old Economy and the Internet Economy.** *Journal of Consumer Policy*, **29**:3–36, 2006. 98

REFERENCES

- [122] A. PERREAU DE PINNINCK, C. SIERRA, AND M. SCHORLEMMER. **Distributed Norm Enforcement: Ostracism in Open Multi-Agent Systems.** In *Computable Models of the Law*, 4884 of *Lecture Notes in Computer Science*, pages 275–290. Springer Berlin / Heidelberg, 2008. 98
- [123] J. BRZOSTOWSKI AND R. KOWALCZYK. **Experimental Evaluation of Possibilistic Mechanism for Negotiation Partners Selection.** In *Rational, Robust, and Secure Negotiations in Multi-Agent Systems*, 89, pages 127–145. Springer Berlin / Heidelberg, 2008. 99, 100
- [124] J. SCHAFER, D. FRANKOWSKI, J. HERLOCKER, AND S. SEN. **Collaborative Filtering Recommender Systems.** In *The Adaptive Web*, 4321 of *Lecture Notes in Computer Science*, pages 291–324. Springer Berlin / Heidelberg, 2007. 99
- [125] S. KOC-MÉNARD. **Team performance in negotiation: a relational approach.** *Team Performance Management*, 15(7-8):357–365, 2009. 100
- [126] I. RAHWAN, S.D. RAMCHURN, N.R. JENNINGS, P. MCBURNEY, S. PARSONS, AND L. SONENBERG. **Argumentation-based negotiation.** *The Knowledge Engineering Review*, 18(4):343–375, 2003. 100, 102, 105
- [127] P. MCBURNEY, D. HITCHCOCK, AND S. PARSONS. **The eightfold way of deliberation dialogue.** *International Journal of Intelligent Systems*, 22(1):95–132, 2007. 100, 102, 105
- [128] G. BREWKA AND T. EITER. **Argumentation Context Systems: A Framework for Abstract Group Argumentation.** In *Logic Programming and Non-monotonic Reasoning*, 5753 of *Lecture Notes in Computer Science*, pages 44–57. Springer Berlin / Heidelberg, 2009. 100, 102, 105
- [129] S. HERAS, V. BOTTI, AND V. JULIAN. **Argument-based Agreements in Agent Societies.** *Neurocomputing*, 75(1):156–162, 2012. 100, 102, 105
- [130] S. SA AND J. ALCANTARA. **Cooperative Dialogues with Conditional Arguments.** In *Proceedings of the 11th International Joint Conference on Autonomous Agents and Multiagent Systems*. ACM, 2012. 100, 102, 105
- [131] H. NURMI. **Voting systems for social choice.** *Handbook of Group Decision and Negotiation*, pages 167–182, 2010. 100, 105, 132, 140, 207
- [132] K.P. SYCARA AND M. LEWIS. **Modeling teams of specialists.** In *Proceedings of the 27 Hawaii International Conference on System Sciences (HICSS 1994)*, 4, pages 271–280, 1994. 102
- [133] R. BOOTH. **Social contraction and belief negotiation.** *Information Fusion*, 7(1):19–34, 2006. 102
- [134] G. QI, W. LIU, AND D. BELL. **Combining multiple prioritized knowledge bases by negotiation.** *Fuzzy Sets and Systems*, 158(23):2535–2551, 2007. 102
- [135] S. KONIECZNY AND R. PINO PÉREZ. **Logic based merging.** *Journal of Philosophical Logic*, 40(2):239–270, 2011. 102, 103
- [136] P. EVERAERE, S. KONIECZNY, P. MARQUIS, ET AL. **The strategy-proofness landscape of merging.** *Journal of Artificial Intelligence Research*, 28(1):49–105, 2007. 103
- [137] J. M. SUCH, A. GARCÍA-FORNES, AND A. ESPINOSA. *Enhancing Privacy in Multi-agent Systems: Avoiding undesired Information Collection and Information Processing.* LAP LAMBERT Academic Publishing, 2012. 103
- [138] M. KLEIN, P. FARATIN, H. SAYAMA, AND Y. BAR-YAM. **Protocols for Negotiating Complex Contracts.** *IEEE Intelligent Systems*, 18(6):32–38, 2003. 105, 106
- [139] G. LAI, C. LI, AND K. SYCARA. **Efficient Multi-Attribute Negotiation with Incomplete Information.** *Group Decision and Negotiation*, 15:511–528, 2006. 105, 106
- [140] T. MILLER AND P. MCBURNEY. **Annotation and matching of first-class agent interaction protocols.** *Argumentation in Multi-Agent Systems*, pages 141–158, 2009. 105
- [141] J. HULSTIJN, M. DASTANI, AND L. VAN DER TORRE. **Negotiation protocols and dialogue games.** In *Proceedings of the Belgian-Dutch AI Conference (BNAIC-2000)*, PUBLISHER =, 2000. 106
- [142] CHRIS REED, TIMOTHY J. NORMAN, AND NICHOLAS R. JENNINGS. **Negotiating the semantics of agent communications languages.** *Computational Intelligence*, 18(2):229–252, 2002. 106
- [143] R. NAIR, T. ITO, M. TAMBE, AND S. MARSELLA. **Task Allocation in the RoboCup Rescue Simulation Domain: A Short Note.** In *RoboCup 2001: Robot Soccer World Cup V*, pages 751–754, London, UK, 2002. Springer-Verlag. 108
- [144] R. NAIR, M. TAMBE, AND S. MARSELLA. **Role allocation and reallocation in multiagent teams: Towards a practical analysis.** In *Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, pages 552–559. ACM, 2003. 108
- [145] M. HOOGENDOORN AND C. JONKER. **Formation of Virtual Organizations Through Negotiation.** In *Multiagent System Technologies*, 4196 of *Lecture Notes in Computer Science*, pages 135–146. Springer Berlin / Heidelberg, 2006. 108
- [146] A. CAMPBELL AND A. WU. **Multi-agent role allocation: issues, approaches, and multiple perspectives.** *Autonomous Agents and Multi-Agent Systems*, 22:317–355, 2011. 108
- [147] B. J. GROSZ, S. KRAUS, D. G. SULLIVAN, AND S. DAS. **The influence of social norms and social consciousness on intention reconciliation.** *Artificial Intelligence*, 142(2):147 – 177, 2002. 108

