Practical Distributed Coalition Formation via Heuristic Negotiation in Social Networks

Sarvapali D. Ramchurn, Enrico Gerding, Nicholas R. Jennings, Hu Jun
Agents, Interaction, and Complexity Group
School of Electronics and Computer Science
University of Southampton
{sdr,eg,nrj,jh09v}@ecs.soton.ac.uk

ABSTRACT

We present a novel framework for decentralised coalition formation in social networks, where agents can form coalitions through bilateral negotiations with their neighbours. Specifically, we present a practical negotiation protocol and decision functions that enable agents to form coalitions with agents beyond their peers. Building on this, we establish baseline negotiation strategies which we empirically show to be efficient (agreements are reached in few negotiation rounds) and effective (agreements have high utility compared to a centralised approach) on a variety of network topologies. Moreover, we show that the average degree of social networks can significantly affect the performance of these strategies.

1. INTRODUCTION

Coalition formation (CF) is one of the fundamental approaches in multi-agent systems for establishing collaboration between groups or networks of self-interested agents. Traditionally, CF has been studied in terms of its algorithmics and economics [13]. For example, while [13] introduced algorithms to form the best coalitions from a set of agents, [3] study how to divide payoffs obtained as a result of coalitional actions. However, existing solutions tend to be centralised and (or) not scalable (dealing with < 30 agents), and typically assume that agents are cooperative and (or) have complete access to the information about other agents' preferences. Hence, such solutions are not readily applicable to problems involving large numbers of self-interested (i.e., individual utility maximising) agents that are ignorant of their neighbours' preferences. In turn, in the economics literature, studies of CF with self-interested agents are typically focused on determining equilibrium outcomes (based on some allocation rules imposed on the population or network of agents) of the game rather than on the reasoning that each individual agent has to perform (in the absence of information about other agents' payoffs) to come to such equilibrium outcomes [6]. In short, most existing approaches ignore the practical process by which large numbers of self-interested autonomous agents can form coalitions in a fully decentralised manner and decide on the divisions of payoff without complete information. Crucially, there exists no framework to regiment the communication and negotiation interactions that agents need to engage in to form coalitions in such situations.

To address this shortcoming, in this paper we develop the first decentralised coalition formation (DCF) framework that allows agents to form coalitions through negotiation in a completely decentralised manner, and without knowing other agents' preferences. In so doing, we further assume that agents are connected through a *social network*, where interactions can only occur between neighbours in the network. This is common in markets, for example, where an agent may only form groups with those agents it knows or trusts [6].

In more detail, this paper advances the state of the art as follows. First, we propose a novel bargaining protocol whereby agents can reach agreements on which coalitions to form, which coalitional actions to execute, and how to divide the resulting surplus. While agents can only negotiate with their neighbours, this protocol ensures that coalitions can be formed across the network and negotiations do not result in deadlocks. Second, we devise novel negotiation decision functions that effectively exploit the network structure and build on heuristic negotiation techniques that have been designed for two-agent negotiations [7, 12]. We use these techniques because they have been shown to be effective in practical conflict settings where there is limited or no access to statistical models of agents' preferences. Third, we empirically evaluate the system (involving a hundred agents) and establish negotiation efficiency benchmarks for DCF on three types of social networks. In so doing, we show how the new negotiation strategy we propose can improve social welfare by up to 10% (compared to a baseline) and how the average degree of the network can significantly constrain or enlarge the negotiation space, hence impacting the sizes of coalitions formed and the payoffs of the agents. Taken altogether, our results provide new insights on practical DCF within large scale systems and open up several new areas of research in the area of DCF.

The rest of this paper is structured as follows. Section 2 discusses relevant literature building upon which Section 3 defines the model involving the social network, the negotiation protocols, and the offers exchanged. Then, Section 4 defines the negotiation decision functions that include the negotiation strategy of an agent. Section 5 empirically evaluates the performance of the DCF process and Section 6 concludes.

2. RELATED WORK

Our approach to coalition formation is inspired from the seminal work of Myerson (summarised in [9]), who argued agents will form coalitions through some negotiation process where they are assumed to be able to compute the exact payoff they should get in their respective coalitions. Following on from this, a number of works have addressed the payoff distribution problem among self-interested agents [9, 2, 14, 11]. While these approaches present interesting stability concepts (i.e., payments to agents and assignments to coalitions) they do not consider the negotiation process that happens for the agents to reach such outcomes when the agents have no information about other agents' preferences or cannot communicate with every other agent in the system.

The problem where coalitions are constrained by the network which connect the agents (i.e., when agents interact in a peer-to-peer fashion within a social network) has also been studied within the network or group formation literature in the field of economics [6, 16, 4]. These approaches consider a number of bargaining protocols and allocation rules that allow the formation of coalitions but only study the equilibrium properties of such protocols rather than provide the decision functions needed by the agents to come to agreements as to what coalitions to form. Hence, despite the numerous possible applications of coalition formation (e.g., grid computing, web services, or sensor networks [5]), existing work on this topic has remained largely theoretical and applicable to small systems (typically not larger than 30 agents), and the process by which coalitions can be formed in practice is not adequately addressed in the current literature. An exception is the work by [8], which considers the dynamic coalition formation problem where agents can create, negotiate, and form new coalitions as and when needed. However, [8] only introduces a high-level framework for studying such interactions and this framework is not evaluated, nor does it provide specific negotiation protocols or strategies that agents could use to form coalitions.

