

A Baseline for Non-Linear Bilateral Negotiations

The full results of the agents competing in ANAC 2014

Reyhan Aydoğan^{1,2}, Catholijn M. Jonker², Katsuhide Fujita³, Tim Baarslag⁴,
Takayuki Ito⁵, Rafik Hadfi⁵, and Kohei Hayakawa⁵

¹ Computer Science, Özyeğin University, Istanbul, Turkey

² Interactive Intelligence Group, Delft University of Technology, The Netherlands

³ Faculty of Engineering, Tokyo University of Agriculture and Technology, Tokyo, Japan

⁴ Agents, Interaction and Complexity Group, University of Southampton, Southampton, UK

⁵ Department of Computer Science and Engineering, Nagoya Institute of Technology, Japan

reyhan.aydogan@ozyegin.edu.tr, C.M.Jonker@tudelft.nl,
katfuji@cc.tuat.ac.jp, tblm13@ecs.soton.ac.uk,
ito.takayuki@nitech.ac.jp, rafik@itolab.nitech.ac.jp,
hayakawa.kohei@itolab.nitech.ac.jp

Abstract. In the past few years, there is a growing interest in automated negotiation in which software agents facilitate negotiation on behalf of their users and try to reach joint agreements. The potential value of developing such mechanisms becomes enormous when negotiation domain is too complex for humans to find agreements (e.g. e-commerce) and when software components need to reach agreements to work together (e.g. web-service composition). Here, one of the major challenges is to design agents that are able to deal with incomplete information about their opponents in negotiation as well as to effectively negotiate on their users' behalves. To facilitate the research in this field, an automated negotiating agent competition has been organized yearly. This paper introduces the research challenges in ANAC 2014 and explains the competition set up and competition results. Furthermore, a detailed analysis of the best performing five agent has been examined.

Keywords: Automated Negotiation, Nonlinear Utility Functions, Agent Competition

1 Introduction

Conflict is an omnipresent phenomenon in human society [42, 38, 15]. It spans from daily situations like discussing a holiday plan with friends and arranging a meeting between colleagues to complex scenarios like politics and business. Automated negotiation tools provide an important mechanism for decision makers to resolve their conflicts and to reach mutually acceptable agreements. There is a growing interest and need for automated negotiation mechanisms [15, 30]. To facilitate the research in automated negotiation, an international competition namely Automated Negotiating Agent Competition (ANAC)⁶ [6] is yearly organized. This competition challenges researchers to

⁶ <http://ii.tudelft.nl/anac>

design and develop fully-automated negotiation agents that can negotiate under certain protocols and conditions. To compete, one has to develop a negotiation agent that can negotiate across a variety of negotiation scenarios. In all variants of the competition, the agents have to negotiate with incomplete knowledge - agents do not know neither their opponents preferences nor their negotiation strategy.

From May 2010 to May 2013 four instances of the ANAC competition [4, 7, 9, 48] have been held in conjunction with the International Conference on Autonomous Agents and Multiagent Systems (AAMAS). This competition follows in the footsteps of a series of successful competitions that aim to advance the state-of-the-art in artificial intelligence (other examples include the Annual Computer Poker Competition and the various Trading Agent Competitions (TAC) [46]). ANAC focuses specifically on the design of practical negotiation strategies. In particular, the overall aim of the competition is to advance the state-of-the-art in the area of bilateral, multi-issue negotiation, with an emphasis on the development of successful automated negotiators in realistic environments with incomplete information (where negotiators do not know their opponent's strategy, nor their preferences) and continuous time (where the negotiation speed and number of negotiation exchanges depends on the computational requirements of the strategy). One of the successes of ANAC lies in the development of state-of-the-art negotiation strategies that co-evolve every year; some notable examples include [1, 8, 10–12, 16, 17, 26, 27, 34, 40, 47]. The previous four incarnations of ANAC already yielded more than 60 new strategies and scenarios [4] which can be used as benchmarks to test the efficacy of subsequent work in this area. This is in line with the goal of ANAC to build a community in which work on such negotiating agents can be compared by standardized negotiation benchmarks and performance metrics.

Each year, the organizers bring up a new challenge for the participants. This year, ANAC 2014's focus is on negotiating with nonlinear utility functions as well as dealing with large-scale outcome spaces [18]. In ANAC 2014, negotiating agents were not allowed to access the structure of the nonlinear utility functions directly; therefore, they needed to explore the outcome space smartly to generate their bids. The main challenge in ANAC 2014 is to explore a large-scale outcome space effectively. In this paper, we explain the competition setup and results of the qualification and final rounds. Furthermore, we analyze the performance of the best performing five agents in an additional experimental set up. Following the competition, we systematically generated 108 different negotiation scenarios and tested the performance of the best performing five agents elaborately. We studied the effect of the domain size, the number of constraints and issue-constraint distribution on the performance of the agents.

Our experimental results show that the agent applying a Genetic Algorithm namely *Gangster* outperforms the other agents in terms of the individual utility gained by the agents. The performance of the agents using Simulated Annealing in their strategy, namely *Agent M* and *Whale Agent* is closed to the performance of *Gangster* agent. On average, the performance of *Agent M* with respect to the distance to Pareto optimal outcome and Nash outcome is slightly better than *Gangster* agent. Their overall performance regarding to the utilitarian social welfare metrics is almost the same. Moreover, our evaluation results also show that the performance of the agents highly depends on their opponents and the negotiation scenarios played. For example, *DoNA* agent, which

was ranked second place in ANAC 2014, was outperformed by all agents in the second experiment. In the second experiment DoNA had less opponents and in the second experiment a greater variety of scenarios was used.

The rest of this paper is organized as follows: Section 2 gives an introduction on ANAC 2014 rules and competition setup. Section 3 provides an overview of agents participated in ANAC 2014. Section 4 explains ANAC 2014 qualification and final round results while Section 5 provides a detailed analysis of best performing five ANAC 2014 agents. Finally, Section 6 concludes the paper with directions to future work.

2 ANAC 2014

Each year, the ANAC organizers bring a new challenge to the participants (e.g. discount factors in ANAC 2011 [7, 19], private reservation value for each agent in ANAC 2012 [48], learning from past negotiation sessions in ANAC 2013 [50]). The focus of ANAC 2014 is on bilateral multi-issue negotiation, in which each agent's preferences are represented in the form of nonlinear utility functions as well as dealing with large-scale negotiation domains. In such complex domains, exploring the outcome space is not as straightforward as it was in the former competition, in which linear additive utility functions were used on relatively small-size domains. To sum up, the main challenge in ANAC 2014 is to find efficient exploration strategies for interdependent preferences, particularly on large-size domains.

As in previous years, the General Environment for Negotiation with Intelligent multi-purpose Usage Simulation (GENIUS)⁷ [36] has been used in ANAC 2014. GENIUS is a research tool for automated multi-issue negotiation that enables ANAC participants to develop and to test their negotiating agents. It also provides an easily accessible framework to develop negotiating agents via a public API. Each participating team has to design and build a negotiation agent using the GENIUS. For ANAC 2014, we extended the GENIUS framework with nonlinear utility functions in the form of weighted hypercubes [28, 38].

2.1 ANAC 2014 Rules

The aim for the entrants to the competition is to develop an autonomous negotiation agent. Performance of the agents is evaluated in a tournament setting, where each agent is matched with all other submitted agents, and each pair of agents will negotiate in a number of negotiation scenarios. There is an initial qualifying round, and the top 8 performing agents will continue to the finals, which is held at the AAMAS 2014 conference. All teams that make it through to the finals have a representative attending the AAMAS 2014 conference, who has the opportunity to give a brief presentation describing their agent.

In the previous years, additive utility functions were used to represent negotiating agents' preferences in ANAC. Since realistic negotiation scenarios involve bundles of interdependent issues under one or more specific constraints, as a main update with respect to the previous year, ANAC 2014 extends the agents' utility model to *nonlinear*

⁷ <http://ii.tudelft.nl/genius>

utility functions (see Section 2.2). When generating scenarios, we should not restrict the interdependencies to low-order constraints but allow multidimensional interdependencies to be parametrically defined. One way of doing it is to generate a utility spaces as an hypergraph with connectivities defined as hyperedges connecting a number of issues [25]. Such approach allows a quantitative assessment of the complexity of any nonlinear utility space using information entropy [23]. Moreover, using a graph-theoretic analysis of the complexity gives a good estimate of the performance of any optimization algorithm. Therefore, we use weighted hypercubes to represent agents' preferences.

