Architecture of an Agent-Based Negotiation Mechanism

Jen-Hsiang Chen
DKERG
School of MIS,
Coventry University,
UK
j.chen@
coventry.ac.uk

Rachid Anane
DKERG
School of MIS,
Coventry University,
UK
r.anane@
coventry.ac.uk

Kuo-Ming Chao
DKERG
School of MIS,
Coventry University,
UK
k.chao@
coventry.ac.uk

Nick Godwin
DKERG
School of MIS,
Coventry University,
UK
a.n.godwin@
coventry.ac.uk

Abstract

One of the central issues in facilitating mobile teamwork is the creation and establishment of teams from autonomous agents. It is widely accepted that team building assumes an expression of, and agreement on, common interests. This paper describes a new approach to the provision of mechanisms to facilitate the creation of teams and to help resolve conflict through automated negotiation. The negotiation mechanism is implemented by a combination of a game theory approach and a co-evolutionary approach. This scheme involves a process that iterates over the generation of a set of strategies by the co-evolutionary approach, the encoding of these strategies into a payoff matrix, and the reasoning on the matrix by the game theory approach in order to find an optimised point. The process terminates when the game theory approach finds an optimised point that satisfies both agents. The main advantage of this system is that agents, without knowing each other's strategies, agree on an optimised solution that conforms to Nash equilibrium and Pareto efficiency.

1. Introduction

Collaboration between software agents, and subsequently team building and teamwork, are increasingly seen as viable strategies for dealing with the complexity of distributed tasks. This approach seeks to take advantage of the nature of agents as autonomous, proactive and social entities. Some of the benefits that accrue form this approach include the use and reuse of services already offered by agents, their complementary tasks and the setting up of vertical processes such as the supply chain. In this context, automated negotiation plays a crucial role in resolving conflict and in bringing together different agents for work on collaborative projects.

One useful perspective on negotiation is to view it as a distributed search on a space of potential agreements [1]. It is a process that is supported by the selection and communication of strategies expressed by means of a negotiation protocol. The protocol is designed to ensure orderly and meaningful communication between agents, whereas strategies are formulations of exhaustive plans of action that also cover contingency situations [2]. The identification of effective negotiation strategies in implementing negotiation mechanisms has presented workers in the field with a significant challenge. In their search for a solution many researchers have opted exclusively either for a co-evolutionary approach such as Genetic Algorithms (GAs) or for a game theory approach. Whilst each approach has its own merits they both suffer from serious limitations. In the GA approach global optimised solutions are difficult to find, a limitation that leads to sub-optimal solutions. In game theory the complexity of the search space for many practical application domains preempts the generation of the payoff matrix during negotiation.

The aim of this paper is to present the design and implementation of a negotiation mechanism that combines an evolutionary approach with a game theory approach. The symbiosis between the two approaches leads to a system that takes advantage of their strengths whilst avoiding their weaknesses.

The paper is organised as follows. In section 2 the game theory approach is presented and illustrated with an example. Section 3 deals with a co-evolutionary model and its implementation. Section 4 gives an architectural description of the proposed automated negotiation mechanism. Section 5 describes the application of the system to the establishment of a partnership. The last

Proceedings of the 22 nd International Conference on Distributed Computing Systems Workshops (ICDCSW'02)
0-7695-1588-6/02 $17.00 © 2002 IEEE

section presents some conclusions and directions for further work.

2. Game theory approach: concepts and example

Game theory has been for a long time the subject of study by researchers interested in competitive decision-making, bargaining and negotiation. Game theory models have provided a useful framework for the coordination of rational agents in conflict resolution. Game theory has formed the basis of most of the work on automated negotiation [3, 4, 5]. Two fundamental concepts have presided over the application of game theory to automated negotiation, Nash Equilibrium and Pareto efficiency. Nash equilibrium obtains when two agents have no longer any incentive in deviating from their strategies [6]. The convergence afforded by Nash bargaining game equilibrium makes it the most popular solution to the bargaining problem [7], in contrast to the Bayesian equilibrium concept, which was adapted to games with incomplete information [8]. This approach has been applied to various application domains. The design of negotiation mechanisms by Rosenschein and Zlotkin [3] is noteworthy for its generality and relevance to various domains. Kraus introduced non-cooperative models that incorporate the notion of time and resource restriction into worth-oriented domains [9].

