The Fifth Automated Negotiating Agents Competition (ANAC 2014)

Katsuhide Fujita, Reyhan Aydoğan, Tim Baarslag, Takayuki Ito and Catholijn Jonker

Abstract In May 2014, we organized the Fifth International Automated Negotiating Agents Competition (ANAC 2014) in conjunction with AAMAS 2014. ANAC is an international competition that challenges researchers to develop a successful automated negotiator for scenarios where there is incomplete information about the opponent. One of the goals of this competition is to help steer the research in the area of bilateral multi-issue negotiations, and to encourage the design of generic negotiating agents that are able to operate in a variety of scenarios. 21 teams from 13 different institutes competed in ANAC 2014. This chapter describes the participating agents and the setup of the tournament, including the different negotiation scenarios that were used in the competition. We report on the results of the qualifying and final round of the tournament.

K. Fujita (✉)
Faculty of Engineering, Tokyo University of Agriculture and Technology,
Tokyo, Japan
e-mail: katfuji@cc.tuat.ac.jp

R. Aydoğan
Computer Science Department, Özyeğin University, Istanbul, Turkey
e-mail: R.Aydogan@tudelft.nl

T. Baarslag
Agents, Interaction and Complexity Group at the University of Southampton,
Southampton, UK
e-mail: tb1m13@ecs.soton.ac.uk

T. Ito
Techno-Business Administration (MTBA), Nagoya Institute of Technolog,
Aichi, Japan
e-mail: ito.takayuki@nitech.ac.jp

C. Jonker
Man Machine Interaction Group, Delft University of Technology, Delft, Netherlands
e-mail: C.M.Jonker@tudelft.nl

N. Fukuta et al. (eds.), *Recent Advances in Agent-based Complex Automated Negotiation*, Studies in Computational Intelligence 638,
DOI 10.1007/978-3-319-30307-9_13

1 Introduction

Success in developing an automated agent with negotiation capabilities has great advantages and implications. In order to help focus research on proficiently negotiating automated agents, we have organized the first automated negotiating agents competition (ANAC).

The results of the different implementations are difficult to compare, as various setups are used for experiments in ad hoc negotiation environments [7]. An additional goal of ANAC is to build a community in which work on negotiating agents can be compared by standardized negotiation benchmarks to evaluate the performance of both new and existing agents. Recently, the analysis of ANAC becomes important fields of automated negotiations in multi-agent systems [1].

In designing proficient negotiating agents, standard game-theoretic approaches cannot be directly applied. Game theory models assume complete information settings and perfect rationality [9, 10]. However, human behavior is diverse and cannot be captured by a monolithic model. Humans tend to make mistakes, and they are affected by cognitive, social and cultural factors [8]. A means of overcoming these limitations is to use heuristic approaches to design negotiating agents. When negotiating agents are designed using a heuristic method, we need an extensive evaluation, typically through simulations and empirical analysis.

We employ an environment that allows us to evaluate agents in a negotiation competition: GENIUS [7], a **G**eneral **E**nvironment for **N**egotiation with **I**ntelligent multi-purpose **U**sage **S**imulation. GENIUS helps facilitating the *design* and *evaluation* of automated negotiators' strategies. It allows easy development and integration of existing negotiating agents, and can be used to simulate individual negotiation sessions, as well as tournaments between negotiating agents in various negotiation scenarios. The design of general automated agents that can negotiate proficiently is a challenging task, as the designer must consider different possible environments and constraints. GENIUS can assist in this task, by allowing the specification of different negotiation domains and preference profiles by means of a graphical user interface. It can be used to train human negotiators by means of negotiations against automated agents or other people. Furthermore, it can be used to teach the design of generic automated negotiating agents.

