

# Informing Coalition Structure Generation in Multi-Agent Systems Through Emotion Modelling

Martyn Lloyd-Kelly and Luke Riley

Department of Computer Science  
University of Liverpool, UK

mlk5060@liverpool.ac.uk and L.J.Riley@liverpool.ac.uk

---

## Abstract

We propose a hybrid coalition formation method for multi-agent systems that combines a rational mechanism and an emotionally-inspired mechanism to reduce the associated computational cost. To initialise coalition formation, the rational mechanism is used and in subsequent iterations, the emotional mechanism (that forms coalitions resulting from emotional reactions to aspects of interactions between agents) is used. The emotions of *anger* and *gratitude* are modelled and used as a basis to model *trust* which is in turn used to restrict the coalition state-space. We offer some discussion as to how this hybrid method offers an improvement over using a method that only considers payoff maximisation and we propose some direction for future work.

**1998 ACM Subject Classification** I.2.11 Distributed Artificial Intelligence

**Keywords and phrases** Multi-Agent Systems, Coalition Formation, Emotion

**Digital Object Identifier** 10.4230/OASISs.ICCSW.2012.81

## 1 Introduction

Multi-agent systems (*MAS*) are systems of autonomous agents capable of interacting with one another in some way [19]. In self-interested MAS, agents will attempt to achieve individual goals whilst maximising individual *payoffs*. However, under certain circumstances, agents may have to temporarily form mutually beneficial partnerships with other agents to achieve goals [3]. These partnerships are known as *coalitions* and a set of such coalitions are known as a *coalition structure*. Forming coalition structures in MAS has been shown to be an important topic that can be applied in many different areas. For example, [12] notes that it has proved useful in: e-commerce, e-business and distributed sensor networks.

In human societies, coalition formation takes into account rational aspects such as expected monetary payoffs as well as various *emotional dispositions* towards individuals. Emotions stemming from one individual's appraisal of another individual's actions, such as *gratitude* and *anger*, appear to be integral in the establishment and maintenance of *trust* [4] and such emotional appraisals, in conjunction with various cognition-based aspects, form the basis that enables one individual to trust another [6]. Trust formation of this kind has been tested using a theoretical framework by [9] who concludes that there is a high importance placed upon understanding the affective qualities of relationships. Therefore, whilst we can posit that coalitions in human societies are informed in part by maximisation of current finances, this does not fully encompass the whole spectrum of reasoning undertaken when forming coalitions as some consideration is also given to emotional aspects.

With regards to MAS, finding the best coalition structure i.e. the one that maximises social welfare, is a computationally complex activity as an exponential amount of coalitions ( $2^n - 1$ ) have to be checked [3, 12]. In this paper we discuss an attempt to model the emotions of anger and gratitude so that they may be used as a basis to form coalitions.



© Martyn Lloyd-Kelly and Luke Riley;  
licensed under Creative Commons License NC-ND  
2012 Imperial College Computing Student Workshop (ICCSW'12).  
Editor: Andrew V. Jones; pp. 81–87



OpenAccess Series in Informatics

OASIS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

# ICCSW

We propose that these relationships will reduce the computation costs placed on agents if coalition formation is highly complex, information about the environment is incomplete or uncertain, the robustness of other agent’s proposed coalition is questionable, or there is a time-bound on coalition formation. We believe that modelling anger and gratitude to inform trust is applicable for use in such a task as trust is utilised when current conditions are highly complicated and uncertain [17] and it has also been proposed that the primary function of trust is to reduce complexity [8].

Therefore, we are not aiming to improve the *quality* of coalitions, instead we aim to reduce the associated computation cost by considering the roles that anger and gratitude can play in developing a mechanism of trust for use in coalition structure generation. In this paper we do *not* detail experimental results of testing this mechanism rather, we endeavour to: make the distinction between *emotional* and *rational* coalition formation processes clear, make the novel coalition formation process explicit and propose avenues for future work.