In these scenarios, the agents no longer have linear utility functions; instead, they can only sample their utility of a bid through a `getUtility()` method. In terms of the agent design, this means that the agents do not have access to methods pertaining to scenarios, such as functionality specifying the weights of the negotiation issues. The participants were not allowed to exploit the structure of the nonlinear preferences, so that it becomes a challenge to explore the outcome space efficiently. The agents have to search the utility space bid-by-bid, using the `getUtility()` method for any bid they are interested in. Another challenge is to deal with large-size domains, with outcome spaces as big as 10^{50} outcomes.

Negotiations are bilateral and based on the alternating-offers protocol. Offers are exchanged in real time with a deadline after 3 minutes. In addition, half of the domains contain a discount factor, which causes the value of an agreement to decrease over time. The challenge for an agent is to negotiate without any prior knowledge of the opponent's preferences and strategy. Negotiations are repeated several times to obtain statistically significant results. Agents can be disqualified for violating the spirit of fair play. The competition rules allow multiple entries from a single institution, but require each agent to be developed independently.

2.2 Negotiation Scenarios

In all ANAC competitions the scenarios form the setting in which the participating agents compete with each other. The agents play against each other over a number of scenarios. In ANAC 2014 setting, a negotiation scenario consists of domain description and two conflicting nonlinear preference profiles. There exists a number of representation used in automated negotiation [36, 2, 45, 3]. For convenience, only integer issues are considered for domain specification and each agent's preferences are represented in the form of weighted hypercubes [28, 29]. According to this representation, a utility function consists of a set of hypercube regions in the outcome space. Each hypercube represents a single constraint c_k . A numeric *weight* or *utility value* $u(c_k)$ is associated to each constraint. The utility of a given bid, o is calculated as sum of the utility values for the hypercubes including this bid, as follows:

$$u(o) = \sum_{c_k \in C | Satisfy(o, c_k)} u(c_k). \quad (1)$$

Figure 1 shows a sample utility function for a two-issue negotiation problem. This utility function consists of a unary constraint $C1$ and two binary constraints $C2$ and $C3$. The corresponding utility values associated to these constraints are 5, 10 and 12

respectively. According to this example, the contract x ($issue_1=2$; $issue_2=3$) would yield a utility value $u(x) = 15$ for the agent, since it satisfies both $C1$ and $C2$ (that is, constraints $C1$ and $C2$ overlap, creating a region of higher utility). The contract y ($issue_1=4$; $issue_2=2$), on the other hand, would yield a utility value $u(y) = 5$, because it only satisfies $C1$.




Fig. 1. Example of a utility space with two issues and three constraints

It is worth noting that we extend this representation to support negated *constraints*, which are satisfied when the given bid remains outside of the given hypervolume. This richer expressiveness allows us to approximate any desired function, by a finite number of constraints[38].

Figure 2 shows an example of a nonlinear utility space. There are two issues, *i.e.*, two dimensions, with domains $[0, 99]$. There are 50 unary constraints (*i.e.*, that relate to 1 issue) as well as 100 binary constraints (*i.e.*, that inter-relate 2 issues). The utility space is, as we can see, highly nonlinear, with many hills and valleys. It will generally be the case, in fact, that agents do not fully know their desirable contracts in advance because each own utility functions are simply too large and complex. If we have 10 issues with 10 possible values per issue, for example, this produces a space of 10^{10} (10




Fig. 2. Utility Space for a Single Agent

billion) possible contracts, too many to evaluate exhaustively. Agents must thus operate in a highly uncertain environment.

In addition, one of the key points in achieving automated negotiation in real life has been the non-linearity and size of the domains. Many real-world negotiation problems sometimes assume the nonlinear and large domains. When an automated negotiation strategy covers the linear function effectively, it is not always possible or desirable in nonlinear situations [35]. Therefore, the constraint-based nonlinear utility function with integer issues has been the one of the important topic in automated negotiations [28, 29]. The existing work analyzes and defines some measures for the constraint-based nonlinear utility function [20], and other existing works try to improve the effectiveness of finding contracts in the bumpy nonlinear utility function [21, 22, 37, 39].

Typically, the characteristics of a negotiation scenario such as size, number of issues, opposition, discount factor and reservation value can greatly influence the negotiation outcome, [5]. Therefore, also in the ANAC 2014 competition, we systematically varied some of these factors. Given our focus on the scale of the domains and the non-linearity, we generated scenarios to differ in terms of number of issues, number of possible proposals, opposition of preference profiles and mean distance of all points in the outcome space to the Pareto Frontier. We decided to use 12 negotiation scenarios for ANAC 2014 as described in Table 1. In more detail, three types of negotiation domains with varying size (e.g., moderate, large and very large) were defined. For each domain category, different discount factors and reservation values have been allocated (See Table 1). For each integer issue, 10 possible integer values are possible. If the number of issues is equal to 50, this domain has 10^{50} possible outcomes. In the rest of the paper, we will use the following domain categories:

- Moderate Domain: consists of 10 issues - 10^{10} outcomes
- Large Domain: consists of 30 issues - 10^{30} outcomes
- Very Large Domain: consists of 50 issues - 10^{50} outcomes

The utility spaces of three different scenarios are graphically represented in Figure 3.

Table 1. Negotiation Domains in ANAC 2014

ID	Number of Issues	Size of Outcome Space	Discount Factor	Reservation Value
1	10	10^{10}	None	None
2	10	10^{10}	0.50	None
3	10	10^{10}	None	0.75
4	10	10^{10}	0.50	0.75
5	30	10^{30}	None	None
6	30	10^{30}	0.50	None
7	30	10^{30}	None	0.75
8	30	10^{30}	0.50	0.75
9	50	10^{50}	None	None
10	50	10^{50}	0.50	None
11	50	10^{50}	None	0.75
12	50	10^{50}	0.50	0.75

2.3 Competition Setup

The tournament platform for running and analyzing a variety of negotiations was (GENIUS) [36]. Some sample negotiation scenarios were provided to the participants. However, the negotiation scenarios used in the competition were not known by the participants in advance. The success of each negotiating agent is measured by considering the results of all negotiations within the underlying tournament.

Given the complexity of the negotiation domains, time limitation (from submission to having the final results available at the AAMAS 2014 conference that hosted our competition, and the available computers for running the tournament, it was infeasible to set up a single tournament for the expected number of agents. The expected number of agents was around 18 based on the emails we received that expressed an interest to participate. Running a full tournament for 19 agents, where each agent negotiates against 18 other agents in all negotiation scenarios, requires at least 1026 hours (19×18 different combinations \times 3 minutes per negotiation \times 12 scenarios \times 5 repetitions = 1026 hours). Therefore, ANAC 2014 consisted of two phases: a qualifying round and a final round. The aim of the qualifying round is to determine the top 10 agents that will compete in the final round. In total 21 agents were submitted from 13 different institutions in 8 different countries. Before starting the qualifying round, all submitted

