Stable Collaboration Patterns of Self-Interested Agents in Iterative Request for Proposal Coalition Formation Environments

Carlos Merida-Campos and Steven Willmott

Universitat Politècnica de Catalunya, Software Department, E-08034 Barcelona, Spain
Email: {dmerida, steve}@lsi.upc.edu

Abstract: This paper explores a general model of economic exchange between heterogeneous agents representing firms, traders, or other socioeconomic entities, that self-organise into coalitions to fulfil specific tasks. In particular, the work addresses *coalition formation* problems in which many tasks are addressed to the same population over time in an iterative fashion. The purpose of the paper is to describe the necessary elements that lead the system to an equilibrium state and asses the impact of coalition size constraints on the type of collaboration patterns established between agents. By using a novel data mining technique called *collaboration graphs* it is possible to see that stable states can be reached using simple iterative protocols and that the number of stable states increases as the coalition size limit decreases.
Keywords: MAS, Coalition formation, Request for Proposal, distributed problem solving, Electronic markets

1. Introduction

In recent years, Multi Agent Systems (MAS) have emerged as a powerful analytical tool to shed light on phenomena from research areas such as sociology, economics and biology, among many others. Work in multi-agent systems aims not only to generate new understanding of complex phenomena in human societies but also to understand the properties which may emerge in entirely electronic environment. Progress has been particularly fast in electronic marketplaces in which trading systems are often already highly automated with autonomous agents working on behalf of human traders to react to market triggers.

In this context, the work presented here focuses specifically on the challenge of coalition formation in environments characterised by iterated sequences of many *request for proposal* (RFP from now on) events and in which agents have little or no a-priori knowledge of the skill sets of others. Such environments are common place in many markets -representing regular calls for tender, contract offers and other requests for tasks to be carried out which may require not only single providers but a consortia of organisations with compatible and complementary skills. Further, in such scenarios, while some information on other agents in the population is available this is unlikely to be complete or accurate. In its generic form, RFP is a protocol by which an entity submits a description of a requirement and different providers issue different proposals that satisfy the requirements with different degrees. The proposals (one or more) which best satisfy the requirements are those which are rewarded (with the contract to carry out the task and its associated payment). This mechanism was first studied for agent systems in [3, 4] and [5] where RFP requests are seen by the population over time and agents must regularly decide whether to stay in the coalition they are in, or to leave. In [6] , agents can further decide which coalitions they would like to join and whether to accept or reject new members. This paper shows results on the *equilibrium dynamics* of such iterative RFP environments where agents may form coalitions and as more tasks arrive make choices on how to adapt their coalition structures based on success in the market.

Equilibrium convergence is usually a desirable property of a system and its establishment depends upon the market properties. Further, in a given system, there may be more than one equilibrium and market properties can affect the number of equilibria present. In the Iterated RFP model studied, stability is reached when a society is partitioned into coalitions in such a way that no agent changes the coalition it is in, either because it is not desirable unilaterally for this agent, or because it is not desirable for the coalition the agent wants to move to. This form of stability corresponds to the classical *Nash* equilibrium concept in game theory. In this paper it is shown that assuming an individually rational and pure strategic behaviour among the agents (score maximization), at least one equilibrium state is reachable. It is also shown that the number of stable states is dependant on the coalition size constraints imposed in the system. More specifically, there is an inverse relationship between coalition size limit and number of possible stable states. These results are proven analytically as well as experimentally by observing collaboration patterns across different experiments.

Section 2 explains the concrete market mechanism under study as well as the agent based system designed to model the iterative RFP coalition formation method, section 3 analyses the equilibrium properties of the system, proving a set of propositions that are later verified experimentally in section 4, which also explains the *collaboration graph* analysis technique used to check the collaboration patterns established throughout the set of experiments. Finally section 5 presents conclusions.

2. A Model for the Iterative RFP System

RFP Systems are those where an individual or an entity submits an invitation (a *Request For Proposal*) for providers of a product or service to bid on the right to supply that product or service to the individual or entity that issued the call. Calls are an iterative processes that repeats periodically over time - issuing the same or different tasks at regular intervals. This way a certain individual that participated in a certain call will likely participate in the subsequent calls for similar tasks over and over again, and adapting to the variations in the requirements and in the environment in order to be more competitive. Further, depending on the market rules, both the winning coalition in each cycle and others may be rewarded in ranked order. Given the competitive nature of the process, proposals tend to be formed by many partners that complement each other.