Section 2 of this paper introduces the background on coalition games, some initial discussion on the distinction between rationality and emotion and how emotions are modelled/used in artificial intelligence. Section 3 outlines the rational and emotional coalition formation methods used, how they are hybridised and a brief example illustrating how we propose the process works. Finally, section 4 concludes with an overview of the contributions made and some proposals for future work.

## 2 Background

Coalition formation takes place in an  $n$ -person cooperative game originally defined in [18] and is denoted:  $\mathcal{G} = \langle N, v \rangle$  where  $N$  represents the set of agents ( $N = \{1, \dots, n\}$ ) and  $v$  is the *characteristic function* which assigns every coalition a real number that represents its payoff ( $v(2^N) \rightarrow \mathbb{R}$ ). The *outcome* of an  $n$ -person cooperative game is a coalition structure  $\{C^1, \dots, C^k\}$  (where the individual coalitions are distinct and exhaustive) and a *payoff vector* that divides the gains of the coalition structure between all the agents. The payoff vector is fully denoted as:  $x = (x_1, \dots, x_n) \in \mathbb{R}^n$  where  $x_i$  denotes the individual payoff for agent  $i$  and  $x_i \geq 0$  for all  $i \in N$  (see [3]). A payoff vector *satisfies* the agents of the system if it is said to be *stable* i.e. a payoff vector where no subset of agents in the system have an incentive to deviate from the current coalition structure.

Solutions of coalition games focus on what is mathematically optimal to do but as previously noted in section 1, the most mathematically rational solution is not the only consideration made in human societies. Such a solution may not be found especially if coalition formation is time constrained and/or attempted in an environment that is quite complex [16]. This paper seeks to outline a coalition formation method which simplifies the equilibrium computations required in such coalition games. As most scenarios where humans form coalitions e.g. the workplace, friendships, marriages etc. are not one-shot games (like the standard coalition game formalism) but iterative interactions, we will take inspiration from how emotions are formed and used in such iterative interactions so as to affect coalition formation. The coalition formation method that we are proposing here is an example of the “good enough, soon enough” design paradigm as we put no guarantees on if optimal solutions are found; this is left for future work.

Before continuing, we distinguish between the notions of *emotional* and *rational* coalition formation methods as understanding this difference is integral to understanding the approach outlined. As in [7], we use Axelrod’s tournament [1] as a basis for the distinction. In Axelrod’s tournament, some submitted strategies considered what action to perform in the

current round based upon an assessment of past/present/future payoffs. It is this payoff-based reasoning that we term as *rational* reasoning whereas *emotional* reasoning simply takes into consideration the *emotional disposition* of the agent towards another i.e. there is no explicit reasoning regarding past/present/future payoffs when identifying agents to form coalitions with. Emotional dispositions are concise histories that can take into account a multitude of interaction features but can be represented easily. These dispositions can be consulted quickly to not only determine whether or not to form a coalition with an agent but to also rank potential coalitions. Such an approach is inspired by the way emotions are used by Nawwab et. al to alter the preference ordering of actions when different emotions are activated [10]. We should make clear here that in no way at all do we intend for *emotion* to be interpreted as *irrational*, as is the usual dictum.

We model the emotions of anger and gratitude using the Ortony, Clore and Collins model (OCC) [11] as a basis. We take the view that anger and gratitude play a functional role, following [5] and use them to inform agents about whether to trust another. The exact implementation of this emotional approach is detailed in section 3.

### 3 The Coalition Formation Method

In this section we first outline the rational and emotional coalition formation processes in sections 3.1 and 3.2 before providing an outline of how we combine these two processes in section 3.3. Furthermore, in section 3.3 we present mathematical evidence asserting that such a hybrid technique is capable of buying advantages in computation time over the standard rational approach. Finally, we talk through a brief example illustrating how we propose the technique will work in section 3.4.