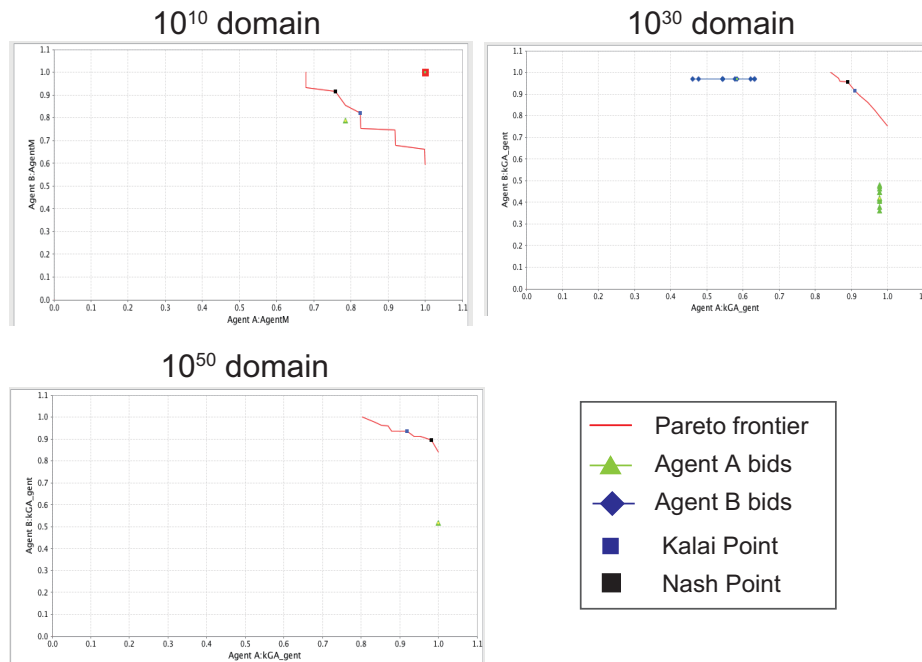


Fig. 3. Utility Spaces of Three ANAC 2014 Negotiation Scenarios

agents were tested and two of them were disqualified as their code did not run correctly. Therefore, 19 agents were considered in the qualifying round. These agents were distributed to three groups (pools) randomly for the qualifying round. It took two weeks to complete three tournaments in which approximately 6 or 7 agents negotiated with each other. From each pool, the best three agents with respect to social welfare and individual utility were qualified for the final round. From one pool, we picked the best four agents; that makes 10 in total. It is worth mentioning that in ANAC 2014, the participants were not allowed to update their agents after the qualifying round.

In order to complete such an extensive set of tournaments within a limited time frame, we used five high-spec computers, made available by Nagoya Institute of Technology and Tokyo University of Agriculture and Technology, Japan. Specifically, each of these machines contained an *Intel Core i7* CPU, at least 16GB of DDR3 memory, and a hard drive with at least 2TB of capacity.

In the final round, we ran a single tournament for the qualified 10 agents. We considered the utilitarian social welfare criterion, the sum of each agent's individual utility. The same negotiation scenarios were used in both qualifying and final rounds. The single tournament consisted of $10 \times 9 \times 12 = 1080$ negotiation sessions. Each agent pair negotiated twice (i.e., switching their roles) since the role of the negotiators might change the negotiation outcome. Each negotiation session was repeated five times.

3 ANAC 2014 Agents

In total 21 agents were submitted from 13 different institutions in 8 different countries. Two of them were disqualified as their code does not run correctly. Among 19 agents, the best 10 agents with respect to social welfare and individual utility were qualified for the final. Table 2 lists the ANAC 2014 finalists. To prevent the manipulation, the participants were not allowed to submit more than one agent. Multiple submissions from a single institution were allowed if only each agent would be developed independently by different team members.

Table 2. Finalists of ANAC 2014

Agent Name	Affiliation	Team Members
<i>AgentM</i>	Nagoya Institute of Technology	Makoto Niimi
<i>AgentYK</i>	Shizuoka University	Yoshiaki Kadono Farhad Zafari
<i>BraveCat</i>	University of Isfahan	Faria Nassiri-Mofakham
<i>DoNA</i>	Bar Ilan University	Eden Shalom Erez Inon Zuckerman
<i>E2Agent</i>	Ariel University	
<i>E2Agent</i>	Nagoya Institute of Technology	Yuichi Enoki
<i>Gangster</i>	IIIA-CSIC Barcelona	Dave de Jonge Balint Szollosi-Nagy Marta Skarzynska
<i>Group2Agent</i>	Delft University of Technology	David Festen
<i>kGAgent</i>	Tokyo University of Agriculture and Technology	Shinji Kakimoto Katsuhide Fujita
<i>Sobut</i>	University of Electro-Communications	Satoshi Takahashi
<i>WhaleAgent</i>	Nagoya Institute of Technology	Motoki Sato

In the rest of this section, we provide, in alphabetical order, a brief description of the individual strategies of the finalists based on the provided descriptions by the teams.

***AgentM* [41]** has a cooperative and compromising strategy considering the negotiation time and the difference between the best and worst bid offered by the opponent. The bidding strategy is aim to improve the social welfare using Simulated Annealing (SA) considering the frequency of the opponent's bids for each issue. While generating its bid, this agent considers its own best bid, the best bid offered by its opponent and the most frequently asked bid by its opponent so far. The acceptance of the opponent's bid is decided whether or not the utility of the opponent's bid is higher than the lowest utility asked by the agent itself during the negotiation.

***AgentYK* [31]** aims to generate bids increasing the social welfare by using *BidElement* and *PairBidElement*, so that they are more likely to be accepted by the opponent. According to its bidding strategy, this agent evaluates to what extent the given bid would be preferred by the opponent as well as considering its own past compromises. Typically,

it evaluates the similarity of a given bid with the opponent's previous bids. This agent tries to make bids that are more likely to be accepted by the opponent by considering the co-occurrences in the history of the negotiation. It is also inclined to make more compromised proposals at most three times through the end of the negotiation session.

BraveCat [51] employs a hybrid bidding strategy called RBT strategy (R: Random, B: Behavioral, Time dependent), which is a mixture of a random strategy, a time dependent strategy, and a behavioral strategy. *BraveCat* uses a new distance based opponent model to estimate the utility of a candidate bid to be sent to the opponent in each round of the negotiation. Moreover, by using iterative deepening search, *BraveCat* overcomes the limitations imposed by the huge amount of memory needed in the large domains.

DoNA [14] employs a domain-based approach using behavioral strategies based on solely two domain parameters: reservation value and discount factor. These two parameters are used to divide the class of all possible domains into different regions, in each of which employs a predefined strategy relying on a behavioral intuition. It considers the state space of domains divided into a 3-by-3 grid of combination of values. DoNA agent applies a cognitive model based on the analysis regarding the time allocated to the negotiations and the concession stance for the given reservation value and discount factor. It computes the minimum acceptable utility value by random sampling of offers where the distribution of opponent's bids is estimated by a normal distribution as follows: $argmin = 1.2 * \mu + 2.4 * \sigma$ (where σ is the standard deviation, and μ is the average).

E2Agent is an extension of *Agent K* [33], winner of the ANAC 2010 competition. When creating a counter offer *Agent K* calculates a target utility U_t based on the previous offers made by the opponent and the remaining negotiation time. *Agent K* then makes random bids above the target utility. If there is no bid with the target utility, the target utility is lowered slightly. The target utility U_t at time t is calculated using the following formula:

$$U_t = 1 - (1 - E_{\max}(t)) \cdot t^\alpha, \quad (2)$$

where $E_{\max}(t)$ is the estimated maximum value the opponent will present in the future based on the average and variance of previous bids, and α is a parameter which controls the concession speed. *E2Agent* uses a sophisticated acceptance mechanism, where it will use the average and variations of the previous bid utilities presented by the opponent in order to determine the best possible bid it can expect in the future. It will either accept or reject the offer based on the probability that the opponent will present a better offer in the future. If it has already received an offer from the opponent with the same utility or higher, it will offer that bid instead. In addition, it searches for bids with Simulated Annealing to find the bids effectively since it difficult for *AgentK* in the non-linear domains.

Gangster [13] applies a Genetic Algorithm (GA) to explore a large-scale bid space associated with nonlinear utilities. It calculates a target utility for the agent. It decides

whether to accept its opponent's bid by comparing it with the target utility. It applies a global GA to get the samples of the agreement space and stores 10 of the bids with the highest utilities and it also applies a local GA similarly. It proposes the best proposal in the current round or any of the previous rounds. *Gangster* gets some advantages of employing a GA with Manhattan distance and the combination of local and global GAs.

Group2Agent [44] *Group2Agent* tries to find bids with the high social welfare in the complex domains. It uses the Greedy Coordinate Descent (GCD) algorithm[49], which can scale linearly with the number of issues and is shown to be effective, locating meaningful middle ground between negotiating parties. It can gain the high average social welfare of 1.79, being only 0.03 below the optimal social welfare solution and found the optimal solution itself 3 out of 25 times. Furthermore, GCD scales better than algorithms such as SA used in other agents.

k-GAgent [32] applies the *SPEA2*[52], a method based on GA for multiple objective optimization, to explore the Pareto Frontier effectively. To estimate the opponent's utility function, the opponent's bid is divided into small parts considering the combinations of issues, and it considers how many times of these parts occurs in opponent's past bids.