The First Automated Negotiating Agents Competition (ANAC 2010) was held in May 2010, with the finals being run during the AAMAS 2010 conference. Seven teams had participated and three domains were used. *AgentK* generated by the Nagoya Institute of Technology team won the ANAC 2010 [2]. The Second Automated Negotiating Agents Competition (ANAC 2011) was held in May 2011, with the AAMAS 2011 conference. 18 teams had participated and eight domains were used. The new feature of ANAC 2011 was the discount factor. *HardHeaded* generated by the Delft University of Technology won the ANAC 2011 [3]. The Third Automated Negotiating Agents Competition (ANAC 2012) was held in May 2012, with the AAMAS 2012 conference. 17 teams had participated and 24 domains were used. The new feature of ANAC 2012 was the reservation value. *CUHKAgent* generated by the Chinese University of Hong Kong won the ANAC 2012 [12]. The Forth Automated Negotiating Agents Competition (ANAC 2013) was held in May 2013, with the AAMAS 2013 conference. 19 teams had participated and 24 domains were used. The new feature of ANAC 2013 was that agents can use the bidding history. *The Fawkes* generated by the Delft University of Technology won the ANAC 2013 [4].

ANAC organizers have been employing some of the new feature every year to develop the ANAC competition and the automated negotiations communities. One of the key point in achieving automated negotiation in real life is 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 is effective to the linear function effectively, it is not always possible or desirable in the nonlinear situations [5]. In ANAC 2014, we used the constraint-based nonlinear utility function with integer issues. In addition, the domains deal with large-size domains, with outcome spaces as big as 1050 outcomes.

The remainder of this paper is organized as follows. Section 2 provides an overview over the design choices for ANAC, including the model of negotiation, tournament platform and evaluation criteria. In Sect. 3, we present the setup of ANAC 2014 followed by Sect. 4 that layouts the results of competition. Finally, Sect. 5 outlines our conclusions and our plans for future competitions.

2 Set up of ANAC

2.1 Negotiation Model

Given the goals outlined in the introduction, in this section we introduce the set-up and negotiation protocol used in ANAC. In this competition, we consider *bilateral* negotiations, i.e. negotiation between two parties. The interaction between negotiating parties is regulated by a *negotiation protocol* that defines the rules of how and when proposals can be exchanged. In the competition, we use the alternating-offers protocol for bilateral negotiation as proposed in [11], in which the negotiating parties exchange offers in turns. The alternating-offers protocol conforms with our criterion to have simplicity of rules. Moreover, it is a protocol which is widely studied and used in literature, both in game-theoretic and heuristic settings of negotiation (a non-exhaustive list includes [6, 9, 10]).

Now, the parties negotiate over a set of *issues*, and every issue has an associated range of alternatives or *values*. A negotiation *outcome* consists of a mapping of every issue to a value, and the set, Ω of all possible outcomes is called the negotiation *domain*. The domain is common knowledge to the negotiating parties and stays fixed during a single negotiation session. In addition to the domain, both parties also have privately-known preferences described by their *preference profiles* over Ω. These preferences are modeled using a utility function U that maps a possible outcomes $\omega \in \Omega$ to a real-valued number in the range [0, 1]. While the domain (i.e. the set

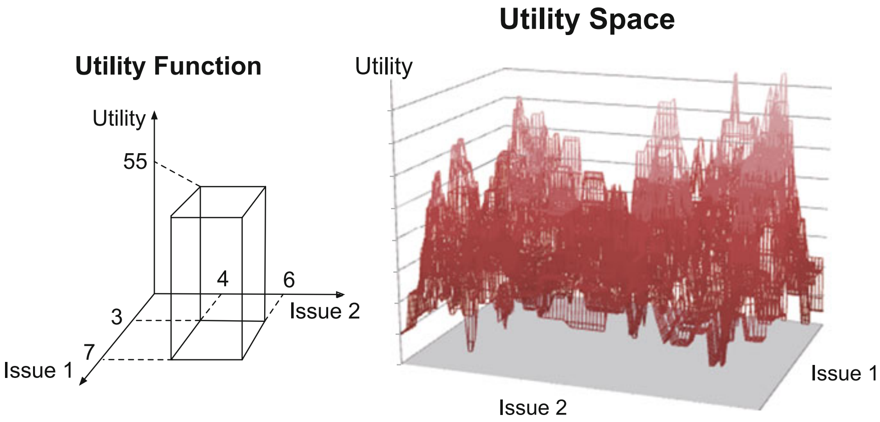

Fig. 1 Example of a nonlinear utility space

of outcomes) is common knowledge, the preference profile of each player is private information. This means that each player has only access to its own utility function, and does not know the preferences of its opponent.[1] Moreover, we use the term *scenario* to refer to the domain and the pair of preference profiles (for each agent) combined.