Turning to the multi-agent bargaining literature, since the seminal work of Faratin et al. [7], a number of negotiation mechanisms have been developed to solve resource allocation problems in multi-agent systems [12, 15]. While these constitute robust decentralised mechanisms for conflict resolution, none of them have attempted to instantiate the type of negotiation process needed to form coalitions where both the actions and the sharing of payoffs among the coalition members need to be agreed upon. Hence, our work aims to bridge the gap between the bilateral negotiation and coalition formation domains by specifically developing a set of negotiation protocols and decision making functions for automated negotiation for coalition formation.

3. THE MODEL

In this section, we present our model of DCF. Let $I = \{1, 2, \ldots, |I|\}$ denote the set of agents. Furthermore, let $S = \{C_1, \ldots, C_n\}$ denote a coalition structure (i.e., a partition of I), where $C_k \subseteq I$, $C_k \neq \emptyset$ denotes a coalition and each agent $i \in I$ is part of *exactly* one coalition $C \in S$. That is, $C_k \cap C_l = \emptyset$ for all $C_k, C_l \in S$ whenever $k \neq l$, and $\cup_{C \in S} C = I$ (e.g., if C is the grand coalition then $|S| = 1$). Coalitions with only a single agent are called singletons. Importantly, agents within the same coalition can perform some joint tasks. Thus, each agent i can perform an atomic action $\alpha_i \in D_i$ where $D_i = \{\alpha_i^1, \alpha_i^2, \ldots, \alpha_i^{|D_i|}\}$ is the set of all possible actions for agent i. Then, the agents' joint actions in a coalition C is given by $\boldsymbol{\alpha}_C = \cup_{i \in C} \alpha_i$ (we omit brackets around α_i for clarity). Furthermore, each agent i incurs a cost $\beta_i(\boldsymbol{\alpha}_C) \in \mathbb{R}^+$ for performing a given action, which depends on the actions taken by other agents in the coalition and/or the size of the coalition. However, we assume that, *ceteris paribus*, these costs are non-increasing with the size of the coalition. Thus, we assume coalitions typically enable synergies (e.g., agents sharing resources or services to perform their individual actions). Specifically, for an agent i:

$$\beta_i(\boldsymbol{\alpha}_C) \geq \beta_i(\boldsymbol{\alpha}_C \cup \alpha_j) \text{ for all } j \notin C, \alpha_j \in D_j \quad (1)$$

Also, as is standard, the coalition as a whole receives a payoff or value $V(\boldsymbol{\alpha}_C)$ for executing the joint actions. The purpose of the negotiation process is to both generate coalitions and for each agent i to vie for a share $v_i \in \mathbb{R}^+$ of the coalition payoff, where $\sum_{i \in C} v_i = V(\boldsymbol{\alpha}_C)$. Given this, an agent i has the following utility function, which it tries to maximise through negotiation:

$$u_i(\boldsymbol{\alpha}_C, v_i) = v_i - \beta_i(\boldsymbol{\alpha}_C) \quad (2)$$

We assume that agents' cost functions are *private* information (i.e., unknown to other agents), while the value function is common knowledge. Furthermore, as motivated in Section 1, we assume that not all coalitions are possible due to constraints that may prevent some agents from directly negotiating with each other. To this end, in the next section we discuss how these constraints are modelled.

3.1 The Social Network

Formally, we assume that agents are connected by a network $G = (I, E)$ where I is the set of vertices (agents) and E is a set of undirected edges. For example, an edge $\epsilon = (1, 2) \in E$ between agents 1 and 2 means that these agents can negotiate and form a coalition together. Note that agents that are not peers can be part of the same coalition, *as long as there is a path connecting the agents in the coalition*. However, those agents in the coalition that are not directly connected cannot directly engage in negotiation, and so an agent with more peers within the coalition would typically have more negotiating power to request a larger payoff.

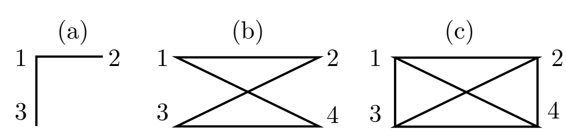

Figure 1: Simple social networks.

Consider the example social networks in Figure 1. In Figure 1(a), the possible coalitions are $\{1\}, \{2\}, \{3\}, \{1, 2\}, \{1, 3\}, \{1, 2, 3\}$. Thus, $\{2, 3\}$ cannot exist because there is no direct link between 2 and 3 but $\{1, 2, 3\}$ can exist because there is a common node $\{1\}$. Similarly, in Figure 1(b), $\{1,3\}$ and $\{2,4\}$ cannot exist. Finally, in the fully connected graph, Figure 1(c), all subsets are possible. Now, given the social network, let $N(i)$ be the set of neighbours of i. That is, $N(i) = \{i' | (i', i) \in E, i' \neq i\}$. Now, suppose an agent i invites agent $j \in N(i)$ to join a coalition $C \subseteq N(i) \cup \{i\} \setminus \{j\}$, to perform action α_j in $(\boldsymbol{\alpha}_C \cup \alpha_j)$ for a share v_j in $V(\boldsymbol{\alpha}_C)$. From this information agent j can compute its utility (as per equation (2)) and can then negotiate for a better share of the payoff $v'_j > v_j$ within the coalition. Note that we assume here that agent i knows the domain of actions D_j of its neighbour $j \in N(i)$ (otherwise, the agents would not know what to negotiate over). Next, we detail the negotiation process.