#### 3.1 The Rational Coalition Formation Process

The rational coalition formation process finds an optimal coalition structure and a stable payoff vector for the system, i.e. it exclusively considers payoffs. The issue with this process is that it is computationally expensive: the asymptotically fastest algorithm to solve the coalition structure generation problem runs in  $O(3^n)$  time [15]. For our research we will use the rational model detailed in [14], which is a distributed dialogue game that finds both an optimal coalition structure and a stable payoff vector. Theoretically however, any rational coalition structure generation model could be combined with the emotional approach outlined in sections 3.2 and 3.3.

In the model of [14], communication only occurs between agents when they offer a proposal to form a coalition. Proposals can be simply viewed as a three part tuple:  $\langle i, C, x(C) \rangle$ , where agent  $i$  proposes that coalition  $C$  forms with the coalition payoff vector denoted  $x(C)$ . If an agent in  $C$  cannot object to this coalition and payoff vector with a counter proposal then this coalition and payoff vector is said to be stable. Once a stable coalition for every agent in the system exists then a stable coalition structure also exists which entails the completion of the rational coalition formation process.

#### 3.2 The Emotional Coalition Formation Process

The emotional coalition formation process enables an agent to restrict the state-space of coalitions it has to search based upon its emotional disposition towards others. This emotional disposition is informed by the performance of those agents in past coalitions. To

achieve this, each agent  $i$  is endowed with the following that are inspired by Reilly's model of emotion in [13]:

- Anger/gratitude variable: In accordance with the OCC we model opposite pairs of emotions so there is one variable that represents anger/gratitude. This variable is denoted by  $AngGrt_j$  where  $AngGrt$  is the current value of the variable and  $j$  indicates the agent that the variable applies to. An agent can have  $AngGrt \cdot (N - 1)$  variables where  $N$  is the total number of agents in the system ( $N - 1$  is used as anger/gratitude may not be felt towards the agent experiencing them). If the  $AngGrt$  for  $i$  is negative towards another agent  $j$  then it can be inferred that  $i$  is angry with  $j$  and does not trust it. If the  $AngGrt$  for  $i$  is positive towards another agent  $j$  then it can be inferred that  $i$  feels gratitude towards  $j$  and trusts it. If the  $AngGrt$  of  $i$  towards another agent  $j$  is 0 then  $i$  is neither grateful or angry with  $j$ . Therefore, an agent may only be grateful, angry or neutral towards another, it may never activate any combination of these emotions in tandem with respect to the same agent.
- Anger/gratitude activation threshold: two constant values ( $Ang$ ,  $Grt$ ) that applies to an agent's  $AngGrt$  variable. If  $i$ 's  $AngGrt_j \leq Ang$ , the effect of anger is manifested in  $i$  towards  $j$ . If  $i$ 's  $AngGrt_j \geq Grt$ , the effect of gratitude is manifested in  $i$  towards  $j$ . The values of these variables could be varied in order to make agents more/less trusting.
- Anger/gratitude effect: Prescribes what happens when an agent's  $AngGrt$  value towards another agent is  $\leq Ang$  or  $\geq Grt$ . Anger and gratitude have opposing effects: if  $i$ 's  $AngGrt_j \leq Ang$  then  $i$  will not include  $j$  in its coalition state-space search. Conversely, if  $AngGrt_j \geq Grt$  then  $i$  will include  $j$  in its coalition state-space search. The behaviour of agents during the course of their interactions are not modified due to the activation of emotions, only the coalition state-space is affected.

For the purposes of this paper, the value of  $AngGrt$  is altered according to whether or not the coalition succeeds or fails however, the inputs to this emotional disposition alteration are context-dependent and can vary. If an agent  $i$  joins a coalition  $C$  with another agent  $j$  then the following two situations may occur in context of this paper:

1. The coalition  $C$  succeeds – all agents  $i \in C$  increase all  $AngGrt_j$  variables (where  $j \in C \setminus \{i\}$ ) by some amount.
2. The coalition  $C$  fails – all agents  $i \in C$  decrease all  $AngGrt_j$  variables (where  $j \in C \setminus \{i\}$ ) by some amount.