In ANAC 2014, we focus on *nonlinear* domains settings with a finite set of integer values per issue. An agent's utility function, in the formulation, is described in terms of constraints. There are l constraints, $c_k \in C$. Each constraint represents a region in the contract space with one or more dimensions and an associated utility value. In addition, c_k has value $v_a(c_k, \mathbf{s})$ if and only if it is satisfied by contract $\mathbf{s}$. Every agent has its own, typically unique, set of constraints. An agent's utility for contract $\mathbf{s}$ is defined as the weighted sum of the utility for all the constraints it satisfies, i.e., as $u_a(\mathbf{s}) = \sum_{c_k \in C, \mathbf{s} \in x(c_k)} v_a(c_k, \mathbf{s})$, where $x(c_k)$ is a set of possible contracts (solutions) of c_k. This expression produces a "bumpy" nonlinear utility function with high points where many constraints are satisfied and lower regions where few or no constraints are satisfied. This represents a crucial departure from previous efforts on multi-issue negotiation, where contract utility is calculated as the weighted sum of the utilities for individual issues, producing utility functions shaped like flat hyperplanes with a single optimum.

Figure 1 shows an example of a utility space generated via a collection of binary constraints involving Issues 1 and 2. In addition, the number of terms is two. The example, which has a value of 55, holds if the value for Issue 1 is in the range [3, 7] and the value for Issue 2 is in the range [4, 6]. The utility function is highly nonlinear with many hills and valleys. This constraint-based utility function representation

[1]We note that, in the competition each agent plays *both* preference profiles, and therefore it would be possible in theory to learn the opponent's preferences. However, the rules explicitly disallow learning *between* negotiation sessions, and only *within* a negotiation session. This is done so that agents need to be designed to deal with unknown opponents.

allows us to capture the issue interdependencies common in real-world negotiations. The constraint in Fig. 1, for example, captures the fact that a value of 4 is desirable for issue 1 if issue 2 has the value 4, 5 or 6. Note, however, that this representation is also capable of capturing linear utility functions as a special case (they can be captured as a series of unary constraints). A negotiation protocol for complex contracts can, therefore, handle linear contract negotiations.

Finally, we supplement it with a deadline, reservation value and discount factors. The reasons for doing so are both pragmatic and to make the competition more interesting from a theoretical perspective. In addition, as opposed to having a fixed number of rounds, both the discount factor are measured in *real time*. In particular, it introduces yet another factor of uncertainty since it is now unclear how many negotiation rounds there will be, and how much time an opponent requires to compute a counter offer. In ANAC 2014, the discount factors and reservation value depend on the scenario, but the deadline is set to 3 min. The implementation of discount factors in ANAC 2014 is as follows.

A negotiation lasts a predefined time in seconds *(deadline)*. The time line is normalized, i.e.: time $t \in [0, 1]$, where $t = 0$ represents the start of the negotiation and $t = 1$ represents the deadline. When agents can make agreements in the deadline, the individual utilities of each agent are the *reservation value*. Apart from a deadline, a scenario may also feature *discount factors*. Discount factors decrease the utility of the bids under negotiation as time passes. Let d in [0, 1] be the discount factor. Let t in [0, 1] be the current normalized time, as defined by the timeline. We compute the discounted utility U_D^t of an outcome ω from the undiscounted utility function U as follows:

$$U_D^t(\omega) = U(\omega) \cdot d^t \tag{1}$$

At $t = 1$, the original utility is multiplied by the discount factor. Furthermore, if $d = 1$, the utility is not affected by time, and such a scenario is considered to be undiscounted.