Pareto efficiency is a major criterion for assessing the degree of agreement in a negotiation [10]. In this respect, Pareto efficiency is achieved by a strategy combination that increases the payoff of one agent without reducing the payoff of another agent. The wider application of game theory models to automated negotiation mechanisms has been marked by a debate over their efficiency, and their relevance in guiding the negotiation process in multi-agent system [11, 12, 13].

As an example of a game theory based approach an extended version of the traditional Trusted Third Party (TTP) game that incorporates a negotiation mechanism is presented [14]. It is called the Trusted Third Party mediated game. It is a game that is supervised by a domain independent agent, with no involvement in the game and trusted by both agents. Under this scheme, agents perform a dual role: they reason on the payoff matrix and seek equilibrium through TTP negotiation [15]. This has the advantage that it can deal with simple games as well as difficult ones. (Since a game may have no equilibrium or many equilibria in pure strategy, these cases arise with games that are considered as difficult)

TTP is supported by a communication protocol implemented by an iterative process over the making of offers and counter offers, or the acceptance or rejection of offers. Furthermore, the negotiation process is sustained by a commitment by agents to their agreements. This pledge is designed to prevent agents from reneging on their commitments in order to maximise their payoffs. It is formulated in terms of a No-Fear-of-Deviation (NFD) equilibrium, which was proposed so that agents are bound by their commitments and avoid deviation from their commitments. To that effect the protocol includes two communication actions: Guarantee and Compensation. The communication actions are designed to allow for the transformation of a difficult game into a simple one, by trading payoffs. The Trusted Third Party (TTP) agent has the ability to perform Guarantee and Compensation actions and ensures that agents honour their commitments. The Guarantee communication action is designed to ensure that agents do not choose a strategy that will lead to a worse result. The Compensation communication strategy is an incentive for another agent to play a given strategy that can lead to a desirable state. Agents make use of these communication actions in order to find an NFD equilibrium. When a stabilised state is reached it is Pareto efficient.

3. Co-evolutionary approach: concepts and implementation

The field of economic and management sciences has seen one of the first applications of evolutionary approaches to conflict resolution [16]. The repeated game case study provided a framework for the investigation of the refinement of Nash equilibria, and their elucidation as learned equilibria [17]. Other approaches saw the introduction of thresholds as evaluation criteria in the refinement of strategies [18]. Although evolutionary approaches to automated negotiation provides an alternative to game theory approaches in conflict resolution, they suffer from some limitations. Co-evolutionary approaches assume the existence of a Nash equilibrium that can be refined. Because they seek approximate solutions, co-evolutionary approaches often produce sub-optimal solutions.

The starting point of a co-evolutionary approach to negotiation is the availability of a set of initial strategies, which are subsequently manipulated in order to find an optimised solution to a specific problem. This involves the development of a utility function for evaluating strategies, so each strategy will have an assigned single value. The fitness functions will use this value to determine the degree of its usefulness as a solution [19]. The generation of optimised strategies or solutions involves the selection of the best solutions, including the parents, and the creation of offspring through mutation and crossover. The next generation is selected by means of a tournament process [20]. The weakest solutions are removed from the population and replaced by better solutions.