The value of  $AngGrt$  implies a notion of anger/gratitude intensity as it is possible that  $i$ 's  $AngGrt_j$  and  $AngGrt_k$  variables may infer anger towards both agents but  $i$  is less angry with  $j$  than it is with  $k$  as  $AngGrt_j$  may be equal to -30 whilst  $AngGrt_k$  may equal -40. Variations in intensity create preference orderings as agent  $i$  will propose a coalition with the agent who has the highest  $AngGrt$  value first. If the proposal is refused, the agent will then propose a coalition with the agent who has the next highest  $AngGrt$  value and so on.

### 3.3 Hybridisation of Rational and Emotional Processes

Initially, all  $AngGrt$  variables for each agent are set to 0 i.e. all agents are emotionally neutral towards all others. So to make a decision about who to form a coalition with, agents consult the rational coalition formation process detailed in section 3.1 *only*. However, in subsequent rounds, if any of agent  $i$ 's  $AngGrt$  values equal  $Ang$  or  $Grt$  then the emotion-based coalition formation approach is used in order to determine who  $i$  will work with, with

no input from the rational coalition formation approach unless all *AngGrt* variables for  $i$ 's coalition choices are equal.

Agent  $i$  informs other agents of its anger/gratitude by communicating the tuple  $\langle i, \{+, -, =\}, j \rangle$  where  $+$  represents gratitude,  $-$  represents anger and  $=$  represents emotion neutrality. As stated, it is initially assumed that  $\forall i, j \in N$  the tuple  $\langle i, =, j \rangle$  holds.

Now for all agents  $k \in N$ , agent  $k$  will know that a proposal for a coalition has to include  $i$  and not  $j$  if  $\langle i, -, j \rangle$  holds, which restricts the state space to search. Alternatively, both  $\langle i, +, j \rangle$  and  $\langle j, +, i \rangle$  have to hold to restrict the state space as both  $i$  and  $j$  have to want to be in a coalition together before it is fair to force them into one. If it is the case that  $i$ 's *AngGrt* is negative/positive towards all agents then agent  $i$  will try to form a coalition with the agent(s) it is least angry with/most grateful to, respectively (explained in section 3.2).

The advantage of these restrictions are clear: consider a coalition search space of  $2^n - 1$  coalitions (for  $n$  agents) and one anger constriction of  $\langle i, -, j \rangle$  then there will be  $(2^n - 1) - (2^{n-2})$  coalitions to check as  $2^{n-2}$  is the amount of potential coalitions that any two agents share. Likewise, given the same coalition search space and one reciprocal gratitude relationship between agent  $i$  and  $j$ , then the amount of coalitions to search is reduced to  $(2^n - 1) - (2^{n-1})$  as  $2^{n-1}$  is the amount of potential coalitions for  $n$  agents that hold one of  $i$  and  $j$  but not the other. Table 1 shows a 4-person coalition game where there are  $2^4 - 1$  possible coalitions that need to be checked. With the addition of one anger tuple  $\langle 1, -, 2 \rangle$  the underlined coalitions do not need to be searched. The underlined coalitions are equal to  $2^{4-2}$  so the full amount of possible coalitions to search is:  $(2^4 - 1) - (2^{4-2}) = 11$ , which is a reduction from the full 15 used in the rational model. The benefit of this approach increases as more anger and gratitude relationships hold, especially if  $n$  becomes unmanageable for standard rational coalition structure generation techniques.

■ **Table 1** State space reduction using emotional coalition formation process.