Sobut employs a simple martingale bidding method used in the casino. If it wins the previous game, it tries to choose the minimum bid utility using the probability of winning (P) in next game. If it loses, it tries to choose the default value in the next game. In other words, the P and the utility of the minimum bid are decided according to the following equations.

$$P = \max(\text{previous outcome}, \text{reserve value})$$

$$\text{Utility of minimum bid} = \max(P, r)$$

It uses a random function to decide r , and the range of r is specified as $(0, 0.5]$.

WhaleAgent [43] has two bidding strategies: Hard headed and Conceder. According to the hardheaded strategy, it offers a bid whose utility is greater than 0.9 at the beginning of negotiations. Through the end of the negotiation, it decreases the threshold of accepting its opponent's bids. If the utility of the opponent's offer is greater than this threshold, it accepts that bid. The threshold is equivalent to the value used in Hard-headed strategy and Conceder strategy. The agent is inclined to accept any offer in the end of the negotiation in order avoid to fail the negotiations. To search the bid space with nonlinear utility values, the agent employs a search strategy that combines Simulated Annealing and Hill Climbing. First, it explores bids whose utility is greater than the threshold by using the Hill Climbing strategy. When it cannot find such a bid, it uses the Simulated Annealing using a random starting point.

Agent Name	Mean (Individual)	Rank (Individual)	Mean (Social welfare)	Rank (Social)
E2Agent	0.60449771	1	1.467013776	1
GROUP2Agent	0.569022057	2	1.309827507	2
kGA_gent	0.567855409	3	1.253293459	4
Sobut	0.514388859	4	1.25826658	3
ArisawaYaki	0.502270746	5	1.216825393	6
Simple ANAC2013	0.498502294	6	1.219932496	5

Fig. 4. Average scores of each agent in the qualifying round (pool1)

Agent Name	Mean (Individual)	Rank (Individual)	Mean (Social welfare)	Rank (Social)
Gangster	0.694014347	1	1.596774909	2
WhaleAgent	0.682003332	2	1.606191324	1
AgentYK	0.659956126	3	1.59018266	3
Flinch	0.649326182	4	1.573767682	4
AgentQuest	0.639303883	5	1.472990004	6
Simpaico	0.595949075	6	1.509024094	5

Fig. 5. Average scores of each agent in the qualifying round (pool2)

4 Results of ANAC 2014 Competition

Qualifying Round First, the *qualifying round* was played in order to select the finalists from the 19 agents that were submitted by the participating teams (2 agents were disqualified from the trial tests) 19 agents was divided to three groups (pools) randomly, and the best three agents in social welfare and individual utility in each pool proceeded to the final round. Each tournament wasn't repeated to prohibit the learning from the previous tournaments.

In order to complete such an extensive set of tournaments within a limited time frame, we used five high-spec computers, made available by Nagoya Institute of Technology and Tokyo University of Agriculture and Technology. Specifically, each of these machines contained an *Intel Core i7* CPU, at least 16GB of DDR3 memory, and a hard drive with at least 2TB of capacity.

Fig.4 - 6 show the results of each agent in the qualifying round (pool1, pool2, and pool3). The finalists are selected from all pools by considering the individual utilities and social welfare. The individual utility means the average of utility of the individual

Agent Name	Mean (Individual)	Rank (Individual)	Mean (Social welfare)	Rank (Social)
DoNA	0.668464329	1	1.285703724	2
AgentM	0.542950221	2	1.28268408	3
BraveCat v0.3	0.518940747	3	1.422961239	1
AgentTRP	0.484535552	4	1.119857699	5
Aster	0.479403688	5	1.112286696	6
AgentTD	0.464952079	6	1.168321409	4
Atlas	0.410946126	7	0.947732281	7

Fig. 6. Average scores of each agent in the qualifying round (pool3)

agent in the tournaments. The social welfare means the average of the sum of utilities of two agents in the tournaments. As figures showing, the best three or four agents are selected by considering the individual utility and social welfare. As a results, *kGAgent*, *E2Agent*, *GROUP2Agent*, *Sobut* are selected as finalists from the pool1; *Gangster*, *WhaleAgent*, *AgentYK* are selected as finalists from pool2; *DoNA*, *AgentM*, *BraveCat* are selected as finalists from pool3. They are the best three in each pool considering the individual utility or the social welfare.

Rank	Agent	Score	Variance
1	<i>Agent M</i>	0.754618239	3.12×10^{-5}
2	<i>DoNA</i>	0.742245035	9.31×10^{-6}
3	<i>Gangster</i>	0.740674889	3.49×10^{-6}
4	<i>WhaleAgent</i>	0.710740252	3.90×10^{-5}
5	<i>GROUP2Agent</i>	0.708401404	6.38×10^{-5}
6	<i>E2Agent</i>	0.703955008	2.85×10^{-5}
7	<i>kGAgent</i>	0.676595111	5.02×10^{-5}
8	<i>AgentYK</i>	0.666450943	2.38×10^{-5}
9	<i>BraveCat</i>	0.661940343	2.84×10^{-5}
10	<i>Sobut</i>	0.627684701	1.71×10^{-5}

Table 3. Tournament results in the final round (Individual Utility)

Final Round The final round consisted of 10 agents that were selected from the qualifying round. For each pair of agents, under each preference profile, we ran a total of some negotiations. By averaging over all the scores (individual utility and social welfare) achieved by each agent (against all opponents and using all preference profiles),

Rank	Agent	Score	Variance
1	<i>Agent M</i>	1.645412137	4.12×10^{-5}
2	<i>Gangster</i>	1.627451908	1.21×10^{-5}
3	<i>E2Agent</i>	1.608936143	1.39×10^{-5}
4	<i>WhaleAgent</i>	1.603199277	3.55×10^{-5}
5	<i>AgentYK</i>	1.569877186	1.16×10^{-4}
6	<i>GROUP2Agent</i>	1.56154598	8.46×10^{-5}
7	<i>BraveCat</i>	1.545384774	3.11×10^{-5}
8	<i>DoNA</i>	1.473686528	3.89×10^{-5}
9	<i>Sobut</i>	1.469972333	1.12×10^{-4}
10	<i>kGAgent</i>	1.463168543	4.32×10^{-4}

Table 4. Tournament results in the final round (Social Welfare)

the final ranking were decided. Formally, the average score $U_{\Omega}(p)$ of agent p in scenario Ω is given by:

$$U_{\Omega}(p) = \frac{\sum_{p' \in P, p \neq p'} U_{\Omega}(p, p') + U_{\Omega}(p', p)}{2 \cdot (|P| - 1)} \quad (3)$$

where P is the set of players and $U_{\Omega}(p, p')$ is the utility achieved by player p against player p' when player p is playing side A of Ω and player p' is under the side B of Ω . For the final round, we matched each pair of finalist agents, under each preference profile, for a total of 5 times.

It is notable that *AgentM* was the clear winner of both categories (see Table 3 and 4). However, the differences in utilities between many of the ranked strategies are small, so the ranking of several of the agents was decided by a small margin. Finally, the first places in the individual utility and social welfare categories were awarded to *AgentM*; The second place in the individual category was awarded to *DoNA*; The second place in the social welfare was awarded to *Gangster*.

In more detail, we can discuss the relationships between the social welfare and other measures. As figure and showing, the percentage of agreements and the pareto distance are important features of obtaining the high social welfare. Especially, the correlation coefficient of the percentage of agreements is about 1.0 and, the average of pareto distance is about -1.0 . In other words, the effective strategy of obtaining the social welfare is that finding the pareto frontiers with the high percentage of agreements.

5 In Depth Evaluation of ANAC 2014 Agents

After the competition, a detailed analysis of best performing five ANAC 2014 agents, namely *Agent M*, *DoNA*, *Gangster*, *WhaleAgent*, and *E2Agent*, has been carried out. We first explain our experimental setup, then present the results, after which we analyze the negotiation results with respect to several characteristics of the negotiation scenarios.