A co-evolutionary system based on the above model was implemented in order to support negotiation between agents that can handle multiple issues. Under this scheme, the system is provided with a generic strategy from which it generates randomly a number of specific strategies that will make up the initial population. Each strategy is encoded into its corresponding chromosome. An agent is able to specify its own utility function and fitness function. Strategies can be refined by means of the fitness function, which can be adjusted by changing its parameters. The further away from the threshold is the strategy the fitter it is. Furthermore, the mutation rate and the crossover rate can be set arbitrarily by the agent. The negotiation process is initiated by one agent, and sustained by exchanges of offers and counter offers which make up the underlying protocol. Although convergence towards a common solution is desirable and acceptable it may not represent a global optimised solution. It is this realisation that preempts the incorporation of acceptance or rejection of offers into the protocol. This, however, allows an agent to make an offer that was worse that a previous one. These characteristics enable the system to search the whole space of potential agreements. Agents negotiate over a number of issues, and strategies are created and selected in order to maximise their gain with respect to these issues.

4. An automated negotiation mechanism

We propose a system that will support the negotiation process by enabling agents to generate and select effective strategies that lead to high payoffs. Each agent in the system incorporates two major components: a Genetic Algorithm and a No-Fear-of-Deviation equilibrium algorithm. The Genetic Algorithm was described earlier, in section 3, whereas the NFD was presented in section 2. The two components exchange information internally by means of a payoff matrix, generated by the co-evolutionary approach. Each agent implements a loop over the generation and selection of strategies, the encoding of a payoff matrix and the determination of an optimised point (Figure 1). Agents exchange information externally by means of a protocol. Furthermore, the negotiation process is guided by a set of actions defined by strategies and supported by a protocol, under some constraints. Agents must negotiate on the same issues and when one agent makes an offer, the other agent must make a counter offer, accept the offer or reject it.

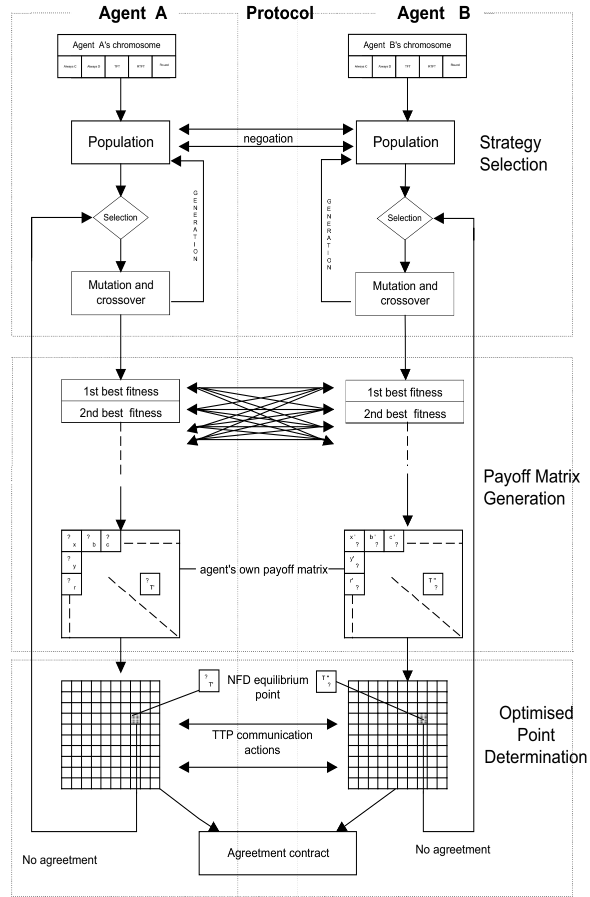

Figure 1: Automated negotiation mechanism

The negotiation system possesses a number of characteristics:

1. Agents have only limited knowledge of each other's internal information. An agent is only aware of the encoded form of a strategy not its internal structure. Furthermore, an agent has only access to its payoff matrix and utility function. It has no knowledge of the other agent's payoffs and utility function.
2. Since the co-evolutionary process is responsible for finding high payoffs for both agents by searching the space of potential agreements, deals cannot be made at that level. This ensures that the space of potential agreements is explored thoroughly, and that agents act in a rational manner by, for example, not accepting an offer that is worse than a previous one. An agent makes use of the distance between the NFD point and its threshold in order to accept or reject a deal.
3. The NFD equilibrium algorithm is applied to the payoff matrix, when required, to obtain a Nash equilibrium that is also Pareto-efficient. This condition can be satisfied by a number of solutions.