REFERENCES

- [148] P. PARUCHURI, M. TAMBE, S. KAPETANAKIS, AND S. KRAUS. **Between collaboration and competition: An initial Formalization using Distributed POMDPs.** In *Proceedings of the GTDT workshop at AAMAS-2003*, 2007. 108
- [149] N. MATOS AND C. SIERRA. **Evolutionary computing and negotiating agents.** In *Workshop on Agent Mediated Electronic Trading (AMET'98)*, pages 91–111, 1998. 109, 112
- [150] RAZ LIN, SARIT KRAUS, TIM BAARSLAG, DMYTRO TYKHONOV, KOEN HINDRIKS, AND CATHOLIJN M. JONKER. **Genius: An Integrated Environment for Supporting the Design of Generic Automated Negotiators.** *Computational Intelligence*, 2012. 109, 216
- [151] T. BAARSLAG, K. V. HINDRIKS, AND C. M. JONKER. **Towards a Quantitative Concession-Based Classification Method of Negotiation Strategies.** In *Agents in Principle, Agents in Practice. Lecture Notes of The 14th International Conference on Principles and Practice of Multi-Agent Systems*, pages 143–158, 2011. 109, 218
- [152] B.I. SPECTOR. **Creativity heuristics for impasse resolution: Reframing intractable negotiations.** *The Annals of the American Academy of Political and Social Science*, 542(1):81–99, 1995. 111
- [153] L. BABCOCK AND G. LOEWENSTEIN. **Explaining bargaining impasse: The role of self-serving biases.** *The Journal of Economic Perspectives*, 11(1):109–126, 1997. 111
- [154] R.D. TOLLISON AND T.D. WILLETT. **An economic theory of mutually advantageous issue linkages in international negotiations.** *International Organization*, 33(4):425–449, 1979. 111
- [155] J.K. SEBENIUS. **Negotiation arithmetic: adding and subtracting issues and parties.** *International Organization*, 37(2):281–316, 1983. 111
- [156] T.C. MORGAN. **Issue linkages in international crisis bargaining.** *American Journal of Political Science*, 34(2):311–333, 1990. 111
- [157] C. HOU. **Predicting agents tactics in automated negotiation.** In *Proceedings of the IEEE/WIC/ACM International Conference on Intelligent Agent Technology (IAT 2004)*, pages 127–133. IEEE, 2004. 112
- [158] C.R. WILLIAMS, V. ROBU, E.H. GERDING, AND N.R. JENNINGS. **Using gaussian processes to optimise concession in complex negotiations against unknown opponents.** In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI'11)*, pages 432–438, 2011. 112
- [159] C. YU, F. REN, AND M. ZHANG. **An Adaptive Bilateral Negotiation Model Based on Bayesian Learning.** In *The 4th International Workshop on Agent-based Complex Automated Negotiations (ACAN 2011)*, pages 40–47, 2011. 112
- [160] YANN CHEVALEYRE, PAUL E. DUNNE, ULLE ENDRISS, JÉRÔME LANG, MICHEL LEMAÎTRE, NICOLAS MAUDET, JULIAN PADGET, STEVE PHELPS, JUAN A. RODRÍGUEZ-AGUILAR, AND PAULO SOUSA. **Issues in Multiagent Resource Allocation.** *Informatica*, 30:2006, 2006. 167
- [161] T. BAARSLAG, K. HINDRIKS, C. JONKER, S. KRAUS, AND R. LIN. **The first automated negotiating agents competition (ANAC 2010).** *New Trends in Agent-Based Complex Automated Negotiations*, pages 113–135, 2010. 216
- [162] S. KAWAGUCHI, K. FUJITA, AND T. ITO. **Compromising Strategy Based on Estimated Maximum Utility for Automated Negotiation Agents Competition (ANAC-10).** In *Modern Approaches in Applied Intelligence*, 6704, pages 501–510. Springer Berlin / Heidelberg, 2011. 218
- [163] K.V. HINDRIKS, C.M. JONKER, AND D. TYKHONOV. **The Benefits of Opponent Models in Negotiation.** In *Proceedings of the 2009 IEEE/WIC/ACM international joint conference on Web Intelligence and Intelligent Agent Technology*, pages 439–444, 2009. 219
- [164] V. SANCHEZ-ANGUIX, T. DAI, Z. SEMNANI-AZAD, K. SYCARA, AND V. BOTTI. **Modeling power distance and individualism/collectivism in negotiation team dynamics.** In *45 Hawaii International Conference on System Sciences (HICSS-45)*, pages 628–637, 2012. 239