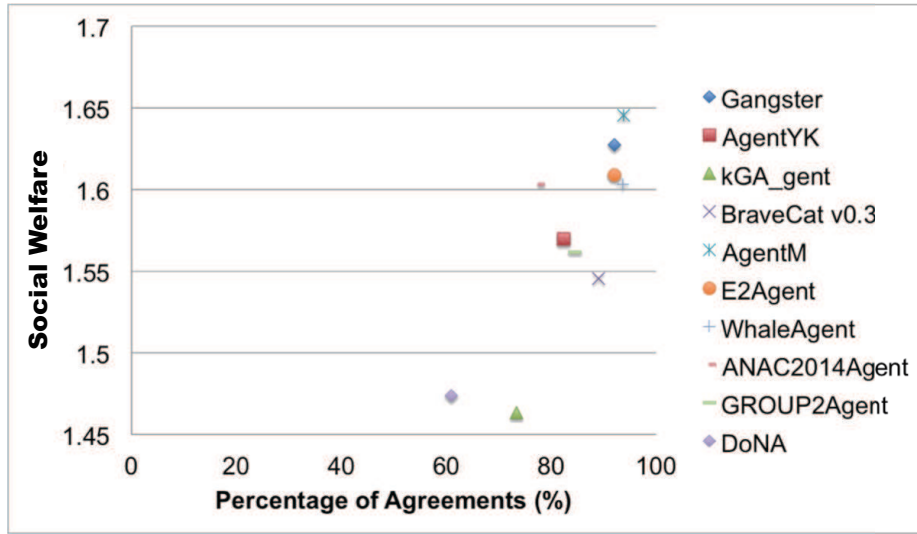


Fig. 7. Plotting graph between the percentage of agreements and the social welfare (A correlation coefficient = 0.8735)

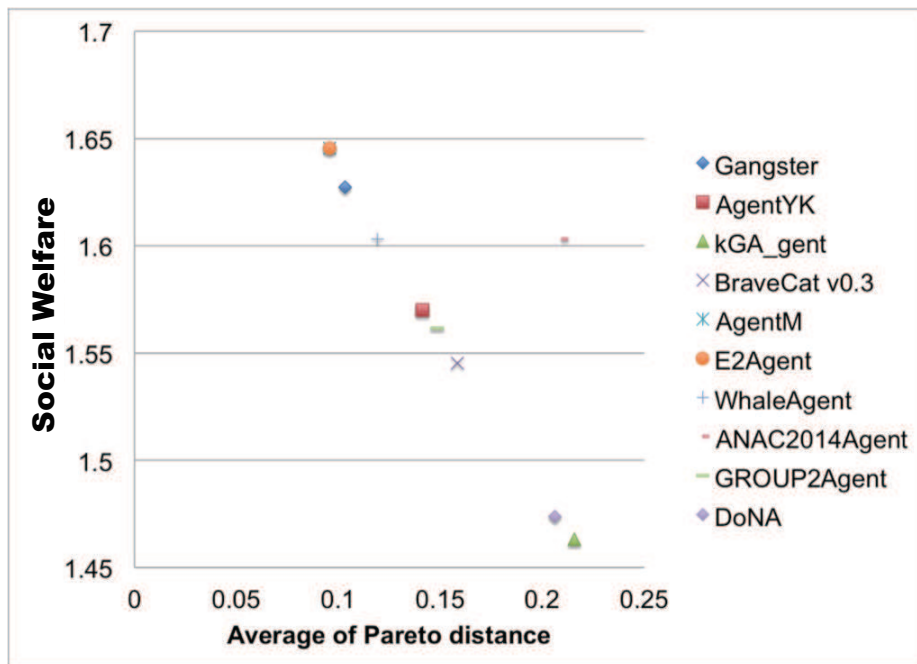


Fig. 8. Plotting graph between the average of pareto distance and the social welfare (A correlation coefficient = -0.9994)




Fig. 9. Cubic, bell and linear constraints

5.1 Experimental Setup

For the in depth evaluation of ANAC 2014 agents, a variety of negotiation scenarios was generated. For the generation of these negotiation scenarios, we consider the following factors:

- The number of issues (n).
- The number of constraints (m).
- The constraint-issue distribution (π).

Each distribution π assigns a predefined number of issues $[1, n]$ to each constraint. For example, $\pi(c_k)$ denotes the number of issues involved in constraint c_k , where $k \in [1, m]$. Our scenario generator supports three types of distributions:

1. Uniform distribution: $\pi(c_k) = p, \forall k, p \in [1, n]$. All constraints have the same cardinality.
2. Power-law distribution: $\pi(c_k)$ is scale-free in terms of $k, k \in [1, m]$. In other words, few constraints are dense (involve most of the issues) while most of the other constraints have only fewer issues.
3. Random distribution: $\pi(c_k)$ is randomly chosen from $[1, n]$, for all k .

In our experiments, an agent i 's preference profile is therefore generated randomly based on a tuple (n, m, π_i) . The idea behind the constraint-issue distribution is to have a generative model allowing us to specify the complexity of each profile. For instance, two negotiating agents 1 and 2 will respectively have two utility spaces (n, m_1, π_1) and (n, m_2, π_2) . The complexities of π_1 and π_2 are reflected through their entropies $H(\pi_1)$ and $H(\pi_2)$ [24].

When generating the scenarios, we attempted to diversify the agent profiles along the following lines. First, we assume that an agent's utility space is built using three types of basic parametric constraints. As shown in figure 9, we distinguish linear (hyperplane), bell-shaped and conic-shaped utility functions [24]. Secondly, it is possible to adjust the complexity of any given utility space by specifying the constraint-issue connectivity using the π distribution. Such connectivity affects the computational complexity required to optimise over any given utility space [23]. Utility spaces with randomised sets in particular, render the search for optimal bids more difficult.

It is important to note that the used scenarios do not account for the level of competitiveness or collaboration between two agents. In particular, we did not differentiate between the cases where the constraints are disjoint sets, and thus giving rise to a zero-sum game, and the cases where the constraints are nested, which yields a collaborative behaviour. Instead, we used generic randomised sets, which reinforces complexity and is useful for any temporal (not utilitarian) assessment of the agents performances.

For our experiments, 108 different negotiation scenarios were generated by taking the number of issues n to be in $\{2, 10, 50\}$, the number of constraints m in $\{1, 10, 50, 100\}$, and the constraint-issue distributions are drawn from $\{Random, Uniform, Power - law\}$. Note that 9 different combinations exist for constraint-issue distribution in bilateral negotiations (e.g. Random vs. Random, Random vs. Uniform and so forth). As a result, 3 different number of issues * 4 different number of constraints * 9 different combinations of constraint-issue distributions in bilateral negotiation makes 108 different negotiation scenarios. The deadline for each negotiation was set to 3 minutes, as it was in the competition. We picked the five best agents from ANAC 2014, each agent negotiates with the other agents and plays in both sides ($5 \times 4 \times 2 = 40$ negotiations). Each negotiation is repeated 5 times and we have 108 negotiation scenarios, which makes 21600 negotiation runs in total.

5.2 Experiment Results

We evaluated the performance of the agents according to the following criteria:

- The average utility gained by the agent (i.e., individual utility).
- The average distance to the closest Pareto outcome.
- The average distance to Nash outcome.
- The average utilitarian social welfare (i.e., sum of all agents' utilities).

Table 5 shows the performance of each agent with respect to each aforementioned criterion over 21600 negotiations. According to these results, *Gangster* outperforms other agents in terms of the individual utility gained by the agents on average. The performance of *Agent M* is close to *Gangster* (0.511 vs 0.543). As far as optimality of a negotiation outcome is concerned, it is desired to reach agreements on the Pareto Frontier and/or to reach a Nash solution. Therefore, it is desired to find agreements whose distance to Pareto outcome and/or Nash outcome is as small as possible. Table 5 shows that *Agent M* performs best according to these criteria followed by *WhaleAgent* and *Gangster*. With respect to the average utilitarian social welfare, the performance of *Agent M* and *Ganster* is almost the same and they outperform other agents.

It is worth mentioning that *DoNA*, that was ranked in second place in ANAC 2014 (see Section 4), is outperformed by all agents in this experimental setup. This inconsistency can be explained by the fact that the performance of an agent highly depends on its opponents. In the final round of ANAC 2014, *DoNA* negotiates against nine different agents while it negotiates against the best four agents among those agents in this analysis. Furthermore, the negotiation scenarios used in this section are different than the scenarios used in the final round of ANAC 2014. This clearly influences the performance of the agents.

Performance Criterion	E2Agent	DoNA	AgentM	Gangster	WhaleAgent
Average Individual Utility:	0.478	0.353	0.511	0.543	0.490
Average Pareto Distance:	0.180	0.349	0.112	0.117	0.117
Average Nash Distance:	0.252	0.434	0.163	0.180	0.182
Average Utilitarian Social Welfare:	0.936	0.652	1.056	1.057	1.046

Table 5. Overall Performance of Best Five ANAC 2014 Agents over 21600 Negotiations

In the following sections, we investigate the effect of domain size, constraint size and constraint-issue distribution on the performance of these agents.