3.2 The Negotiation Protocol

The agents negotiate to form coalitions, establish the joint actions, and agree on a share of the surplus. This problem is complex as, in practice, multiple agents can interact simultaneously and asynchronously. While related works have suggested variants of Rubinstein's bargaining protocol to achieve similar goals [10, 6], they focus on equilibrium outcomes and do not consider many practical issues that arise in designing a system of negotiating agents for DCF (such as the rules dictating the sequence of exchanges between agents). To address these, here we identify key desiderata that a DCF protocol needs to satisfy and define a novel negotiation protocol that does indeed meet our desiderata. We start by defining the offers that agents can propose.

We denote a *simple* offer from agent i to agent j to join a coalition $C, j \notin C$ as the pair $o_{ij} = (\boldsymbol{\alpha}_C \cup \alpha_j, v_j)$, which specifies the actions of the agents in the coalition C, the action that agent j should perform, and the share of the payoff (in absolute terms) received by agent j. Since the proposing agent does not know the utility function of the other agent (because the costs of all its actions are private information), we allow agents to send a *compound* offer, O_{ij} containing multiple simple offers. For example, agent i may propose a different share for each action $\alpha_j \in D_j$. Note that the share v_i of i is not revealed in the offer to j but i privately knows it as well as its cost. Hence, in the rest of the paper, we note the utility of i in an offer to j as $\hat{u}_i(o_{ij}) = u_i(\boldsymbol{\alpha}_C \cup \alpha_j, \hat{v}_i)$ where $o_{ij} = (\boldsymbol{\alpha}_C \cup \alpha_j, v_j)$ and $\hat{v}_i$ is the privately known share of i in $V(\boldsymbol{\alpha}_C \cup \alpha_j)$. Moreover, to help in selecting the best agent to send compound offers to, we overload $\hat{u}_i(\cdot)$ to compute the maximum utility obtainable from O_{ij} as $\hat{u}_i(O_{ij}) = \max_{o \in O_{ij}} \hat{u}_i(o)$. Now, given the offer, the receiving agent may decide to propose an offer (i.e., invite an agent), counter-propose an existing offer, accept, or reject. Thus, agents can take the following illocutionary actions: $Propose(O_{ij})$, $Counter(o_{ji})$, $Accept(o_{ji})$, or $Reject()$. Alternatively, an agent can simply *do nothing* or wait to receive a response (we elaborate on this in section 4). Note that, when accepting or counter offering, the agent must specify a (simple) offer (instead of a compound one) from the set of received offers, $o_{ji} = (\boldsymbol{\alpha}_C, v_i) \in O_{ji}$. Moreover, the agent can also accept an *enhanced* offer (see Section 4.1 for more details), $o'_{ji} = (\boldsymbol{\alpha}'_{C'}, v'_i)$ that merges other coalitions with C in an attempt to improve i's payoffs (and as a consequence improves social welfare). To ensure that the enhanced offer does not make agent j or any other agent in coalition $C \setminus \{i\}$ worse off, it has to satisfy two conditions. First, $V(\boldsymbol{\alpha}_{C'}) - v'_i \geq V(\boldsymbol{\alpha}_C) - v_i$ (i.e., the total share received by agents in $C \setminus \{i\}$ cannot decrease). Second, $\boldsymbol{\alpha}_C \subset \boldsymbol{\alpha}'_{C'}$ (i.e., it can only add agents and the other agents' actions remain the same). Note that, due to the assumption given by equation (1), such an offer guarantees to be no worse in terms of utility for all agents in C.

Now, the negotiation protocol specifies a number of rules which agents must follow (and which should be enforced by the system) to ensure that the offers and agreements that are made can be honoured, and that no inconsistencies arise within the network (e.g., where two different coalitions share some agents). To date, however, negotiation protocols have mainly been designed in order to avoid deadlocks and ensure termination [12]. But in the DCF domain, we expect the negotiation protocol to *also satisfy* the following desiderata: (i) enable the formation of large coalitions (of size ≥ 2) to try and improve social welfare by exploiting synergies among agents, and (ii) minimise the information exchange needed across the network to try and minimise communication overheads (to improve scalability) and the risk of agents losing their negotiation power (to improve participation in the process). For example, if each agent is expected to broadcast the coalition it would like to form (assuming the network can handle this) and the individual payoffs the members (including itself) should get, it risks getting a smaller payoff (than when contacting peers) since all other agents may strategically use this information to demand a larger payoff. Hence, here we present the first DCF protocol specifically designed to meet all these desiderata:

Rule 1: Offer Commitment. Once an agent i proposes an offer O_{ij}, and agent j accepts, then i is committed to any $o_{ij} = (\boldsymbol{\alpha}_C, v_j) \in O_{ij}$ and any enhanced offer $o'_{ij} = (\boldsymbol{\alpha}'_C, v'_j)$ where $\boldsymbol{\alpha}_C \subset \boldsymbol{\alpha}'_C$ and $V(\boldsymbol{\alpha}_{C'}) - v'_i \geq V(\boldsymbol{\alpha}_C) - v_i$. If instead the recipient j counter offers $(\boldsymbol{\alpha}_C, v'_j)$ to i, then i is automatically decommitted from its previous offer. Hence, only the latest offer between the two agents counts. Moreover, if agent j rejects i's offer, i is no longer committed to its offer.

Rule 2: Proposal Consistency. Agent i can propose an offer $o_{ij} = (\boldsymbol{\alpha}_C \cup \alpha_j, v_j)$ only if i has either received an offer containing $\boldsymbol{\alpha}_C$, or has received offers which, when merged (to create an enhanced offer as desired), result in $\boldsymbol{\alpha}_C$. This prevents agents from involving other agents in a coalition they do not agree to. Note that merging offers is consistent with rule 1 since agents are committed to these offers.