2.1. Definitions

A *population* I consists of a finite number of n individuals or *agents*. Agents are able to form *coalitions* to address a certain *task* in a more efficient way. $\sigma_i = \{\sigma_{i1}, \sigma_{i2}, \dots, \sigma_{im}\}$ denotes a coalition of size $|\sigma_i| = m$, where σ_{ij} represents an agent from population I forming part of coalition σ_i. This model assumes that an agent can just belong to one coalition. Agents have heterogeneous capabilities, thus having different performance levels in different skills. A finite number of k *skills*, indexed from 1 to k is set for which each agent σ_{il} has a fixed value: $\sigma_{il} = \langle \sigma_{il}^1, \sigma_{il}^2, \dots, \sigma_{il}^k \rangle$. In this manner it is possible to define a continuum of possibilities between agents that are *specialised* in the performance of a certain skill being unskilled for the rest of them, and agents that are *versatile*, being averagely apt for the performance of all the skills defined. A Task T is specified by a set of k skill requirements: $T = \langle T^1, T^2, \dots, T^k \rangle$. Those requirements are modelled in the form of a number. Each skill T^i could be considered as a necessary subtask for performing task T. The number of skills for which there is some requirement is noted as: $|T| = \sharp T^i : T^i > 0$. In a coalition, skills of agents are aggregated in such a way that each agent gives the best of itself in a join effort to create a group as competitive as possible under the requirements of the Task. The coalition has a value in each skill representing the aggregated effort of its members:

$$\sigma_i^l = \max(\sigma_{ij}^l) : 1 \le j \le m \tag{1}$$

In this manner, the agent in a coalition which is the best fit for performing a certain subtask would be the one that performs it. The set of maximal values corresponding to the best value in the coalition for each skill is noted as: $\widehat{\sigma_i} = \langle \sigma_i^1, \sigma_i^2, \dots, \sigma_i^k \rangle$. The aggregated effort of agents in equation 1 is used to measure an score $scr(\sigma_x, T)$ that indicates how well the agents in coalition σ_x perform together for accomplishing a task specification T. The score of a coalition is computed as the scalar product between σ_i and T. Amongst many possible choices, this metric is chosen because it captures in a simple way the different importance of subtasks T^l, and the additive valuation of all the required skills. Moreover, the function is monotonically increasing in the size of the coalition, upper bounded (reflecting the fact that a maximum level of quality is desired and possible for any given task - a perfect proposal) and deterministic (i.e. there are no objective or random elements to the valuation - it is based purely on the competence of the coalition):

$$scr(\sigma_x, T) = \sum_{j=0}^{k} (\sigma_x^j * T^j) \tag{2}$$

2.2. Agent Actions and Strategies

Each player's strategic variables are its coalition choice to join σ_j and a set of agents in this coalition $\phi_{jk} : \{(\phi_{jk} \subset \sigma_j) \vee (\phi_{jk} = \emptyset)\}$ to eliminate. The possibility of optimisation that an agent has, responds to the change of value that certain agents can experiment when in their coalition they are outskilled by a new member and so they become redundant because they are no longer complementary with the other coalition members . The new membership together with the optimisation proposal is not accepted straightforward, it is evaluated by the members of the target coalition. Just those actions accepted by a majority (more than the half) of members in the affected coalition, are performed. An agent that is requested to take an action can submit a finite number of requests in an specific order: $\psi_i = \langle (\sigma_a, \phi_{ab}), ..., (\sigma_p, \phi_{pq}) \rangle$. Those actions are evaluated sequentially until one is accepted or none of them is. If none of its action proposals is accepted, the agent stays in the same coalition where it was. Each of these actions potentially changes the coalition structures in the society and with many agents taking such actions each turn leads to an evolution of state.