5.3 Effect of Domain Size

Table 6 shows the negotiation results by grouping the negotiations according to their domain size (number of issues). Considering the average individual utility gained by the agents, it can be seen that *Gangster* agent is the best performer. According to this performance criterion, the performance of *WhaleAgent* and *Agent M* is similar to each other and they perform better than *E2Agent* and *DoNA*. It can be seen that the performance of *DoNA* drops drastically when the outcome space becomes increases to 10^{50} , based on 50 issues with 10 possible values. In general, the performance of all agents reduces when we increase the domain size to 10^{50} . In the table, domain size is indicated by the number of issues, as the number of possible values per issue is always 10. So domain size indication 2 actually means 10^2 , and indication 10 stands for 10^{10} .

In terms of utilitarian social welfare, *Gangster* is the winner in the negotiations where the domain size is 2 or 10; however, when the domain size is 50, *WhaleAgent* gains higher social welfare on average. According to the results showing the average distance to Nash outcome, we can say that *Agent M* and *WhaleAgent* perform better than *Gangster* agent. That is, *Gangster* agent achieves higher utility for himself while *Agent M* and *WhaleAgent* consider other agents' utility and act more collaboratively in their negotiation so that that negotiation outcomes seems to be more fair. As far as the results regarding to the average Pareto distance are considered, the performance of *Gangster* agent drops when the outcome space becomes huge. *WhaleAgent* performs better than *Gangster* agent when the domain size is 50.

5.4 Effect of Constraint Size

In this section, we study whether the number of constraints in a given negotiation scenario has any effect on the performance of the agents. Table 7 shows the negotiation results by grouping them according to the number of constraints in the underlying negotiation scenarios. According to the average individual utility gained by the agents, the winner is *Gangster* agent, irrespective of the constraint size.

When the number of constraints is extremely low or high, using *Gangster* agent seems a promising approach in terms of the average distance Pareto outcome. However, when the number of constraints is 10 or 50, then *Agent M* and *WhaleAgent* perform better. This is also valid for utilitarian social welfare. If we investigate the performance

Performance Criterion	Domain Size	E2Agent	DoNA	AgentM	Gangster	WhaleAgent
Average Individual Utility:	2	0.492	0.366	0.557	0.572	0.492
	10	0.503	0.431	0.512	0.582	0.509
	50	0.439	0.261	0.463	0.475	0.469
Average Pareto Distance:	2	0.214	0.384	0.091	0.084	0.117
	10	0.127	0.217	0.087	0.088	0.088
	50	0.198	0.445	0.158	0.177	0.145
Average Nash Distance:	2	0.329	0.516	0.171	0.188	0.228
	10	0.210	0.322	0.147	0.163	0.158
	50	0.217	0.465	0.172	0.190	0.159
Average Utilitarian Social Welfare:	2	0.938	0.652	1.136	1.142	1.088
	10	0.993	0.803	1.093	1.104	1.083
	50	0.879	0.501	0.941	0.926	0.968

Table 6. Overall Performance of Best Five ANAC 2014 Agents wrt Varying Domain Size

Performance Criterion	Constraint Size	E2Agent	DoNA	AgentM	Gangster	WhaleAgent
Average Individual Utility:	1	0.474	0.188	0.466	0.537	0.459
	10	0.496	0.397	0.553	0.556	0.521
	50	0.480	0.417	0.531	0.558	0.510
	100	0.462	0.409	0.493	0.522	0.469
Average Pareto Distance:	1	0.199	0.513	0.144	0.120	0.151
	10	0.196	0.370	0.125	0.166	0.124
	50	0.173	0.275	0.089	0.092	0.088
	100	0.150	0.237	0.091	0.089	0.103
Average Nash Distance:	1	0.307	0.673	0.213	0.210	0.244
	10	0.285	0.454	0.189	0.240	0.198
	50	0.225	0.331	0.132	0.140	0.141
	100	0.191	0.278	0.120	0.130	0.143
Average Utilitarian Social Welfare:	1	0.875	0.330	1.006	1.052	0.986
	10	1.001	0.728	1.125	1.072	1.121
	50	0.955	0.779	1.088	1.087	1.083
	100	0.915	0.771	1.007	1.018	0.995

Table 7. Overall Performance of Best Five ANAC 2014 Agents w.r.t. Varying Constraint Size

of the agents with respect to the average distance to the Nash solution, *Agent M* is the best choice except for the extreme case in which a negotiation scenario consists of only a single constraint.

5.5 Effect of Constraint-Issue Distribution

In this section, we study the effect of the complexity of negotiation scenarios on the performance of the agents. Complexity is defined in terms of the entropy of each constraint-issue distribution. The higher the entropy, the more complex the distribution is. The complexity of the negotiation scenarios are as follows: $H(\pi_{uniform}) > H(\pi_{random}) > H(\pi_{pl})$.

This presupposes that entropy correlates with the bumpiness of the utility space and the computational time needed to explore it [23]. Table 8 shows the negotiation results grouped by their constraint-issue distribution. Based on these results, we can say that *Agent M* is good at complex domains. According to all evaluation criteria, *Agent M* is the winner when the scenarios are generated with normal distribution. For relatively less complex negotiation scenarios generated by random distribution, *Gangster* agent is the best choice.

We note that having more constraints implies more marginal utility ($\sum_{k=1}^m w(c_k)$) but not necessary better agreement zones. In this sense, one constraint provides low utility when it is isolated in the overall utility space. Moreover, having many large constraints give better utility for the agents since such large constraints give more acceptable agreements.

Performance Criterion	Constraint-Issue Distribution	E2Agent	DoNA	AgentM	Gangster	WhaleAgent
Average Individual Utility:	Uniform	0.509	0.264	0.570	0.565	0.531
	Power-Law	0.602	0.488	0.654	0.727	0.633
	Random	0.372	0.330	0.403	0.436	0.376
Average Pareto Distance:	Uniform	0.229	0.561	0.137	0.159	0.181
	Power-Law	0.205	0.336	0.123	0.139	0.137
	Random	0.123	0.190	0.082	0.068	0.080
Average Nash Distance:	Uniform	0.260	0.606	0.159	0.186	0.208
	Power-Law	0.255	0.407	0.164	0.185	0.189
	Random	0.187	0.255	0.121	0.115	0.131
Average Utilitarian Social Welfare:	Uniform	1.014	0.510	1.152	1.119	1.083
	Power-Law	1.188	0.927	1.368	1.379	1.345
	Random	0.738	0.614	0.817	0.845	0.817

Table 8. Overall Performance of Best Five ANAC 2014 Agents wrt Constraint-issue Distribution

DoNA agent gets the lowest average individual utility and utilitarian social welfare. Similarly, its performance is dropping as the complexity of the negotiation domain in terms of distributions type is increasing from power-law towards uniform. *DoNA* gets better utility on average for power-law, but less for random and lesser for uniform.

This decreasing performance is due to the increasing complexity and the usage of a behavioral strategy that only focuses on satisfiability (instead of optimality) with respect to the discount factor and the reservation value. The increasing complexity reduces the marginal utility and renders any satisfactory parameter value as performing well enough for the agent.

Agents adopt different metaheuristics when sampling their utility spaces (Simulated Annealing, Genetic Algorithms, Heuristic Search, etc.). To understand the effect of domain complexity on the agents, we should understand the nature of the used metaheuristics and how they perform given different distributions (π). Depending on the domain, such metaheuristics can be for instance randomized or not, global or local, use the opponent bids or not, etc. When comparing the agents' performances, it is important to classify the agents based on their metaheuristics. For instance, metaheuristics that rely on randomization (Simulated Annealing, Genetic Algorithms) work well when there is no prior knowledge of the constraints shapes and connectivity.

The complexity of the utility spaces affects the exploration time. For example, if the agent is sampling a power-law distribution, he can find the optimal bids quicker than if the utility space was distributed according to a uniform distribution. An agent with a power-law distribution and a sampling metaheuristic that works well with power-law topologies will quickly know what his best bids are. He can consequently focus more on his acceptance/concession strategies. An agent with a complete distribution will have to search the utility space and propose at the same time.