If offers are allowed to be made by all agents at all times, the system is not likely to converge to a solution. To guarantee convergence (e.g., avoid deadlocks and infeasible solutions), the following rules (graphically illustrated in Figure 2) impose restrictions messages sent the system:[1]

Rule 3: Message Sending. All agents start in the *Ready* state (see Figure 2). *Ready* agents can propose an offer or counter offer, but to a *single* peer at a time (this limits communication overheads as desired). Once it has sent the message, it goes to the *Wait* state where it waits for a reply (either an accept offer, a counter offer, or a reject) before sending another message. Furthermore, a waiting agent can-

[1]Some of these rules are based on those that help achieve convergence on graphical models [1]. Due to lack of space, a proof of correctness and termination is not included in this version of the paper.

not receive any messages except from the agent it proposed the offer to (in practice, this can be achieved by making the state of the agent public). This rule prevents deadlocks, by preventing cycles of offers, and conflicting coalitions from being formed (i.e., where one coalition is not a strict subset of another) at the same time since an agent can only commit to form one coalition at a time. Note that, instead of proposing an offer, an agent can choose to remain *Ready*. As shown in Figure 2, it remains *Ready* even if it has received an offer. This is important as it allows an agent to receive multiple offers and merge them to form a bigger coalition (discussed in Section 4).

Rule 4: Negotiation Termination. Once an agent has accepted an offer or received an accept, it moves to the *Done* state. At this point, it has to send a message to all its peers from which it has received an offer (and who are, therefore, in the *Wait* state) informing them of the outcome. If the accepted offer subsumes offers received by neighbouring agents, it sends an accept to these agents, thereby propagating the accepted offer through the network. To all other agents, it sends a reject message. By doing so, these agents are taken out of their *Wait* state, allowing them to negotiate with the remaining agents. In addition, to ensure termination, we assume a deadline after which the agents that have not reached an agreement each form a singleton coalition. Given the above rules, we next describe ways to generate meaningful offers. Note that, contrary to most work on CF, we assume that an agent does not know other agents' utility functions, nor does it have time to form models of these functions. Hence, it is difficult for an individual agent to calculate equilibrium offers. Moreover, given that the computation of stable solutions is usually NP-Hard, we turn to bargaining heuristics that have typically been applied in bilateral negotiation [7]. Hence, by extending such heuristics to our DCF protocol, we provide a new framework to analyse the properties of the DCF process in a practical setting.

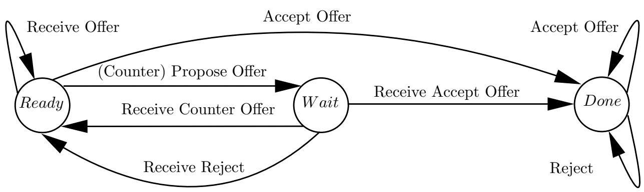

Figure 2: The negotiation protocol.

4. NEGOTIATION DECISION FUNCTIONS

Given the rules from the previous section, an agent can make a number of decisions. We separate the decision making into two parts. The negotiation *strategies* (discussed later in Section 4.2), determine the share, v_i, for any new offers or counter offers as well as the timing of such offers. Given the share, the *offer generation procedure* then chooses whether to accept an offer or produce a set of (counter) offers, and which (enhanced) offers to produce or accept. We start with the latter part, which is largely an optimisation problem that considers all the choices that respect our protocol.

4.1 Offer Generation Procedure

We let $\mathcal{O}_i$ denote the set of current *standing offers* received by agent i, i.e., offers that have not been rejected or counter offered. For convenience, we include the singleton coalitions which we denote by $O_{ii} = \{(\langle\alpha_i\rangle, V(\alpha_i))|\alpha_i \in D_i\}$. If agent i has not received any offers from any of its neighbours, then $\mathcal{O}_i = O_{ii}$. Otherwise, each offer $O_{ji} \subseteq \mathcal{O}_i$ from an agent $j \in N(i)$ to agent i consists of a number of simple offers. As we will see, agent j can offer a different share and actions $\boldsymbol{\alpha}_C$ for each of agent i's actions $\alpha_i \in D_i$ (it is possible to include even more combinations but this makes the offer impractically large). Therefore, $|O_{ji}| \leq |D_i|$.

Now, we can use Rule 1 from section 3.2 to expand the set of received offers by including all the *enhanced* offers that can be obtained by merging two or more offers. When doing so, agent i can claim any surplus resulting from the merge (while keeping the share of the other agents the same, as required by Rule 1). Consider an example where agent 1 receives $O_{21} = \{(\boldsymbol{\alpha}_C = \langle\alpha_1^1, \alpha_2^3\rangle, v_1 = 0.2), (\boldsymbol{\alpha}_C = \langle\alpha_1^2, \alpha_2^5\rangle, v_1 = 0.3)\}$ from agent 2, and $O_{31} = \{(\boldsymbol{\alpha}_C = \langle\alpha_1^1, \alpha_3^2\rangle, v_1 = 0.15), (\boldsymbol{\alpha}_C = \langle\alpha_1^2, \alpha_3^4\rangle, v_1 = 0.4)\}$ from agent 3. Then, agents 2 and 3 are also committed to the merged offers $O'_{21} = \{(\boldsymbol{\alpha}_C = \langle\alpha_1^1, \alpha_2^3, \alpha_3^2\rangle, v'_1), (\boldsymbol{\alpha}_C = \langle\alpha_1^2, \alpha_2^5, \alpha_3^4\rangle, v''_1)\}$, where v'_1 and v''_1 are maximised subject to the commitments to other agents in the coalition as set out in Rule 1:

$$v'_1 = V(\alpha_1^1, \alpha_2^3, \alpha_3^2) - (V(\alpha_1^1, \alpha_2^3) - 0.2) - (V(\alpha_1^1, \alpha_3^2) - 0.15)$$
$$v''_1 = V(\alpha_1^2, \alpha_2^5, \alpha_3^4) - (V(\alpha_1^2, \alpha_2^5) - 0.3) - (V(\alpha_1^2, \alpha_3^4) - 0.4)$$

Note that any received offer may have been merged with other offers by the sending agent, allowing the possibility to create large coalitions. We let $\mathcal{E}$ denote the expansion function, and $\mathcal{E}(\mathcal{O}_i)$ the set of all standing offers including all enhanced offers.