2.2 Running the Tournament

As a tournament platform to run and analyse the negotiations, we use the GENIUS environment (General Environment for Negotiation with Intelligent multi-purpose Usage Simulation) [7]. GENIUS is a research tool for automated multi-issue negotiation, that facilitates the design and evaluation of automated negotiators' strategies. It also provides an easily accessible framework to develop negotiating agents via a public API. This setup makes it straightforward to implement an agent and to focus on the development of strategies that work in a general environment.

GENIUS incorporates several mechanisms that aim to support the design of a general automated negotiator. The first mechanism is an analytical toolbox, which provides a variety of tools to analyse the performance of agents, the outcome of the negotiation and its dynamics. The second mechanism is a repository of domains and utility functions. Lastly, it also comprises repositories of automated negotiators. In addition, GENIUS enables the evaluation of different strategies used by automated agents that were designed using the tool. This is an important contribution as it allows researchers to empirically and *objectively* compare their agents with others in different domains and settings.

The timeline of ANAC 2014 consists of two phases: the qualifying round and the final round. The domains and preference profiles used during the competition are not known in advance and were designed by the organizers. An agent's success is measured using the evaluation metric in all negotiations of the tournament for which it is scheduled.

First, a *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). Since there were 19 agents, which each negotiate against 18 other agents, in the different domains, a total pair-wise tournament in the qualifying round is impossible. Therefore, 19 agents was divided to three groups (pools) randomly, and the best three agents in social welfare and individual utility in each pool proceed to the final round. It took two weeks to finish the all pools of the qualifying round. In ANAC 2014, we didn't allow the updating agents between the qualifying round and the final round.

The final round was played among the agents that achieved the best scores (individual utility and social welfare) in each pool during qualifying. The domains and preference profiles are same as the qualifying round. The entire pairwise matches played among 10 agents, and the final ranking of ANAC 2014 was decided. In the final, a single tournament consists of $10 \times 9/2 \times 2 \times 12$ $(domains) = 1080$ negotiation sessions.[2] Again, each single tournament was repeated five to prohibit the learning from the previous tournaments. To reduce the effect of variation in the results, the tournament was repeated 5 times, and the final score means the average of the five trials.

3 Competition Domains and Agents

3.1 Scenario Descriptions

The ANAC is aimed towards modeling multi-issue negotiations in uncertain, open environments, in which agents do not know what the preference profile of the opponent is. The various characteristics of a negotiation scenario such as size, number of issues, opposition, discount factor and reservation value can have a great influence on

[2]The combinations of 10 agents are $10 \times 9/2$, however, agents play each domain against each other twice by switching the roles.

Table 1 The domains used in ANAC 2014

ID	Number of issues	Size	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

the negotiation outcome. Therefore, we generated three types of domains and profiles in the competition because the nonlinear domains are generated easily. Especially, in the qualifying round and final round, we used all 12 scenarios by allocating different discount factors and reservation values to three types of domains and profiles. In other words, they have vary in terms of the number of issues, the number of possible proposals, the opposition of the preference profiles and the mean distance of all of the points in the outcome space to the Pareto frontier (see Table 1). The shapes of the outcome spaces of each scenario are represented graphically in Fig. 2.

In generating the domains for the competition, the agents can negotiate across a variety of negotiation scenarios. In addition, the challenge of ANAC 2014 is on negotiating with nonlinear utility functions as well as dealing with large-scale outcome space. Up to this year, additive utility functions have been employed to represent agents' preferences in ANAC. Although this type of functions is compact and easy to process, they cannot represent preferential interdependencies where in many real life problem we have. Therefore, we employed the constrains-based and large-sized utility functions in ANAC 2014.

3.2 Agent Descriptions

ANAC 2014 had 21 agents, registered from 13 institutes from 8 countries: GWDG, Germany; Sun Yat-sen University, China; Bar-Ilan University, Israel; University of Isfahan, Iran; Tokyo University of Agriculture and Technology, Japan; IIIA-CSIC Barcelona, Spain; Shizuoka University, Japan; Royal Holloway University of London, U.K.; University of Electro-Communications, Japan; Delft University of Technology, The Netherlands; Hunan University, China and Nagoya Institute of Technology, Japan ($\times 9$).