6 Conclusion

In this paper, we describe the fifth international automated negotiating agents competition. The main challenge in ANAC 2014 is to design intelligent negotiating agents that can explore a large-scale outcome space effectively under time constraints. In this competition, we used negotiation scenarios using nonlinear utility functions and agents were not allowed to access the structure of their utility function. Therefore, they need to use advance search techniques such as Simulated Annealer, Hill-Climber, Genetic Algorithms and so on.

Following the competition, we conducted additional experiments in order to analyze the performance of the best performing five agents elaborately. To achieve this, we generated 108 different nonlinear negotiation scenarios systematically and ran 21600 negotiations to evaluate the agents with respect to a number of performance criteria. We consider the average individual utility gained by the agents, the utilitarian social welfare, the average distance to Pareto and Nash solution. Our results showed that *Gangster* agent outperforms other agents regarding to the individual utility and social welfare while *Agent M* is better off regarding to other metrics. Furthermore, we studied the effect of domain size, constraint size and issue-constraint distribution on the performance of the agents. An interesting result is that *Agent M* gains higher individual utility than *Gangster* agent in complex negotiation scenarios generated with uniform distribution while *Gangster* agent is better off for other type of scenarios.

In conclusion, we believe that ANAC facilitates the research in automated negotiation. The joint discussions with the participating teams gave us great insights for the

future competitions. There is a great interests in multilateral negotiation and human-agent negotiation. As a future work, we are going to organize a multiparty negotiation competition in which the negotiating agents have more than one opponents. In addition, it is also interesting to investigate human-agent negotiation although its experimental set up might be costly(i.e. we need human negotiators to test the performance of the agents).

Acknowledgments. This work was supported by the ITEA M2MGrids Project, grant number ITEA141011 and partially by the FP7 project IQmulus (FP7-ICT-2011-318787), "A High-volume Fusion and Analysis Platform for Geospatial Point Clouds, Coverages and Volumetric Data sets. Many thanks to Dr. Ivan Marsa-Maestre for his support on scenario generation.

References

1. B. An and V. R. Lesser. Yushu: a heuristic-based agent for automated negotiating competition. In T. Ito, M. Zhang, V. Robu, S. Fatima, and T. Matsuo, editors, *New Trends in Agent-based Complex Automated Negotiations*, volume 383 of *Studies in Computational Intelligence*, pages 145–149. Springer Berlin Heidelberg, 2012.
2. R. Aydođan, T. Baarslag, K. V. Hindriks, C. M. Jonker, and P. Yolum. Heuristics for using cp-nets in utility-based negotiation without knowing utilities. *Knowledge and Information Systems*, 45(2):357–388, 2014.
3. R. Aydogan and P. Yolum. Learning opponents preferences for effective negotiation: An approach based on concept learning. *Journal of Autonomous Agents and Multi-Agent Systems*, 24(1):104–140, 2012.
4. T. Baarslag. *What to Bid and When to Stop*. Dissertation, Delft University of Technology, Sep 2014.
5. T. Baarslag. *Exploring the Strategy Space of Negotiating Agents: A Framework for Bidding, Learning and Accepting in Automated Negotiation*. Springer Theses: Recognizing Outstanding Ph.D. Research. Springer International Publishing, 2016.
6. T. Baarslag, R. Aydođan, K. V. Hindriks, K. Fujita, T. Ito, and C. M. Jonker. The automated negotiating agents competition, 2010-2015. *AI Magazine*, 36(4):115–118, 12/2015 2015.
7. T. Baarslag, K. Fujita, E. H. Gerding, K. V. Hindriks, T. Ito, N. R. Jennings, C. M. Jonker, S. Kraus, R. Lin, V. Robu, and C. R. Williams. Evaluating practical negotiating agents: Results and analysis of the 2011 international competition. *Artificial Intelligence*, 198:73 – 103, May 2013.
8. T. Baarslag, K. V. Hindriks, and C. M. Jonker. A tit for tat negotiation strategy for real-time bilateral negotiations. In T. Ito, M. Zhang, V. Robu, and T. Matsuo, editors, *Complex Automated Negotiations: Theories, Models, and Software Competitions*, volume 435 of *Studies in Computational Intelligence*, pages 229–233. Springer Berlin Heidelberg, 2013.
9. T. Baarslag, K. V. Hindriks, C. M. Jonker, S. Kraus, and R. Lin. The first automated negotiating agents competition (ANAC 2010). In T. Ito, M. Zhang, V. Robu, S. Fatima, and T. Matsuo, editors, *New Trends in Agent-based Complex Automated Negotiations*, volume 383 of *Studies in Computational Intelligence*, pages 113–135, Berlin, Heidelberg, 2012. Springer-Verlag.
10. M. Ben Adar, N. Sofy, and A. Elimelech. Gahboninho: Strategy for balancing pressure and compromise in automated negotiation. In T. Ito, M. Zhang, V. Robu, and T. Matsuo, editors,

Complex Automated Negotiations: Theories, Models, and Software Competitions, volume 435 of *Studies in Computational Intelligence*, pages 205–208. Springer Berlin Heidelberg, 2013.

11. S. Chen, H. B. Ammar, K. Tuyls, and G. Weiss. Optimizing complex automated negotiation using sparse pseudo-input gaussian processes. In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent Systems, AAMAS '13*, pages 707–714, Richland, SC, 2013. International Foundation for Autonomous Agents and Multiagent Systems.
12. S. Chen and G. Weiss. OMAC: a discrete wavelet transformation based negotiation agent. In I. Marsa-Maestre, M. Lopez-Carmona, T. Ito, M. Zhang, Q. Bai, and K. Fujita, editors, *Novel Insights in Agent-based Complex Automated Negotiation*, volume 535 of *Studies in Computational Intelligence*, pages 187–196. Springer, Japan, 2014.
13. D. de Jonge and C. Sierra. Gangster: An automated negotiator applying genetic algorithms. In N. Fukuta, T. Ito, M. Zhang, K. Fujita, and V. Robu, editors, *Recent Advances in Agent-based Complex Automated Negotiation*, pages 225–234. Springer International Publishing, Cham, 2016.
14. E. S. Erez and I. Zuckerman. Dona - a domain-based negotiation agent. In N. Fukuta, T. Ito, M. Zhang, K. Fujita, and V. Robu, editors, *Recent Advances in Agent-based Complex Automated Negotiation*, pages 261–271. Springer International Publishing, Cham, 2016.
15. S. Fatima, S. Kraus, and M. Wooldridge. *Principles of Automated Negotiation*. Cambridge University Press, 2014.
16. R. Fishel, M. Bercovitch, and Y. Gal. Bram agent. In T. Ito, M. Zhang, V. Robu, and T. Matsuo, editors, *Complex Automated Negotiations: Theories, Models, and Software Competitions*, volume 435 of *Studies in Computational Intelligence*, pages 213–216. Springer Berlin Heidelberg, 2013.
17. K. Fujita. Automated negotiating agent with strategy adaptation for multi-times negotiations. In *IEEE 6th International Conference on Service-Oriented Computing and Applications (SOCA)*, pages 333–337, 2013.
18. K. Fujita, R. Aydođan, T. Baarslag, T. Ito, and C. Jonker. The fifth automated negotiating agents competition (anac 2014). In N. Fukuta, T. Ito, M. Zhang, K. Fujita, and V. Robu, editors, *Recent Advances in Agent-based Complex Automated Negotiation*, pages 211–224. Springer International Publishing, Cham, 2016.
19. K. Fujita, T. Ito, T. Baarslag, K. V. Hindriks, C. M. Jonker, S. Kraus, and R. Lin. The second automated negotiating agents competition (ANAC 2011). In T. Ito, M. Zhang, V. Robu, and T. Matsuo, editors, *Complex Automated Negotiations: Theories, Models, and Software Competitions*, volume 435 of *Studies in Computational Intelligence*, pages 183–197. Springer Berlin Heidelberg, 2013.
20. K. Fujita, T. Ito, and M. Klein. A secure and fair protocol that addresses weaknesses of the nash bargaining solution in nonlinear negotiation. *Group Decision and Negotiation*, 21(1):29–47, 2012.
21. K. Fujita, T. Ito, and M. Klein. An approach to scalable multi-issue negotiation: Decomposing the contract space. *Computational Intelligence*, 30(1):30–47, 2014.
22. K. Fujita, T. Ito, and M. Klein. Efficient issue-grouping approach for multiple interdependent issues negotiation between exaggerator agents. *Decision Support Systems*, 60:10–17, 2014.
23. R. Hadfi and T. Ito. Cognition as a game of complexity. In *Proceedings of 12th International Conference on Cognitive Modeling (ICCM)*, 2013.
24. R. Hadfi and T. Ito. On the complexity of utility hypergraphs. In *The Seventh International Workshop on Agent-based Complex Automated Negotiations (ACAN2014)*, 2014.
25. R. Hadfi and T. Ito. Low-complexity exploration in utility hypergraphs. *Journal of information processing*, 23(2):176–184, 2015.