$ C  = 1$	$ C  = 2$	$ C  = 3$	$ C  = 4$
{1}	{3, 4}	{2, 3, 4}	<u>{1, 2, 3, 4}</u>
{2}	{2, 4}	{1, 3, 4}	
{3}	{2, 3}	<u>{1, 2, 4}</u>	
{4}	{1, 4}	<u>{1, 2, 3}</u>	
	{1, 3}		
	<u>{1, 2}</u>		

### 3.4 Example Process

When the *AngGrt* variables for multiple agents in a MAS begin to equal or surpass their *Ang/Grt* values emotions, a natural representation of these relationships, inspired by [2], is a directed graph with two arrow types:

- Pointed arrow heads from agent  $i$  to agent  $j$  for the tuple  $\langle i, +, j \rangle$
- Flat arrow heads from agent  $i$  to agent  $j$  for the tuple  $\langle i, -, j \rangle$

Figure 1 gives an example of such a directed graph for a MAS coalition game. The emotional dispositions asserted in this game are:  $\langle 1, +, 2 \rangle$ ,  $\langle 2, +, 1 \rangle$ ,  $\langle 2, +, 5 \rangle$ ,  $\langle 3, +, 2 \rangle$ ,  $\langle 4, -, 2 \rangle$ ,  $\langle 5, -, 3 \rangle$  and  $\langle 5, +, 4 \rangle$ . If no arrow exists between two agents then these agents are emotionally indifferent to each other.



■ **Figure 1** Directed graph denoting emotional dispositions in a 5 agent coalition game.

In this 5 agent game, the agents know that a 2-person coalition is at least needed to complete a task but a task gets easier with more agents. Let us assume that agent 1 makes the first proposal for a coalition (the mechanism which determines proposal orders is outside the scope of this paper). Agent 1 has to include agent 2 in its proposal as the agents have reciprocal positive emotions (so the current best coalition for agent 1 is  $C_1 = \{1, 2\}$ ).

Agent 1 then has a choice as to the next agent to invite into the existing coalition and so considers both  $C_2 = \{1, 2, 3\}$  and  $C_3 = \{1, 2, 5\}$ . Agent 4 is not considered as agent 2 is currently angry with agent 4 and therefore does not trust it, so agent 2 would reject the coalition of  $\{1, 2, 4\}$ . Agent 1 therefore has a choice between 2 agents: 3 and 5. As agent 1 is emotionally indifferent to both, the rational method is used to decide between the different coalitions. The best possible payoff for each coalition of agent 1's are:  $x_1(C_3) > x_1(C_2) > x_1(C_3)$ , so agent 1 deems  $C_3$  to be the best coalition. Adding agent 5 to the proposal means that agent 3 should not be considered as agent 5 feels anger towards agent 3 and does not trust it. As no more agents can be added to the coalition,  $C_3$  is then proposed, accepted and formed. Notice here that coalition  $C_3$  was chosen by agent 1 out of a possible  $2^{5-1}$  coalitions yet agent 1 only considered and compared 3 different possible coalition proposals ( $C_1$ ,  $C_2$  and  $C_3$ ).

The acceptance of  $C_3$  leaves the remaining two agents (3 and 4 who are indifferent to each other) to form a coalition if they want to complete a task. So,  $C_4 = \{3, 4\}$  is also proposed, accepted and then formed, resulting in a coalition structure of:  $\{\{C_3\}, \{C_4\}\}$ .

#### 4 Conclusion and Future Work

We have discussed the details of a hybrid coalition formation method that uses a previously established rational coalition formation process augmented with an emotionally-inspired coalition formation process. We have outlined our proposal for how this method will work in context of MAS and have outlined the benefits to computation costs that the method imparts. The modelling of anger and gratitude has also been made explicit and how these emotions are used to create a notion of trust between agents in MAS. Finally, we have outlined and discussed an example MAS in which this hybrid method facilitates a reduction in computational time with regards to forming coalitions.