At this point, agent i has three options. It can either *accept* the best (enhanced) offer so far, $o^* = \arg\max_{o \in \mathcal{E}(\mathcal{O}_i)} u_i(o)$, and receive utility $u_i(o^*)$, *counter offer*, or produce a new set of offers to any of its peers $j \in N(i)$. The decision depends on which of these alternatives results in the highest utility. We now discuss the latter two options in turn.

4.1.1 Generating a Counter Offer

A *counter offer* is defined here as a standing offer where only the share claimed by the receiving agent is changed (thus the coalition and their actions remain the same). An agent i can choose to counter any of the offers $O_{ji} \subseteq \mathcal{O}_i, j \in N(i)$ received. Now, as mentioned earlier, the share is determined by the negotiation tactic (discussed in more detail in section 4.2). Let $v_i = T_i(\mathcal{H}_i, \boldsymbol{\alpha}_C, j, v_i^{min}, v_i^{max})$ denote the negotiation tactic which specifies the share requested by i as a function of the *history* $\mathcal{H}_i$ of offers and counter offers received by agent i (note that, while $\mathcal{O}_i$ only consists of the *standing* offers, $\mathcal{H}_i$ retains any offer made over time), $\boldsymbol{\alpha}_C$ is the proposed coalition (since this is a counter offer, this is the same coalition as in the original offer), j is the recipient of the offer, and v_i^{min}, v_i^{max} specify the *negotiation range*. Specifically, v_i^{min} is calculated such that the utility of the offers is at least equal to the utility it can achieve by accepting the best standing offer, $u_i(o^*)$. Formally:

$$v_i^{min} - \beta(\boldsymbol{\alpha}_C) = u_i(o^*) \Leftrightarrow v_i^{min} = u_i(o^*) + \beta(\boldsymbol{\alpha}_C) \quad (3)$$

Furthermore, $v_i^{max} = V(\boldsymbol{\alpha}_C)$ is the best i can hope to get. Note that, if $v_i^{max} < v_i^{min}$, $\boldsymbol{\alpha}_C$ will not be proposed as part of the counter offer, since no suitable offer can be made.

4.1.2 Generating a New Offer

Another option for the agent is to generate a new offer (if the agent received no offers, then this is the only option available). To establish *which* new offer to generate, agent i computes the best compound offer for each possible neighbour $j \in N(i)$, denoted by O^*_{ij}. This is done as follows. First, it calculates the expanded set of offers, excluding any offers received from agent j, $\mathcal{E}(\mathcal{O}_i \backslash O_{ji})$. Recall that this set contains all possible joined offers, as well as the singleton coalition with just agent i. Then, each of these offers is joined with an action $\alpha_j \in D_j$ of agent j, and the share to agent j, v_j, is computed using the negotiation tactic from Section 4.2. A compound offer then consists of the set of best simple offers, $o^*_{ij}(\alpha_j)$, one for each action $\alpha_j \in D_j$, excluding any offers that result in a utility below $u_i(o^*)$.

More formally, let $o' = J(\boldsymbol{\alpha}_C, \alpha_j, v_j)$ define the merged offer, where $o = (\boldsymbol{\alpha}_C, v_i), o' = (\boldsymbol{\alpha}_{C'}, v_j)$, and crucially, $\boldsymbol{\alpha}_{C'} = \boldsymbol{\alpha}_C \cup \alpha_j$. Furthermore, v_j is the share for agent j which is calculated based on i's share as follows: let $V(\boldsymbol{\alpha}_C) - v_i$ denote the total of shares committed to other agents in the C. Then, $v_i^{max} = V(\boldsymbol{\alpha}_{C'}) - (V(\boldsymbol{\alpha}_C) - v_i)$ is the maximum share that agent i can hope to get from the new coalition while $\hat{v}_i = T_i(\mathcal{H}_i, \boldsymbol{\alpha}_{C'}, j, v_i^{min}, v_i^{max})$ is i's privately known share as computed by the tactic. Given this, $v_j = v_i^{max} - \hat{v}_i$, is the share for agent j, and v_i^{min} is again calculated using Equation (3). Then, the best new offer is computed as:

$$o^*_{ij}(\alpha_j) = \underset{o'=J(\boldsymbol{\alpha}_C, \alpha_j, v_j), o \in \mathcal{E}(\mathcal{O}_i \backslash O_{ji})}{\arg\max} \hat{u}_i(o')$$

Furthermore, O^*_{ij} is given by: $O^*_{ij} = \{o^*_{ij}(\alpha_j) | \alpha_j \in D_j \wedge u_i(o^*_{ij}(\alpha_j)) > u_i(o^*)\}$.