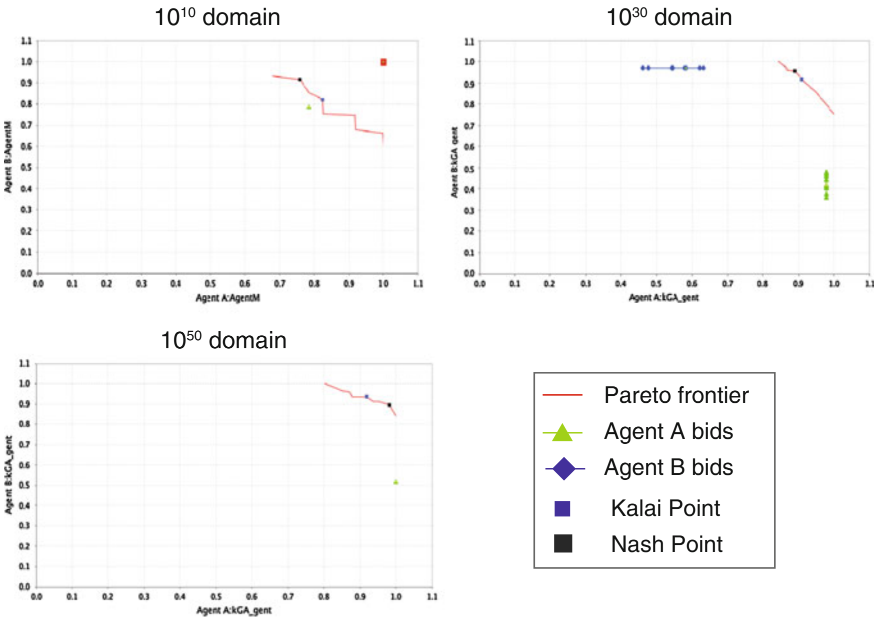

Fig. 2 Acceptance outcome space

The final round in ANAC 2014 had ten teams from eight different universities, as listed in Table 2. They are the winners of the qualifying round. In the rest of the chapter in this book, we provide sections of the individual strategies of the ANAC 2014 finalists based on descriptions of the strategies provided by the teams.

4 Competition Results

4.1 Qualifying Round

First, a *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.

Table 2 Team members and agent names in ANAC 2014

No.	Team members	Affliction	Agent name
1	Eden Shalom Erez	Bar Ilan University	DoNA
	Inon Zuckerman	Ariel University	
2	Farhad Zafari	University of Isfahan	BraveCat
	Faria Nasiri Mofakham		
3	Shinji Kakimoto	Tokyo University of Agriculture and Technology	kGAgent
	Katsuhide Fujita		
4	Motoki Sato	Nagoya Institute of Technology	WhaleAgent
5	Makoto Niimi	Nagoya Institute of Technology	AgentM
6	Dave de Jonge	IIIA-CSIC Barcelona	Gangster
7	Yuichi Enoki	Nagoya Institute of Technology	E2Agent
8	Yoshiaki Kadono	Shizuoka University	AgentYK
9	Satoshi Takahashi	University of Electro-Communications	Sobut
10	Balint Szollosi-Nagy	Delft University of Technology	Group2Agent
	Marta Skarzynska		
	David Festen		
11	Edwin Yaqub	Gesellschaft fur wissenschaftliche Datenverarbeitung (GWDG)	AgentQuest
12	Naiqi Li	Sun Yat-sen University	Flinch
	Zhansong Li		
13	Akiyuki Mori	Nagoya Institute of Technology	Atlas
14	Yoshitaka Torii	Nagoya Institute of Technology	agentTRP
15	Shota Morii	Nagoya Institute of Technology	Aster
16	Yoshihito Sano	Shizuoka University	AgentTD
	Tomohiro Ono		
	Takumi Wakasa		
17	Bedour Alrayes	Royal Holloway University of London	Simpatico
	Paulo Ricca		
	Ozgur Kafali		
	Kostas Stathis		
18	Taniguchi Keiichiro	Nagoya Institute of Technology	ArisawaYaki
19	Kimata	None	Simple ANAC 2013

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. 3 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. 4 Average scores of each agent in the qualifying round (pool2)

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

Figures 3, 4 and 5 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 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.