All the agents follow the same strategy. They are *score maximizers*, hence they aim to be in the most competitive coalition as possible, that is, in the coalition with highest score that accepts them. Each agent σ_{zi} constructs an ordered set of action proposals ψ_i when requested. ψ_i contains all the possible proposals (σ_x, ϕ_{xy}) such that one of the following conditions is satisfied:

$$scr(\sigma_x - \phi_{xy} + \sigma_{zi}, T) > scr(\sigma_z, T) \tag{3}$$

$$\begin{aligned} &scr(\sigma_x - \phi_{xy} + \sigma_{zi}, T) = scr(\sigma_z, T) \wedge \\ &|\{\sigma_x - \phi_{xy} + \sigma_{zi}\}| < |sigma_z| \end{aligned} \tag{4}$$

In words, an agent's action proposal set contain every proposal that either improves the score of the coalition the agent is in, or keeps the same score while reducing the size. Action proposals in a set, are partially ordered having proposal's score as a first ordering criterium, and minimal coalition size as a secondary one.

Agents have information on the score of their current coalition and the potential outcome in terms of score, of any action they want to consider (however they have no information on the evolution of coalitions after its decision, nor the individual skills of agents in other coalitions) .

The economic benefit of being higher ranked is not explicitly represented in this paper, however, in [6] it was shown that score maximising strategy lead to higher benefits than a payoff maximisation strategy.

2.3. RFP Iterative Model

At time 0, every agent is a coalition of just one element ($\sigma_i = \{\sigma_{ii}\}$). A task T is issued and a run of negotiation starts in which every agent, sequentially and following a random order, is asked about an action to take. Agents have no knowledge on the order in which they are requested for an action. The run ends when all agents have been requested for an action. The process last as many runs as necessary until it converges to an stable state. Stability is reached when in a complete run, no agent is willing to leave the coalition is in or none of its actions are accepted by the hosting coalition.

3. Equilibrium Properties of the Iterative RFP Model

In order to analyse the complexity of the coalition formation process, it is necessary to define certain aspects of coalitional structures of the RFP model: A *partition* $\sigma = \{\sigma_1, \sigma_2, \dots, \sigma_m\}$ of population I is an specification of m coalitions of agents. Each agent belongs to one, and only one coalition. The set of coalitions can thus be identified with the index set $M =< 1, 2, \dots, m >$. It is convenient to use the partial order $\leq$ defined in previous section over the set of coalitions, thus having a partially ordered partition $\sigma =< \sigma_1, \sigma_2, \dots, \sigma_m >$ as an element of the space M^n. The total number of partitions in a given population is finite. The concrete amount of partitions of a population of size n corresponds the the n'th Bell number[1] B_n. Those partitions define the *state space* X of the dynamic process of coalition formation. Coalition formation in RFP environments is a dynamic process. For each state σ in X , let $F_X(\sigma)$ represent the set of states achievable by a one-step *move* (one agent move from one coalition to another, replace another agent in a coalition or optimise the coalition it is in). The model is thus *stochastic* as there is uncertainty on which of the achievable states follow a certain state, it will depend on which agent is elected to make a move. Once the concrete agent chosen to move it is known, the concrete state transition is still not deterministic for every case, because although all agents use a pure strategy (see section 2.2), when their best score choice can be obtained by doing two or more different actions ($scr(\sigma_{xj'}, T) = scr(\sigma_{yj''}, T)$) they would decide randomly which one they would try first. Let a *useful member* of a coalition be an element σ_{il} such that $\sigma_{il} \in \widehat{\sigma_i}$ hence $scr(\sigma_i, T) > scr(\sigma_i - \sigma_{il}, T)$. The set of skills of a useful member σ_{il} which value is maximal across the values of the rest of members of the coalition, are called *contributing skills* and are noted as $\widehat{\sigma_{il}} = \{j | \forall \sigma_{iq} : q \neq l : \sigma^j_{il} > \sigma^j_{iq}\}$. Let a coalition σ_i be *leading* in a given partition set if there is no coalition σ_j in the same partition σ such that $\sigma_i < \sigma_j$. Note that there could be more than one *leading* coalition as coalitions are arranged in partial order (σ_k might exist such that $\sigma_i = \sigma_k$). A partition σ is *unstable* if $F_X(\sigma) \neq \{\sigma\}$. That happens when there is at least one individual who would be better off in a different coalition, and that coalition would be better off accepting him or when there is a coalition where one or more of its members is not a *useful member*. A partition is *stable* otherwise, and is noted as σ^*. It is possible to demonstrate the convergence of the system into an stable state σ^* assuming the specified *score-maximising* strategic profile of a fixed population of agents I, as well as a deterministic, upper bounded and monotonically increasing with the number of agents $|\sigma_x|$ valuation function:

PROPOSITION 1 (monotonically increasing value of σ_1).*The* leading *coalition (σ_1) never reduces its value (score).*

Having a monotonically increasing with the number of agents valuation function (equation 1) the value of σ_1 can only decrease by leaving out members or replacing members by less competent ones. This cannot happen, as in the one hand, score maximising strategic profile of agents does not permit the replacement nor the expel of any *useful member*, as it will decrease coalition's score. In the other hand, no agent σ_{1l} has an incentive to leave to a coalition $\sigma_i : i \neq 1$ unless $(\sigma_i + \sigma_{il}) > \sigma_1$, in which case, $(\sigma_i + \sigma_{il})$ would be indexed with value 1, hence increasing the score of previous σ_1. The only case in which an agent might leave σ_1 not to create a coalition that outperforms σ_1, is the case in which it is expelled because it is not a *useful member*, in which case, by definition of useful member, after its move, the score of σ_1 does not change. □

PROPOSITION 2 (maximal σ_1 with unconstrained size).*If coalition size is not constrained to a value smaller than $|T|$, σ_1 always reaches the maximum possible value in a finite number of runs.*

Let σ_1^* be the set of coalitions with maximum possible value and minimum number of members, formally:
$\sigma_1^* = \{\sigma_x | (\arg\max_{y:\sigma_{xy}} scr(\sigma_x, T) \wedge \arg\min_{y:\sigma_{xy}} |\sigma_x|)\}$. Given a fixed T, scr valuation function (equation 2) is maximised with a maximal $\widehat{\sigma_x}$, hence the solution set consists on the set of coalitions with the minimal set of agents with maximal values for each required skill (specialist agents): $\sigma_1^* = \{\sigma_x | (argmax_{y \in I} \sigma^i_{xy} : T^i > 0 \wedge \arg\min_{y:\sigma_{xy}} |\sigma_x|)\}$. This set is never empty and may have more than one coalition if there are different agents with the same value in their contributing skills of a coalition of σ_1^*.

A coalition with a maximal score is always reached in finite time, if it were not like that, there would exist a coalition σ_1 in state of equilibrium such that $scr(\sigma_1, T) < scr(\sigma_1', T) \wedge \sigma_1' \in \sigma_1^*$. This would imply that $\exists i : \sigma_1^i < \sigma_1^i{}'$. Let the agent with contributing skill i of σ_1' be called σ_{1x}'. This agent would have no incentive to stay in a different coalition than σ_1 as is the leading coalition. This agent will be accepted in the coalition as $max(\sigma^i_1, \sigma^i_{1x}') = \sigma^i_{1x}'$ hence $scr(\sigma_{1x}' + \sigma_1, T) > scr(\sigma_1, T)$.

The leading coalition in equilibrium has always minimal size. If it were not like that, a coalition σ_1 in equilibrium would exist, for which $scr(\sigma_1, T) = scr(\sigma_1', T)$ and $|\sigma_1| > |\sigma_1'|$. There are two possible reasons for a leading coalition to be oversized. One is that is contains non useful agents, in which case the coalition would not be stable as those agents would be expelled (see equation 4). The second reasons is that two or more of its members are contributing in less skills than those in coalitions of σ_1^*. That is $\exists \sigma_{xy} : x \neq 1 \wedge \sigma^i_{xy} = \widehat{\sigma^i_{1a}} \wedge \sigma^j_{xy} = \widehat{\sigma^j_{1b}} \wedge i \neq j \wedge a \neq b$. In which case σ_1 would not be stable as σ_{xy} would be willing to submit the following joining+replacement proposal: $< (\sigma_1, \{\sigma_{1a}, \sigma_{1b}\}) >$. This replacement would keep the score of σ_1 while reducing its size, hence it would be accepted and σ_1 would not be stable. □

THEOREM 1 (stability).*A stable partition σ^* always exist and is reached in a finite number of runs*