Proceedings of the 22 nd International Conference on Distributed Computing Systems Workshops (ICDCSW'02)
0-7695-1588-6/02 $17.00 © 2002 IEEE

4. If the NFD equilibrium does not satisfy the requirements of the agents, the co-evolutionary phase is revisited in order to select a new strategy.
5. Once agents agree on the payoffs produced by the NFD algorithm, it is then necessary to determine the strategies that led to these payoffs. To this end the GA process is invoked with the fitness function adjusted to reflect the new NFD equilibrium. As the GA may not be able to find an exact match to the payoff associated with the NFD equilibrium, an approximate solution is acceptable.
6. Compensation and guarantee communication actions are incorporated in order to allow agents to re-evaluate the payoffs in the matrix in their search for an equilibrium.
7. The system stops when a mutually acceptable agreement is achieved following a new equilibrium, or when the resources of the agents are exhausted.

5. Case study: establishing a partnership

The proposed system is examined by considering its application to the establishment of a partnership between a manufacturer and a car dealer. The two parties want to reach an agreement over the sale and purchase of cars that is mutually beneficial, but under some constraints. The dealer wants to buy two types of cars from the manufacturer, and is negotiating over the price and the quantity for each car. Both manufacturer and dealer have decided that they will only close the deal, and therefore cement their relationship, if they can each make a fixed minimum profit from their first transaction. The manufacturer takes into account its production costs, and the dealer refers to the resale price of the cars.

The negotiation process involves therefore four issues, the price P1 and quantity Q1 of the first car and the price P2 and quantity Q2 of the second car. The threshold for the manufacturer is the minimum profit it expects from the transaction, say TM and for the dealer it is TD. For the manufacturer the profit is calculated by deducting the total production cost from the total sale price. If C1 and C2 are the production costs for the first car and the second car respectively, then the profit PM for the manufacturer is:

$$PM = (P1*Q1+ P2*Q2) – (C1*Q1 + C2*Q2)$$

Similarly, if R1 and R2 are the resale prices by the dealer for the first car and the second car respectively, then the profit PD for the dealer is:

$$PD = (R1*Q1+ R2*Q2) – (P1*Q1 + P2*Q2)$$

The main constraint as far as profits are concerned can now be expressed as follows:

PM >= TM for the manufacturer and
PD >= TD for the dealer.

TM and TD acts as thresholds for the two parties and the purpose of the negotiation process is to aim for values for PM and PD as large as possible. The formula for calculating the profit is the utility function, whereas the fitness function is designed to measure the distance of the profit from the threshold.

Issue Value Table

p1	q1	p2	q2
2.7	500	3.85	800
2.5	800	4.1	600
2.3	1000	3.8	900
2.8	450	3.75	950
2.6	700	3.7	1100

Figure 2: Allocation of values to issues

Figure 2 shows the allocation of initial values to the issues under negotiation. As stated earlier, the agents negotiate over four issues, and these issues, either price or quantity, may be assigned values from a set of possible values. These values must, however, conform to the constraints stated above, i.e. the offer price must be higher than the production cost for the manufacturer, and lower than the resale price for the car dealer. These values form the basis of offers and counteroffers.

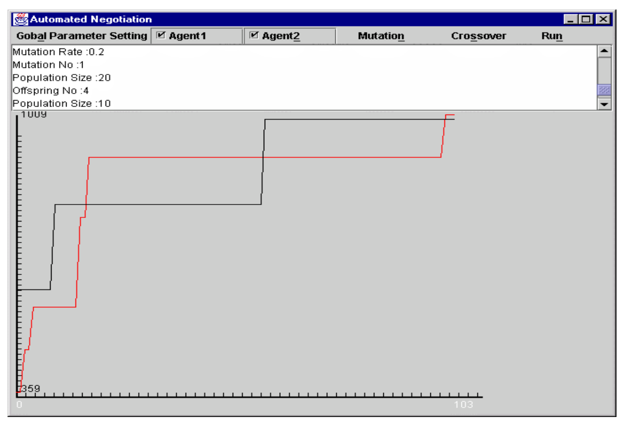

Figure 3: Automated negotiation process

The graph in Figure 3 shows clearly the convergence of the offers by the agents towards an optimised point on which they agree. The optimised point is determined by the cooperation of the co-evolutionary approach and TTP, which results in an NFD equilibrium. The system was able to determine the maximum utilities for each agent. The set of utilities for each agent is close to the set of the other agent without producing an exact match. This slight mismatch is due to the use of discrete data in the experiment.