A number of directions for future work have been outlined in the paper. The obvious continuation of this work would be to implement agents in a MAS who could use this hybrid coalition formation method. Ideally, this would be performed in context of a simulation test-bed so as to investigate whether the method outlined buys us reduced computation time as anticipated. Further questions may also be asked i.e. how scalable is the method and does this approach inadvertently produce coalitions of better quality?

Furthermore, it would be interesting to identify other variables that may affect the emotional disposition of an agent rather than just goal success/failure. Such variables could

include time taken to complete the goal specified, effort expended by other agents, consequences of the actions of another agent in context of goal achievement, shares of payoffs distributed etc. Such considerations could give rise to coalitions that are of better quality as more factors are taken into consideration.

Finally we may also consider the effects of varying anger/gratitude activation thresholds in order to create agents that are more/less trusting. This notion of *emotional characters* may be used to extend simulations so that we may identify those characters that are the most/least successful with respect to the quality of the coalitions formed.

---

## References

- 1 Robert Axelrod. *The Evolution Of Cooperation*. Basic Books, Inc., 1984.
- 2 Claudette Cayrol and Marie-Christine Lagasquie-Schiex. Coalitions of arguments: A tool for handling bipolar argumentation frameworks. *Intelligent Systems*, 25(1):83–109, 2010.
- 3 Georgios Chalkiadakis, Edith Elkind, and Michael Wooldridge. *Computational Aspects of Cooperative Game Theory*. Morgan & Claypool Publishers, 2011.
- 4 Jennifer R. Dunn and Maurice E. Schweitzer. Feeling and believing: The influence of emotion on trust. *Journal of Personality and Social Psychology*, 88:736–748, 2005.
- 5 Nico H. Frijda. *The Emotions*. Cambridge University Press, 1987.
- 6 David J. Lewis and Andrew Weigert. Trust as a social reality. *Social Forces*, 63:967–985, 1985.
- 7 Martyn Lloyd-Kelly, Katie Atkinson, and Trevor Bench-Capon. Emotion as an enabler of co-operation. In *4th International Conference on Agents and Artificial Intelligence, ICAART*, 2012.
- 8 Niklas Luhmann. *Trust and Power*. Chichester: Wiley, 1979.
- 9 Daniel J. McAllister. Affect and cognition-based trust as foundations for interpersonal cooperation in organizations. *The Academy of Management Journal*, 38:24–59, 1995.
- 10 Fahd Saud Nawwab, Trevor Bench-Capon, and Paul E. Dunne. Emotions in rational decision making. In *Lecture Notes in Computer Science*, volume 6057, pages 273–291. Springer, 2010.
- 11 Andrew Ortony, Gerald L. Clore, and Allan Collins. *The Cognitive Structure of Emotions*. Cambridge University Press, 1988.
- 12 Talal Rahwan. *Algorithms for Coalition Formation in Multi-Agent Systems*. PhD thesis, University of Southampton, 2007.
- 13 W. Scott Neal Reilly. *Believable Social and Emotional Agents*. PhD thesis, Carnegie Mellon University, 1996.
- 14 Luke Riley, Katie Atkinson, and Terry Payne. A dialogue game for coalition structure generation with self-interested agents. In *The Fourth International Conference on Computational Models of Argument, COMMA*, 2012.
- 15 Tuomas Sandholm, Kate Larson, Martin Andersson, Onn Shehory, and Fernando Tohmé. Coalition structure generation with worst case guarantees. *Artificial Intelligence*, 111:209–238, 1999.
- 16 Leen-Kiat Soh and Costas Tsatsoulis. Satisficing coalition formation among agents. In *Proceedings of AAMAS*, 2002.
- 17 James D. Thompson. *Organizations in action: Social science bases in administrative theory*. New York: McGraw-Hill, 1967.
- 18 John von Neumann and Oskar Morgenstern. *The Theory of Games and Economic Behavior*. Princeton University Press, 1944.
- 19 Michael Wooldridge. *An Introduction to Multi-Agent Systems Second Edition*. John Wiley & Sons, 2009.