Given Proposition 2, σ_1 converges into a stable state in which the maximum score and minimum size is reached in finite time. Once this coalition has reached its optimal state, Proposition 1 would apply to σ_2, as this coalition score becomes bounded under the score of σ_1. Hence, σ_2 will reach a stable maximal score value. Taking this as step of the inference, the same argument can be applied for the rest of the coalitions so that none of them would change their value. This demonstrates that effectively exists at least one stable state and further, that the system must converge to one. □

PROPOSITION 3 (maximal σ_1 with constrained size).*If coalition size is constrained to a value smaller than $|T|$, σ_1 con-*

[1] This can be calculated by summing the Stirling numbers of the second kind for each number of partitions: $B_n = \sum_{k=1}^{n} S(n, k)$. As an example, the number of partitions of a population of 100 agents is: $4.7585 * 10^{115}$

verges into a stable state (not necessarily with the maximum possible value).

If the coalition is constrained to a limited size below $|T|$, there would be one or more agents with, necessarily, more than one contributing skill: $\forall i : (\exists l : \widehat{\sigma_{il}^{j}} \wedge \widehat{\sigma_{il}^{p}} \wedge j \neq p)$, this fact may create a deadlock in the following situation: If contributing skills for two agents σ_{1x} and σ_{1y} in the leading coalition, are set up in the following way: $\widehat{\sigma_{1x}^{a}}, \widehat{\sigma_{1x}^{b}}, \widehat{\sigma_{1y}^{c}}, \widehat{\sigma_{1y}^{d}}$, assuming that σ_{1x} and σ_{1y} are just contributing in the specified skills. If $|\sigma_1|$ equals the maximum permitted size. If there exists other two agents σ_{iw} and σ_{jz} such that if by doing the following simultaneous replacement: $\sigma_1' = \sigma_1 - \{\sigma_{1x}, \sigma_{1y}\} + \{\sigma_{iw}, \sigma_{jz}\}$ the score is improved $scr(\sigma_1') > scr(\sigma_1)$ and the resultant contributing skills for the new agents in the coalition are $\widehat{\sigma_{1w}^{a}}', \widehat{\sigma_{1z}^{b}}', \widehat{\sigma_{1w}^{c}}', \widehat{\sigma_{1z}^{d}}'$. Under specific circumstances, the previous replacement could not be possible when agents perform sequentially (as iterated RFP protocol imposes). ex: For all of the four possible replacements this situation could happen: $\sigma_1'' = \sigma_1 - \sigma_{1x} + \sigma_{iw}$, having contributing skills set up in the following manner: $\widehat{\sigma_{1w}^{a}}'', \widehat{\sigma_{1l}^{b}}'', \widehat{\sigma_{1y}^{c}}'', \widehat{\sigma_{1y}^{d}}''$ with $scr(\sigma_1'') < scr(\sigma_1)$, hence no sequential replacement would result benefitial and σ_1 would remain suboptimal.

The leading coalition would converge into a stable state because if no deadlock happens, explanation on convergence shown in proposition 2 holds. If coalition is deadlocked in a suboptimal state, the deadlock would prevent σ_1 for changing, hence the coalition would be stable. □

A direct consequence of this proposition is that in a competition where the size of coalitions is constrained to a value lower than the number of skills required, the number of possible stable states is increased by deadlocked states that would be stable although not necessarily optimal.

4. Experimental Analysis

This section reviews the analysed equilibrium dynamics properties from an experimental point of view and assesses the effect of coalition size constraints on the collaboration patterns established between agents. In order to observe the population's patterns of collaboration, the RFP protocol is used as many runs of negotiation as necessary (see section 2.2) until agents reach an *stable state*. At that point, the one to one relationship established between members of the population are saved, and after the final experiment, the frequency of these relationships are compiled across all runs and used to generate weights in a population graph. The data generated in this manner contains information on which agents often work together for the task set. The mechanism used to determine clustering is an implementation of *Kamada-Kawai* [2] algorithm in *Pajek* [1]. This algorithm places nodes in a bi-dimensional space in a close position when they are connected with links of relative high value. Graphs are created from data obtained by the clustering algorithm, and for the sake of the clarity, these also distinguish the frequency of cooperation between agents, not only by of their position in the graph, but also with the tone of the arrows in such a way that they look darker when the relationships are more frequent, and lighter otherwise. Throughout the rest of the paper, these graphs are referred to as *Collaboration Graphs*.