The system initially generates ten strategies that make up the population of each agent. Each strategy from one agent is set against all the strategies in the other agent, in order to assess its overall performance. This process also allows the other agent to select the best strategy in order to express its counter offer in the next round. The system improves its population by removing the least successful strategies. The two best performing strategies are selected as parents and used to generate offspring. The number of offspring in this experiment is four. Thus, the worst four performing strategies in the population will be replaced by the generated offspring in each round. The mutation rate is set to 0.2 and the mutation number is 1 for both agents. The crossover rate and number are 0.1 and 1 respectively for both agents. The number of generations corresponding to the number of rounds where they conduct offer and counter offer is 150. With these parameters the system reaches a satisfactory outcome.

A number of experiments were carried out by using different parameter settings for each agent, such as the mutation rate, crossover rate and generation number. The negotiation process did not produce, however, significant differences from the previous results. It was noted that the size of the population has a large impact on the negotiation process. The larger the population size, the better the chances that the offers will converge quickly and the agents reach an agreement. Although the co-evolutionary process can sometimes find a global optimised point without applying the NDF equilibrium, the TTP is still required in order to ensure that the global optimised point is found, if there is one.

6. Discussions and Conclusion

The work presented in this paper is set against a background on research on automated negotiation, from both perspectives, game theory and co-evolutionary. Foremost among the central issues associated with game theory is the notion of Nash equilibrium. Although a Nash bargaining solution offers an interesting and useful perspective on bargaining behaviour, it fails to deal with bargaining games with multiple issues or incomplete information. Work by Gerding et al [6] shows that the game theory approach can accommodate multiple issues, but within a limited search space and with complete information. Similarly, the introduction of an alternative concept such as the perfect Bayesian equilibrium can provide a solution to dynamic games with incomplete information. Its application to negotiation, however, depends on the ability of the system to identify possibility. The adoption of a co-evolutionary approach to automated negotiation allowed workers to ignore many of the assumptions required for a game theory approach [7]. Random variations were used to encode issues and thresholds in the chromosomes throughout the negotiating process. The major drawback of this approach is that the agreed solution reached by an agent may be worse than the previous one. In an attempt top overcome this limitation Peyman [8] proposed the introduction of a set of family tactics in order to guide the search of optimised generic negotiation strategies. This approach relies on two fundamental assumptions. Firstly, the best generic strategies perform well in all environments and secondly, a unique Nash equilibrium must exist.

The conditions required by GA and game theory approaches when they operate in isolation restrict their range of application. The proposed automated negotiation system offers a symbiosis between a GA approach and a game theory approach that overcomes their inherent limitations. The assumption under game theory of the existence of a payoff matrix in order to reach a Nash equilibrium is often difficult to fulfill. A GA approach is, on the other hand, able to generate a set of effective strategies and related payoffs and to feed them to the game theory approach. A GA approach requires, however, the presence of a Nash equilibrium. In return, the game theory approach can determine an NFD equilibrium from the payoff matrix and thus provide a focal point for the GA approach, which will lead to a refined equilibrium. Some of the advantages that accrue from this combination are that it is possible for negotiating agents to arrive at an agreement without being aware of each other's payoffs or strategies. The transparency afforded by the system facilitates the establishment of links and the creation of teams.

Although the system is able to deal successfully with multiple negotiation issues, one direction for further work points to an exhaustive test with a realistic case study provided by our industrial collaborator. This may lead to the exploration of further aspects of negotiation and to the expansion of the system to reflect new concerns.