Note that an important difference between a counter offer to j and a new offer, is that the latter removes any coalitional actions contained in the offer O_{ji}, and replaces this with a single action by agent j. Therefore, the new offer is not necessarily better in terms of utility than a counter offer.

4.1.3 Negotiation Decision

Having evaluated the received offers, computed counter offers, and new offers, an agent has to choose among them and send messages that respect our protocol. Algorithm 1 describes the steps to do so. Thus, in Step 1, the set of counter offers is generated based on the received offers (using tactic T_i) and in steps 2 and 3, the best new offer and counter offer are generated respectively. These are evaluated against the current best received offer o^* in step 4 and if o^* has the highest utility it is accepted, in which case *Accept* and *Reject* messages are sent to all agents involved (step 5) or not (step 6) in the offer respectively. Otherwise, the agent may send out a counter offer (step 8) or a new offer (step 10) depending on which gives it the best utility (step 7). We next provide the fundamental elements of the negotiation tactics and overarching strategies that are used to generate the share of an agent.

4.2 Negotiation Strategies

A negotiation strategy defines a tactic (i.e., T_i), that determines the share to be offered, and *when* an agent should make the offer.[2] We discuss these in turn. Since our goal is mainly to establish the baseline performance of our DCF framework, our tactic is based on well-known tactics in the automated negotiation literature, in particular [7]. In more detail, we use a family of exponential functions $f_\theta(r) = \frac{e^{\theta \cdot \frac{r}{r^{max}}}-1}{e^{\theta}-1}$ to determine the *concession rate*, where $r \in \{0, 1, \ldots, r^{max}\}$ is the negotiation round, r^{max} is the maximum round (after which no more concession occurs), and θ determines the slope of the curve. Note that r^{max} is not a deadline; negotiation can continue, but the agent will no longer concede. Also note that $f(0) = 0$ and $f(r^{max}) = 1$. Given this, the share is determined by:

$$T_i(\mathcal{H}_i, \boldsymbol{\alpha}_C, j, v_i^{min}, v_i^{max}) = v_i^{min} + (v_i^{max} - v_i^{min}) f_{\theta_i}(r) \quad (4)$$

We consider agents to be conciliatory if $\theta_i < 0$ (as this induces quick initial concessions), aggressive if $\theta_i > 0$ (as this induces slow initial concessions), and passive $\theta_i \simeq 0$ (as this induces concessions linear in the step size). The important point to note here is that we expect θ_i to depend on an agent's properties. For example, if an agent is well connected, it is in a good bargaining position and may choose to set $\theta_i > 0$.

Now, crucially, round r depends on the current offer made, the recipient of the offer, and the history of offers, $\mathcal{H}_i$. Specifically, in our framework, whenever an agent i proposes the same coalition to agent j as before with the same set of actions, and only the share changes, then that offer goes to the next round. If, however, a new offer is made which is not in $\mathcal{H}_i$, then a new counter is created and the round is set to zero. An agent can also first propose a new offer, and then in the next interaction propose an offer that has already been made before. In that case, the round for that offer is retrieved from memory and negotiations continue as before. Maintaining a counter for each unique offer is important because, this way, the *offer generation procedure* will automatically try different types of coalitions. To see this, note that, as a particular coalition is negotiated for several rounds, the agent's utility will decrease (since the agent concedes on that particular coalition), and therefore other options start to become more attractive. When these options

Algorithm 1 Accepting and Counter-offering.

Require: $o^*, O_{ji}, O^*_{ij} \forall j \in N(i)$

1: $O_{counter} = \{(\alpha_C, v_i') \mid o = (\alpha_C, v_i) \in O_{ji}, j \in N(i)\}$ where $v_i' = T_i(\mathcal{H}_i, \boldsymbol{\alpha}_{C'}, j, v_i^{min}, v_i^{max})$.
2: $O^*_{new} = \arg\max_{O \in \{O^*_{ij} | j \in N(i)\}} (\hat{u}_i(O))$
3: $o^*_{counter} = \arg\max_{o \in O_{counter}} (u_i(o))$.
4: **if** $u_i(o^*) \geq \max(u_i(o^*_{counter}), \hat{u}_i(O^*_{new}))$ `//accept offer` **then**
5:     send *Accept*(o^*) to all $j \in C, N(i)$ where $o^* = (\boldsymbol{\alpha}_C, v_i)$.
6:     send *Reject* to all $k \in N(i)$, where $\exists o_{ki} \in O_{ji}$ and $k \notin C$.
7: **else if** $u_i(o^*_{counter}) > \hat{u}_i(O^*_{new})$ **then**
8:     send $Counter(o^*_{counter})$ to agent j `//counter offer.`
9: **else**
10:     send $Propose(O^*_{new})$ to agent j `//new offer.`
11: **end if**

[2]More elaborate strategies could be designed (and part of future work) but we focus on two key features applicable to the protocol.

are exhausted, the agent may decide to go back to a previous offer until an agreement (or the deadline) is reached.

Having provided baseline tactics, we now discuss *when* an agent should make an offer and hence completely define our novel negotiation strategies. From the negotiation decision functions discussed so far, it is clear that an agent can benefit by *doing nothing* and waiting to receive more offers, since this gives an agent more options to merge coalitions and receive a higher payoff. However, if everyone waits, then no offer is ever made. Therefore, we propose two functions to decide the timing of offers, the second of which is specifically designed for our coalition formation protocol:

• **Time-based Strategy(TBS)**— an agent makes offers as soon as the DCF process starts. This is a typical bilateral negotiation strategy and simulates baseline performance for tactics that could be implemented in our framework.
• **Amortised Strategy (AS)**— similar to TBS, but the agent *delays its offers* by τ_i time steps where τ_i is drawn from $U(1, N(i))$. By amortising messages (i.e., staying in the *Ready* state), rather than readily making offers (i.e., going to *Wait*), agents (with more peers) can expect to receive more offers from their peers such that more options become available for merging coalitions to gain a larger share.