4.1. Experimental Set-Up

In order to validate the properties shown, a significant number of different experiment sets have been performed. Those parameters that have a scale effect on the system but do not change the global effects have been fixed. These parameters are: population size ($|I| = 100$), number of skills ($k = 10$), total value of every agent ($\sum_{j=1}^{k} \sigma_{xy}^{j} = 200$) and total value of every task used ($\sum_{j=1}^{k} T^{j} = 100$). The variables under study have been: *coalition size limit* (unlimited size, 8, 6 and 4) and the *distribution of value* (total value of 200 with an specific standard deviation in the way it is distributed across skills) in every agent (stdev equal to: 5,10 or 15). The different distributions among skills create different scenarios of *speciality/versatility* of agents.

4.2. Stability Properties

4.2.1. Convergence and Non Decreasing Value of Leading Coalition

For all the experiments performed, it has been confirmed the convergence of the system into an stable state. Following the defined notion of stability (see section 3). This has been proved monitoring the type of orders that agents submit during the process and observing how in every experiments agents end up not submitting any action proposal or submitting action proposals that are not accepted by the hosting coalition.

Figure 1 shows score values of all the coalitions in every iteration of an experiment. In the graphic it can be appreciated how *leading coalition* rapidly stabilises to a certain value. This coalition is the first one to become stable, as all the agents are willing to join the best coalition that exists, and rapidly, those agents that best complement each other, gather in the same group. Once this coalition has no possibility to increase its score, it rejects every joining order received, and so the next ranking becomes the best chance for the second best group of agents, that find themselves in the second ranking coalition until there is no possible improvement for it. This process repeats until the last ranked coalition is stable. This result illustrates the quick progression to stable states in every coalition from the top ranked to the worse ones as well as the non decreasing value of the leading coalition.

4.2.2. Different Values for Leading Coalition

In figure 2 it can be seen how the leading coalition does not register any variation in the score in the unconstrained coalition size setup across 150 different experiments, whereas in experiments with coalition size constraint equal to 8, 6 or 4, there is more than one value for the leading coalition, revealing the existence of deadlocked states that ended up in suboptimal outcomes in some stable patterns. The optimality insurance of not constraining coalition size is an important property of the system to be considered in order to implement this protocol. Optimality is in any case a desirable outcome for the requester of the proposal. However the fact that size constrained setups present changes in higher ranks, is not the only factor that determines the high number of differences of values for every rank. Next section explores the effect of coalition size constraints not only analysing the scores of coalitions in ev-

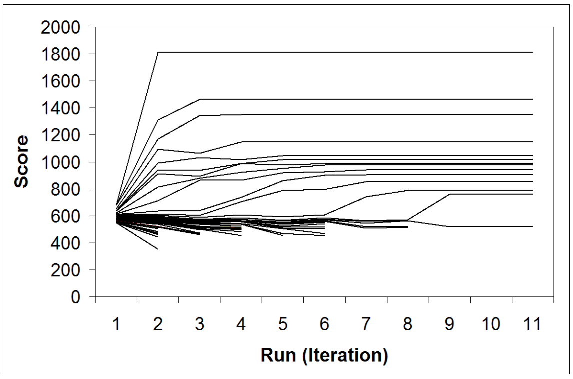

Fig. 1. Progression of score of coalitions during the convergence process.

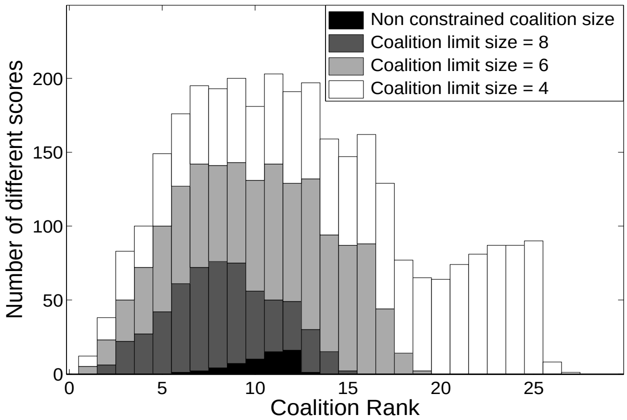

Fig. 2. Number of different values in the score of every coalition in every rank of 150 repeated experiments in equilibrium. Comparing different setups of coalition size constraints we see how constraining size involves more variability in scores.

ery rank but also the frequency of the structures constructed themselves.