References

[1] Jennings. N. R., Parsons. S., Sierra. C., & Payman F., 2000, Automated Negotiation, *Proc of 5th International Conference on the Practical Application of Intelligent Agents and Multi-Agent Systems (PAAM-2000),* Manchester, UK, pp. 23-30.

[2] Barber, K. S., Han, D.C., & Liu, T.H, 2000, Strategy Selection-Based Meta-level Reasoning for Multi-agent Problem-Solving, *Agent-Oriented Software Engineering, First International Workshop AOSE 2000,* Limerick, Ireland, pp. 269-283

[3] Rosenschein, J.S., & Zlotkin, G., 1989, Negotiation and Task Sharing Among Autonomous Agents in Cooperative

Domains, *The Eleventh International Joint Conference on Artificial Intelligence*, Detroit, Michigan, pp. 912-917.
[4] Sun, Y., & Weld, D., 1995, *Automating Bargaining Agents (Preliminary Results). Technical Report,* UW-CSE-95-01-04, University of Washington.
[5] Vulkan. N., & Jennings, N. R., 1998, Efficient Mechanisms for the Supply of Services in Multi-Agent Environments, *Proc of 1st International Conference on Information and Computation Economies*, Charlestown, South Carolina, pp. 1-10.
[6] Rasmusen, E., 1994, *Games & Information: An Introduction to Game Theory*, Blackwell Publishers.
[7] Peyman, F., 2000, *Automated Service Negotiation Between Autonomous Computational Agents*, PhD thesis, University of London.
[8] Sandholm, T & Larson, K., 2000, Deliberation Equilibrium: Bargaining in Computationally Complex Problems, *National Conference on Artificial Intelligence (AAAI)*, Austin, pp. 48-55
[9] Kraus, S., 1997a, Negotiation and Cooperation in Multi-agent Environments. *Artificial Intelligence*, Vol. 94, No. 1–2, pp. 79–97.
[10] Rosenschein, J. S., & Zlotkin, G., 1994, *Rules of Encounter*, Cambridge. USA: The MIT Press.
[11] Castlefranchi, C., & Conte, R., 1997, Limits of Strategic Rationality for Agents and M-A Systems. *Proc of the 4th ModelAge Workshop on "Formal Models of Agents,* pp. 59-70.
[12] Simon, H. A., 1996. *The Sciences of the Artificial, Cambridge*, Massachusetts, MIT Press.
[13] Sycara, K. & Zeng, D., 1996, How Does an Agent Learn to Negotiate, *Intelligent Agents III: Theories, Architectures and Languages, Springer Lecture Notes in Artificial Intelligence 1193*, edited by J. Muller, M. Woolridge & N.R. Jennings, 1996, pp. 233-244.
[14] Wu, S-H., 1999, *Agent Coordination by Negotiation in Trusted-third Party Mediated Games*, PhD thesis, National Tsing Hua University, Taiwan.
[15] Soo, V-W & Wu, S-H., 2000, Negotiation Without Knowing Other Agents Payoffs in the Trusted Third-Party Mediated-Game, *Second workshop on game theoretic and decision theoretic agents*.
[16] Axelrod. R., 1987, The Evolution of Strategies in the Iterated Prisoner's Dilemma*, Genetic Algorithms and Simulated Annealing*, edited by L. Davis, London: Pittman. pp.32-41.
[17] Fudenberg D., Levine D.K., 1998, *The Theory of Learning in Games*, Cambridge: M.I.T. Press.
[18] Oliver, J. R., 1997, *On Automated Negotiation and Electronic Com*merce, PhD thesis, University of Pennsylvania.
[19] Wunderlich, R., 1998, Genetic Strategy Selection for the Finitely Repeated Prisoner's Dilemma, *http://www.roland.wunderlich.com/Papers/GameTheory.*
[20] Goldberg, D.E., 1989, *Genetic Algorithms in Search, Optimisation and Machine Learning*, Addison-Wesley.