We next empirically evaluate TBS and AS on a variety of networks and determine the key factors that impact on the size and utility of coalitions that can be formed using our protocol and decision functions, as well as the efficiency of our strategies in creating such coalitions in a reasonable time. By so doing, we provide the first benchmarks for practical DCF via automated negotiation.

5. EVALUATION

Our aim here is to validate our approach by showing that such strategies allow agents to form coalitions efficiently and effectively. Here, efficiency is measured by the time taken to reach agreements while its effectiveness is quantified by the social welfare (i.e., aggregate utility) of the system. In what follows, we describe the experimental setup. Then we postulate and (in) validate a number of hypotheses with regards to the mechanism's and the agents' performance.

5.1 Experimental Setup

Each agent is randomly attributed between one and ten actions (i.e., $1 \leq |D_i| \leq 10$) and the cost for each action α_i, $\beta_i(\boldsymbol{\alpha}_C \cup \alpha_i)$ is drawn from a uniform distribution $\frac{U(10,100)}{|C|+1}$ (to simulate synergies in C). Coalition values are chosen to mimic decreasing marginal returns from bigger coalitions on average. The following function is used to achieve this by choosing parameters $p, q \in \mathbb{Z}^+$ and randomly picking $\mu^\alpha \in [1, 10]$ for each set of actions: $V(\boldsymbol{\alpha}_C) = e^{(p\times(1-\frac{1}{|C|})^q} \times \sum_{\alpha\in\boldsymbol{\alpha}_C} U(0,1) \times \mu_i^\alpha$.

We evaluate the performance of our strategies on the following networks (similar results were obtained when different parameters were used): (i) Scale free network (SFN) with 100 agents to simulate typical human societies — where the connectivity of agents follows a power law distribution (the maximum degree of each agent is limited to 20 and the average degree is 10), (ii) Random tree (RTN) with 100 agents with maximum degree 20 and the average degree is 2 to simulate hierarchical organisations, and (iii) Fully-connected network (FCN) with 20 agents and the average degree is 19 to simulate a single group. We limit the number of agents in FCN as the high degree of agents in these graphs result an exponential growth in the memory requirements and hence could not be run on a single machine. We experiment with two homogeneous populations of TBS and AS agents (as described in the previous section) with θ_i drawn from $U(-10, 10)$. Furthermore, we set $r_{max} = 10$.

The negotiation deadline is set to 100 time steps and at each time step agents make offers and counter offers to each other until agreements are reached or the deadline for the process is met. To evaluate the performance of our negotiation mechanism, we recorded the number of agreements reached, the time taken to reach such agreements, and the size of coalitions created. Our experiments are repeated 100 times and the results are averaged and the 95% confidence intervals are provided to indicate the statistical significance of the results where needed (equivalent to a t-test with $alpha = 0.05$).

5.2 Coalition Formation Efficiency

Here we evaluate whether the mechanism allows agents to rapidly converge to non-singleton coalitions under different settings. We postulate the following hypotheses:

H 1. *The higher the average degree of the network, the bigger the coalitions formed.*

H 2. *For all networks, AS takes longer to reach agreements than TBS.*

The intuition behind our hypotheses is as follows. We note that, in order to generate large coalitions, RTN and SFN require agents to form a chain of agreements (see Section 4). Instead, in FCN, all agents can directly contact *all* other agents and exchange offers over larger coalitions than they can in sparse networks. Moreover, since AS introduces artificial delays in the negotiation process, we expect AS to be slower.

Our expectations are partially met by the results (see table 1a) as the average size of coalitions in FCN (5.63) is bigger than in RTN (4.37) and SFN (3.62) when TBS is used (hence H1 is partially validated). However, when AS is used, the maximum coalition size is found in RTN (7.40) followed by FCN (5.73) and SFN (3.98). We explain the larger coalitions generated by AS in RTN by the fact that the low average degree of RTN ($2 < 10 < 19$ for RTN, SFN, and FCN respectively) suggests a large number of 'leaf' nodes which make offers earlier than their parent nodes who, in turn, merge received offers to form larger coalitions. The higher degrees of SFN and FCN instead reduce the likelihood of this effect.

Turning to the time taken (see Table 1b) to reach agreements (measured in number of rounds of offers), TBS is seen to take less time (2.67,5.02) than AS (11.32,8.28) in FCN and SFN respectively while, contrary to our expectations, AS takes less time (5.09) than TBS (10.19) in RTN even though AS resulted in the biggest coalitions. These results therefore invalidate H2. On close inspection, this is because, in FCN, agents can evaluate offers for all other agents in one go (since each agent has everyone else as a neighbour), but AS introduces some random delays. SFN presents a similar marked difference between TBS and AS due to delays incurred by AS and the fact that the agents cannot contact all other agents in one go, leading to longer negotiation times than FCN. In RTN, AS outperforms TBS as parent AS agents can assemble large coalitions from leaf agents by delaying their offers while TBS agents tend to propagate offers up the tree and delay the process significantly.

Network	TBS	AS
SFN	3.62±0.05	3.98±0.06
RTN	4.37±0.37	**7.40±0.69**
FCN	**5.63±0.19**	5.73±0.28

(a) Mean size of coalition.