4.3. Effect of Coalition Size Constraint on Collaboration Patterns

4.3.1. Collaboration Patterns in an Unconstrained Coalition Size Setup

Figure 3(a) shows the collaboration graph corresponding to 150 different experiments using the same agent population and task. The graph reflects stable relationships between agents across different runs of experiments, indicating that in spite of the stochastic nature of the process, agents have clear sets of preferred partners to form coalitions with. As it was shown in proposition 2, if coalition size is not constrained, the degree of specialisation of agents becomes the most important factor to determine the rank that an agent will occupy in the population. The only possible difference from one equilibrium structure to another is given by the probability that two or more agents have the same specialisation degree in a certain skill for which those agents provide the maximum value. Depending on that fact, more or less different stable patterns can exist.

(a) Collaboration graph corresponding to the unconstrained coalition size setup

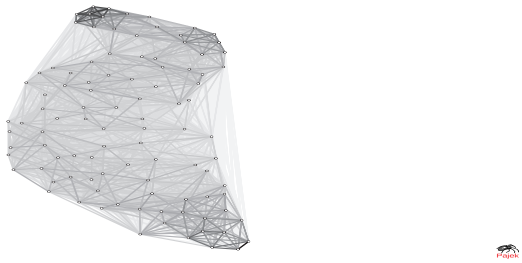

(b) Collaboration graph corresponding to coalition size limit = 8 setup

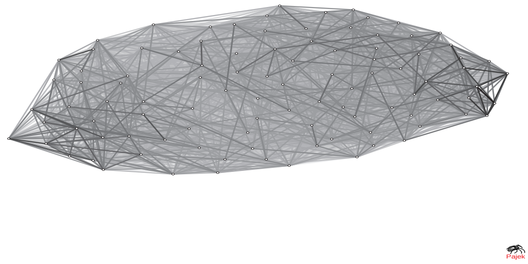

(c) Collaboration graph corresponding to coalition size limit = 6 setup

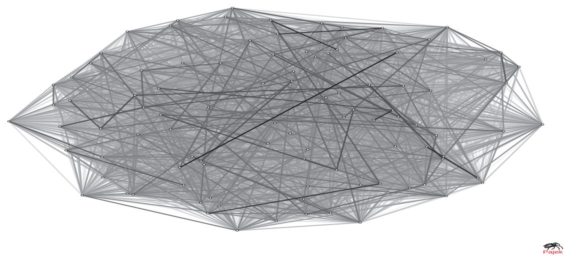

(d) Collaboration graph corresponding to coalition size limit = 4setup

Fig. 3. Collaboration graphs for the same task and the same agent's population accross 150 experiments for 4 different values of coalition size boundness.

4.3.2. Collaboration Patterns in a Constrained Coalition Size Setup

Figures 3(b), 3(c) and 3(d) show collaboration graphs corresponding to different coalition size constraint setup using the same agent population and task across 150 experiments. Figures show that a decreasing degree of clustering in the collaboration patterns corresponds to an increasing constraint in size. This effect is seen for three reasons. The first is the possibility of deadlock effect showed in proposition 3, that predicts an increasing number of stable states given to the possible coexistence of suboptimal and optimal states.The second is because in the bounded size scenario the ranking of coalitions is determined by a tradeoff between *specialisation* and *versatility* of agents abilities. If for the case of unconstrained size, the number of different stable states depended on the number of agents

Fig. 4. Inherent compatibilities of agent's popualtion showed in a collaboration graph for 150 different experiments using the same agent's population but a different task per experiment and no coalition size constraint.

that had the same specialisation degree, in the constrained case it depends on the number of combinations of agents for which their aggregated value of their skills is the same. The more constrained the coalition size is, the lower the values of coalitions. Hence the easier (in probability terms) to get equal aggregated values by forming groups of agents. Finally the third reason of the increment of stable states with the coalition size constraint is given by an increment on the number of possible partitions of the population when coalitions are limited to the maximum size permitted, as the Stirling number of the second kind $S(n, k)$ that counts the total number of partitions of coalitions of size k from a population of n agents, has the following property: $S(100, 10) < S(100, 8) < S(100, 6) < S(100, 4)$. The increment in the state space also increments the probability of possible stable states. The cloud structure of the constrained coalition size demonstrates the high number of stable states reflecting that almost all the agents are collaborating with any other agent in some stable state.