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

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

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 under the 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, a total of 5 times.

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

Table 3 Tournament results in the final round (Individual utility)

Rank	Agent	Score	Variance
1	Agent M	0.754618239	3.12×10^{-5}
2	DoNA	0.742245035	9.31×10^{-6}
3	Gangster	0.740674889	3.49×10^{-6}
4	WhaleAgent	0.710740252	3.90×10^{-5}
5	GROUP2Agent	0.708401404	6.38×10^{-5}
6	E2Agent	0.703955008	2.85×10^{-5}
7	kGAgent	0.676595111	5.02×10^{-5}
8	AgentYK	0.666450943	2.38×10^{-5}
9	BraveCat	0.661940343	2.84×10^{-5}
10	ANAC2014Agent	0.627684701	1.71×10^{-5}

Table 4 Tournament results in the final round (Social welfare)

Rank	Agent	Score	Variance
1	Agent M	1.645412137	4.12×10^{-5}
2	Gangster	1.627451908	1.21×10^{-5}
3	E2Agent	1.608936143	1.39×10^{-5}
4	WhaleAgent	1.603199277	3.55×10^{-5}
5	AgentYK	1.569877186	1.16×10^{-4}
6	GROUP2Agent	1.56154598	8.46×10^{-5}
7	BraveCat v0.3	1.545384774	3.11×10^{-5}
8	DoNA	1.473686528	3.89×10^{-5}
9	ANAC2014Agent	1.469972333	1.12×10^{-4}
10	kGAgent	1.463168543	4.32×10^{-4}

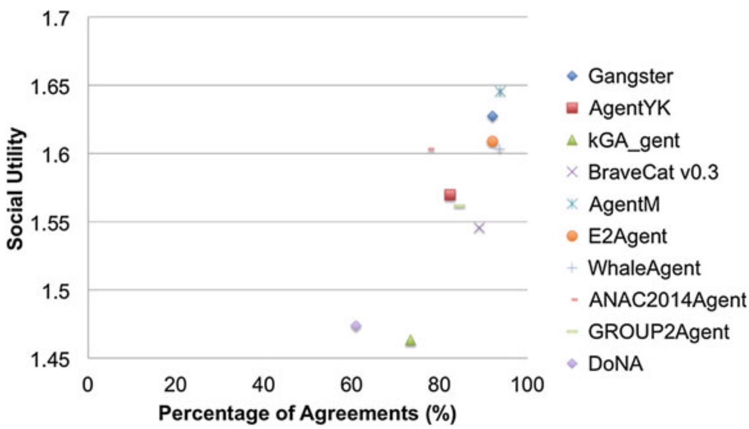

Fig. 6 Plotting graph between the percentage of agreements and the social welfare (A correlation coefficient = 0.8735)

to AgentM agent ($600); The second place in the individual category was awarded to the DoNA ($200); The second place in the social welfare was awarded to the Gangster ($200).

In more detail, we can analyze the relationships between the social welfare and other measures. As figures 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 (Figs. 6 and 7).

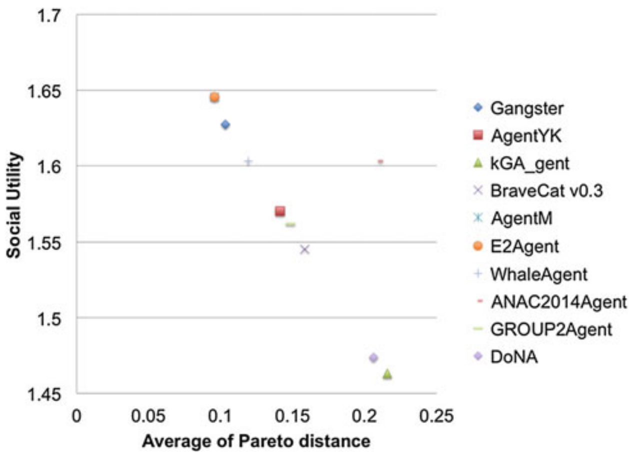

Fig. 7 Plotting graph between the average of pareto distance and the social welfare (A correlation coefficient = −0.9994)