Network	TBS	AS
SFN	5.02±0.29	8.28±0.21
RTN	**10.19±1.23**	**5.09±0.76**
FCN	2.67±0.29	11.32±0.31

(b) Mean time to agreement.

Table 1: Efficiency of CF process.

5.3 Coalition Formation Effectiveness

To compare the average utility AS and TBS obtain in different networks on an equal par, we computed an upper bound on the total utility (i.e., $\sum_{C\in S}\sum_{i\in C} u_i(C)$ where S is a partition of I) of all coalitions in any network using the value of the grand coalition minus the cost of the cheapest actions of all agents (i.e., $V(\boldsymbol{\alpha}_I) - \sum_{i\in I}\min_{\alpha_i\in D_i}(\beta_i(\alpha_i))$). This is to avoid computing the actual optimal coalition structure to form (and actions chosen) which is an NP-Hard problem and hence computationally infeasible for 100 agents in our case. Moreover, we computed the probability of reaching agreements as the ratio number of successful negotiation outcomes to the total number of negotiations started. We then postulate the following hypotheses.

H 3. *The probability of reaching an agreement is higher in networks with higher average degree.*

H 4. *AS always performs better (in utility generated and probability of reaching agreements) than TBS.*

The intuition behind these hypotheses is that in high degree networks, more coalitions can be explored, hence resulting in a larger negotiation space with, possibly, higher maxima. Moreover, since AS agents do actually consider more offers made to them, they also consider a larger negotiation space than TBS when choosing to accept. The results shown in tables 2a and 2b validate these hypotheses as AS outperforms TBS on all network topologies as the probability of reaching agreements and the utility of these agreements are both directly correlated with the average degree of the networks.

Network	TBS	AS
SFN	0.92±0.01	0.93±0.01
RTN	0.49±0.02	0.74±0.02
FCN	**0.94±0.01**	**0.97±0.01**

(a) Agreement probability.

Network	TBS	AS
SFN	0.48±0.01	0.52±0.01
RTN	0.33±0.02	0.43±0.01
FCN	**0.56±0.01**	**0.60±0.02**

(b) Scaled average total utility.

Table 2: Effectiveness of the CF process.

6. CONCLUSIONS

We have developed a novel framework based on heuristic negotiation to support practical decentralised coalition formation in social networks. By so doing, we establish a novel paradigm that addresses the design of DCF protocols and negotiation decision functions and, hence, departs from previous algorithmic and economic approaches. Moreover, we instantiated novel strategies for this framework and empirically evaluated them over general social networks in order to establish benchmarks. Thus, we showed how our amortised strategy outperforms the baseline in terms of the social welfare (by up to 10%) on all networks (trading off time to reach agreements) and also showed that the average degree of the network can significantly impact the performance of strategies (lower degrees can improve delay strategies). Future research will study the analytical properties of the protocol and improve our negotiation strategies.

7. REFERENCES

[1] S. Aji and R. McEliece. The generalized distributive law. *IEEE Trans. on Information Theory*, 46(2):325–343, 2000.

[2] G. Chalkiadakis, E. Markakis, and C. Boutilier. Coalition formation under uncertainty: Bargaining equilibria and the bayesian core stability concept. In *Proc. 6th Intl. Joint Conf. on AAMAS*, 2007.

[3] V. Conitzer and T. Sandholm. Computing shapley values, manipulating value division schemes, and checking core membership in multi-issue domains. In *AAAI*, pages 219–225, 2004.

[4] M. Corominas-Bosch. Bargaining in a network of buyers and sellers. *Journal of Economic Theory*, 115(1):35 – 77, 2004.

[5] V. D. Dang, R. K. Dash, A. Rogers, and N. R. Jennings. Overlapping coalition formation for efficient data fusion in multi-sensor networks. In *Proceedings of The Twenty First National Conference on Artificial Intelligence (AAAI-06)*, pages 635–640, 2006.

[6] G. Demange and M. Wooders. *Group formation in economics: networks, clubs and coalitions*. Cambridge Univ Pr, 2005.

[7] P. Faratin, C. Sierra, and N. R. Jennings. Using similarity criteria to make issue trade-offs in automated negotiations. *AIJ*, 142(2):205–237, 2002.

[8] M. Klusch and A. Gerber. Dynamic coalition formation among rational agents. *IEEE Intelligent Systems*, pages 42–47, 2002.

[9] R. Myerson. *Game theory: analysis of conflict*. Harvard Univ Pr, 1997.

[10] M. Osborne and A. Rubinstein. *Bargaining and markets*. Academic Press, 1990.

[11] M. Perry and P. Reny. A noncooperative view of

coalition formation and the core. *Econometrica: Journal of the Econometric Society*, pages 795–817, 1994.
[12] I. Rahwan, S. Ramchurn, N. Jennings, P. Mcburney, S. Parsons, and L. Sonenberg. Argumentation-based negotiation. *The Knowledge Engineering Review*, 18(04):343–375, 2003.
[13] T. W. Sandholm, K. Larson, M. Andersson, O. Shehory, and F. Tohme. Coalition structure generation with worst case guarantees. *AIJ*, 111(1–2):209–238, 1999.
[14] R. Serrano and R. Vohra. Non-cooperative implementation of the core. *Social Choice and Welfare*, 14(4):513–525, 1997.
[15] R. Smith. The contract net protocol: High-level communication and control in a distributed problem solver. *Computers, IEEE Transactions on*, 100(12):1104–1113, 2006.
[16] A. Watts. A dynamic model of network formation. *Games and Economic Behavior*, 34(2):331 – 341, 2001.