26. J. Hao and H. Leung. ABiNeS: An adaptive bilateral negotiating strategy over multiple items. In *Proceedings of the The 2012 IEEE/WIC/ACM International Joint Conferences on Web Intelligence and Intelligent Agent Technology*, volume 2 of *WI-IAT '12*, pages 95–102, Washington, DC, USA, Dec 2012. IEEE Computer Society.
27. L. Ilany and Y. Gal. Algorithm selection in bilateral negotiation. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence (AAAI 2013)*, 2013.
28. T. Ito, H. Hattori, and M. Klein. Multi-issue negotiation protocol for agents: Exploring non-linear utility spaces. In *The twentieth International Joint Conference on Artificial Intelligence (IJCAI-07)*, volume 7, pages 1347–1352, 2007.
29. T. Ito, M. Klein, and H. Hattori. A multi-issue negotiation protocol among agents with nonlinear utility functions. *Multiagent Grid Syst.*, 4(1):67–83, Jan. 2008.
30. 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(2):199–215, 2001.
31. Y. Kadono and N. Fukuta. Agent yk: An efficient estimation of opponents intention with stepped limited concessions. In N. Fukuta, T. Ito, M. Zhang, K. Fujita, and V. Robu, editors, *Recent Advances in Agent-based Complex Automated Negotiation*, pages 279–283. Springer International Publishing, Cham, 2016.
32. S. Kakimoto and K. Fujita. k-gagent: Negotiating agents considering interdependencies between issues. In N. Fukuta, T. Ito, M. Zhang, K. Fujita, and V. Robu, editors, *Recent Advances in Agent-based Complex Automated Negotiation*, pages 241–247. Springer International Publishing, Cham, 2016.
33. S. Kawaguchi, K. Fujita, and T. Ito. Compromising strategy based on estimated maximum utility for automated negotiation agents competition (anac-10). *Modern Approaches in Applied Intelligence, Lecture Notes in Computer Science*, 6704:501–510, 2011.
34. S. Kawaguchi, K. Fujita, and T. Ito. Compromising strategy based on estimated maximum utility for automated negotiating agents. In T. Ito, M. Zhang, V. Robu, S. Fatima, and T. Matsuo, editors, *New Trends in Agent-based Complex Automated Negotiations, Series of Studies in Computational Intelligence*, pages 137–144, Berlin, Heidelberg, 2012. Springer-Verlag.
35. M. Klein, P. Faratin, H. Sayama, and Y. Bar-Yam. Negotiating complex contracts. *Group Decision and Negotiation*, 12(2):58–73, 2003.
36. R. Lin, S. Kraus, T. Baarslag, D. Tykhonov, K. V. Hindriks, and C. M. Jonker. Genius: An integrated environment for supporting the design of generic automated negotiators. *Computational Intelligence*, 30(1):48–70, 2014.
37. M. A. López-Carmona, I. Marsá-Maestre, M. Klein, and T. Ito. Addressing stability issues in mediated complex contract negotiations for constraint-based, non-monotonic utility spaces. *Autonomous Agents and Multi-Agent Systems*, 24(3):485–535, 2012.
38. I. Marsa-Maestre, M. Klein, C. M. Jonker, and R. Aydoğan. From problems to protocols: Towards a negotiation handbook. *Decision Support Systems*, pages 39–54, 2014.
39. I. Marsá-Maestre, M. A. López-Carmona, M. Klein, T. Ito, and K. Fujita. Addressing utility space complexity in negotiations involving highly uncorrelated, constraint-based utility spaces. *Computational Intelligence*, 30(1):1–29, 2014.
40. S. Morii and T. Ito. Agent’s strategy in multiple-issue negotiation competition and analysis of result. In G. Boella, E. Elkind, B. T. R. Savarimuthu, F. Dignum, and M. K. Purvis, editors, *PRIMA 2013: Principles and Practice of Multi-Agent Systems*, volume 8291 of *Lecture Notes in Computer Science*, pages 486–493. Springer Berlin Heidelberg, 2013.
41. M. Niimi and T. Ito. Agentm. In N. Fukuta, T. Ito, M. Zhang, K. Fujita, and V. Robu, editors, *Recent Advances in Agent-based Complex Automated Negotiation*, pages 235–240. Springer International Publishing, Cham, 2016.

42. V. Sanchez-Anguix, R. Aydođan, V. Julian, and C. Jonker. Unanimously acceptable agreements for negotiation teams in unpredictable domains. *Electronic Commerce Research and Applications*, 13(4):243 – 265, 2014.
43. M. Sato and T. Ito. Whaleagent: Hardheaded strategy and conceder strategy based on the heuristics. In N. Fukuta, T. Ito, M. Zhang, K. Fujita, and V. Robu, editors, *Recent Advances in Agent-based Complex Automated Negotiation*, pages 273–278. Springer International Publishing, Cham, 2016.
44. B. Szollosi-Nagy, D. Festen, and M. M. Skarzynska. A greedy coordinate descent algorithm for high-dimensional nonlinear negotiation. In N. Fukuta, T. Ito, M. Zhang, K. Fujita, and V. Robu, editors, *Recent Advances in Agent-based Complex Automated Negotiation*, pages 249–260. Springer International Publishing, Cham, 2016.
45. W. Visser, R. Aydogan, K. V. Hindriks, and C. M. Jonker. A framework for qualitative multi-criteria preferences. In *The 4th International Conference on Agents and Artificial Intelligence (ICAART 2012)*, pages 243–248. SciTePress, 2012.
46. M. P. Wellman, P. R. Wurman, K. O’Malley, R. Banger, S. de Lin, D. Reeves, and W. E. Walsh. Designing the market game for a trading agent competition. *IEEE Internet Computing*, 5(2):43–51, 2001.
47. C. R. Williams, V. Robu, E. H. Gerding, and N. R. Jennings. Iamhaggler2011: A gaussian process regression based negotiation agent. In T. Ito, M. Zhang, V. Robu, and T. Matsuo, editors, *Complex Automated Negotiations: Theories, Models, and Software Competitions*, volume 435 of *Studies in Computational Intelligence*, pages 209–212. Springer Berlin Heidelberg, 2013.
48. C. R. Williams, V. Robu, E. H. Gerding, and N. R. Jennings. An overview of the results and insights from the third automated negotiating agents competition (ANAC 2012). In I. Marsa-Maestre, M. A. Lopez-Carmona, T. Ito, M. Zhang, Q. Bai, and K. Fujita, editors, *Novel Insights in Agent-based Complex Automated Negotiation*, volume 535 of *Studies in Computational Intelligence*, pages 151–162. Springer, Japan, 2014.
49. T. T. Wu and K. Lange. Coordinate descent algorithms for lasso penalized regression. *The Annals of Applied Statistics*, 2(1):224–244, 2008.
50. K. (Ya’akov) Gal and L. Ilany. *Next Frontier in Agent-based Complex Automated Negotiation*, chapter The Fourth Automated Negotiation Competition, pages 129–136. Springer Japan, Tokyo, 2015.
51. F. Zafari and F. Nassiri-Mofakham. Bravecat: Iterative deepening distance-based opponent modeling and hybrid bidding in nonlinear ultra large bilateral multi issue negotiation domains. In N. Fukuta, T. Ito, M. Zhang, K. Fujita, and V. Robu, editors, *Recent Advances in Agent-based Complex Automated Negotiation*, pages 285–293. Springer International Publishing, Cham, 2016.
52. E. Zitzler, M. Laumanns, L. Thiele, E. Zitzler, E. Zitzler, L. Thiele, and L. Thiele. Spea2: Improving the strength pareto evolutionary algorithm, 2001.