5 Conclusion

This paper describes the fifth automated negotiating agents competition. Based on the process, the submissions and the closing session of the competition we believe that our aim has been accomplished. Recall that we set out for this competition in order to steer the research in the area bilateral multi-issue closed negotiation. 21 teams have participated in the competition and we hope that many more will participate in the following competitions.

ANAC also has an impact on the development of GENIUS. We have released a new, public build of GENIUS[3] containing all relevant aspects of ANAC. In particular, this includes all domains, preference profiles and agents that were used in the competition. This will make the complete setup of ANAC available to the negotiation research community. Not only have we learnt from the strategy concepts introduced in ANAC, we have also gained understanding in the correct setup of a negotiation competition. The joint discussion with the teams gives great insights into the organizing side of the competition.

To summarize, the agents developed for ANAC will proceed the next step towards creating autonomous bargaining agents for real negotiation problems. We plan to organize the next ANAC in conjunction with the next AAMAS conference.

Acknowledgments The authors would like to thank the team of masters students at Nagoya Institute of Technology, Japan for their valuable help in the organization of the ANAC 2014 competition.

[3] http://ii.tudelft.nl/genius.

References

1. Baarslag, T., Fujita, K., Gerding, E.H., Hindriks, K.V., Ito, T., Jennings, N.R., Jonker, C.M., Kraus, S., Lin, R., Robu, V., Williams, C.R.: Evaluating practical negotiating agents: results and analysis of the 2011 international competition. Artif. Intell. **198**, 73–103 (2013)
2. Baarslag, T., Hindriks, K.V., Jonker, C.M., Kraus, S., Lin, R.: The first automated negotiating agents competition (ANAC 2010). In: New Trends in Agent-Based Complex Automated Negotiations, pp. 113–135. Springer (2012)
3. Fujita, K., Ito, T., Baarslag, T., Hindriks, K.V., Jonker, C.M., Kraus, S., Lin, R.: The second automated negotiating agents competition (ANAC 2011). In: Complex Automated Negotiations: Theories, Models, and Software Competitions, pp. 183–197. Springer (2013)
4. Gal, Y.K., Ilany, L.: The fourth automated negotiation competition. In: Next Frontier in Agent-Based Complex Automated Negotiation, pp. 129–136. Springer (2015)
5. Ito, T., Klein, M., Hattori, H.: A multi-issue negotiation protocol among agents with nonlinear utility functions. Multiagent Grid Syst. **4**(1), 67–83 (2008)
6. Kraus, S.: Strategic Negotiation in Multiagent Environments. MIT Press (2001)
7. Lin, R., Kraus, S., Baarslag, T., Tykhonov, D., Hindriks, K.V., Jonker, C.M.: Genius: an integrated environment for supporting the design of generic automated negotiators. Comput. Intell. **30**(1), 48–70 (2014)
8. McKelvey, R.D., Palfrey, T.R.: An experimental study of the centipede game. Econometrica **60**(4), 803–36 (1992)
9. Osborne, M., Rubinstein, A.: Bargaining and Markets. Economic Theory, Econometrics, and Mathematical Economics. Academic Press (1990)
10. Osborne, M.J., Rubinstein, A.: A Course in Game Theory. MIT Press Books. The MIT Press (1994)
11. Rubinstein, A.: Perfect equilibrium in a bargaining model. Econometrica **50**(1), 97–109 (1982)
12. Williams, C.R., Robu, V., Gerding, E.H., Jennings, N.R.: An overview of the results and insights from the third automated negotiating agents competition (ANAC 2012). In: Novel Insights in Agent-based Complex Automated Negotiation, pp. 151–162. Springer (2014)