4.4. Stable States Across Many Varying Tasks

So far experiments assumed a fixed given task, and the same population conforming the same call during a set of experiments. By changing the task after an experiment has reached an stable state, the type of collaboration patterns obtained when coalition size in not constrained still shows a clear structure of preferential partnership. Figure 4 shows those structures. This demonstrates that the inherent compatibilities degree of agents are not heavily affected by the task requirements. The reason for this is that specialisation in a task is positively evaluated no matter what is the task requirement (as long as there is some requirement), as given the aggregation function used (equation 1) coalitions in every rank have the best agents available for every skill no matter the degree of requirement of the skill. Apart from showing that a population of agents has inherent compatibilities when the setup requires a high degree of specialisation, the obtention of those results reveals that the protocol used preserves those compatibilities that could be learned by agents and exploited in order to reduce the search across the space of possibilities every time an agent is requested to make a movement.

5. Conclusions and Future Work

As electronic trading mechanisms become more common place and increasingly feasible for complex environments, the study of the dynamics of such systems will be increasingly important. In this context, the work presented here addresses Request for Proposal style environments which iterate over time and allow agents to gradually adapt their coalition structures between calls. Specifically, results for score maximising agents show that despite limited information, equilibrium states can still be found in such protocols, that the number of stable states varies with coalition size constraints and that traces of the stable states can be seen not just in one task but across many diverse tasks. The paper shows the convergence of the system modelled as a Nash equilibrium when players use a pure score maximizing strategy, proving also that the score of the best coalition is monotonically increasing during the convergence process. The existence of Nash stable states is a desirable property for any system in general, and service composition mechanism in particular. For providers in a service composition process, the Nash stable configuration is a fair outcome, as given the configuration of the system, they would not be willing to change their situation. Moreover, we demonstrate that in such a stable state, the best coalition is optimal (maximum possible value) as long as the coalition size is not constrained, whereas when it is constrained, a suboptimal configuration may occur as a deadlock situation. Finally, we show how coalition size constraint determines the collaboration patterns established between agents. More specifically, size constraints control the importance given to the speciality or versatility facets of agent's heterogeneous capabilities and, depending on this tradeoff, radically different collaboration patterns are established.

Further research is required to analyse the imperfect/ incomplete information scenario by limiting the social awareness of agents in the population. Further research is also required to analyse the system properties when there is more than one simultaneous task requesting proposals to the same population. Those extensions will create more realistic assumptions for the iterative RFP environment in order to consider it as a valid e-commerce mechanism.

References

[1] V Batagelj and A Mrvar, Pajek-program for large network analysis, http://vlado.fmf.uni-lj.si/pub/networks/pajek/.

[2] T Kamada and S Kawai, An algorithm for drawing general undirected graphs. Inf. Process. Lett., Vol. 31, No. 1, 1989, pp. 7–15.

[3] S Kraus, O Shehory and G.Tasse, Coalition formation with uncertain heterogeneous information. Proc. 2nd Int. Joint Conf. On Autonomous Agents & Multi Agent Systems (AAMAS'03) 2003, pp. 1–8, Melbourne, Australia.

[4] S Kraus, O Shehory and G.Tasse, The advantages of compromising in coalition formation with incomplete information. Proc. 3rd. Int. Joint Conf. On Autonomous Agents & Multi Agent Systems (AAMAS'04) 2004, New York, USA.

[5] C Merida-Campos and S Willmott, Modelling coalition formation over time for iterative coalition games. Proc. 3rd Int. Joint Conf. On Autonomous Agents & Multi Agent Systems (AAMAS'04) 2004, pp. 572–579, New York, USA.

[6] C Merida-Campos and S Willmott, Agent compatibility and coalition formation: Investigating two interacting negotiation strategies. Agent Mediated Electronic Commerce VIII (TADA/AMEC). LNAI 4452 2007, pp. 75–90, Hakodate, Japan.