

Decision Making in Multi-agent Coalitions in a Dynamic Hostile World

Madhu Goyal
Faculty of Information Technology
University of Technology Sydney, Australia
madhu@it.uts.edu.au

Abstract

One of the main underpinning of the multi-agent systems community is how and why autonomous agents should cooperate with one another. Several formal and computational models of cooperative work or coalition are currently developed and used within multi-agent systems research. The coalition facilitates the achievement of cooperation among different agents. In this paper, a mental construct called *attitude* is proposed and its significance in coalition formation in a dynamic fire world is discussed. This paper presents ABCAS (Attitude Based Coalition Agent System) that shows coalitions in multi-agent systems are an effective way of dealing with the complexity of fire world. It shows that coalitions explore the attitudes and behaviours that help agents to achieve goals that cannot be achieved alone or to maximise net group utility.

Keywords: multi-agent, coalition formation, attitudes

1 Introduction

Coalition formation has been addressed in game theory for some time. However, game theoretic approaches are typically centralised and computationally infeasible. MAS researchers [7] [8] [9] [10] [11], using game theory concepts, have developed algorithms for coalition formation in MAS environments. However, many of them suffer from a number of important drawbacks like they are only applicable for small number of agents and not applicable to real world domains. A coalition, is a group of agents who join together to accomplish a task that requires joint task execution which otherwise be unable to perform or will perform poorly. It is becoming increasingly important as it increases the ability of agents to execute tasks and maximise their payoffs. Thus the automation of coalition formation will not only save considerable labour time, but also may be more effective at finding beneficial coalitions than human in complex settings. To allow agents to form coalitions, one should devise a coalition formation mechanism that includes a protocol as well as strategies to be implemented by the agents given the protocol.

This paper will focus on the issues of coalitions in dynamic multi-agent systems: specifically, on issues surrounding the formation of coalitions among possibly among heterogeneous group of agents, and on how coalitions adapt to change in dynamic settings. Traditionally, an agent with complete information can rationalise to form optimal coalitions with its neighbours for problem solving. However, in a noisy and dynamic environment where events occur rapidly, information cannot be relayed among the agent frequently enough, centralised updates and polling are expensive, and the supporting

infrastructure may partially fail, agents will be forced to form sub-optimal coalitions. Similarly, in such environments, changes in environmental dynamics may invalidate some of the reasons for the original existence of a coalition. In this case, individual agents may influence the objectives of coalition, encourage new members and reject others and the coalition as a whole adapts as a larger organism. In such settings, agents need to reason, with the primary objective of forming a successful coalition rather than an optimal one, and in influencing the coalition (or forming new coalitions) to suit its changing needs. This paper introduces ABCAS, a novel attitude based coalition agent system in the fire world. The task of fire fighting operations in a highly dynamic and hostile environment is a challenging problem. We suggest a knowledge-based approach to the coalition formation problem for fire fighting missions. Thus the objective of this paper is to design and develop an attitude based approach to the coalition formation for fire fighting problem that would help them to accomplish their tasks during the fire. Owing to the special nature of this domain, developing a protocol that enables agents to negotiate and form coalitions, and provide them with simple heuristics for choosing coalition partners is quite challenging task. The protocol allows the agents to form coalitions, and provide them with simple heuristics that allow the agents to form coalitions in face of time constraints and incomplete information.

2 A Fire Fighting World

We have implemented our formalisation on a simulation of fire world *FFWorld* [5] [6] using a virtual research campus. *FFWorld* is a dynamic, distributed, interactive, simulated fire environment where agents are working together to solve problems,

for example, rescuing victims and extinguishing fire. In a world such as this, no agent can have full knowledge of the whole world. Humans and animals in the fire world are modelled as autonomous and heterogeneous agents. While the animals run away from fire instinctively, the fire fighters can tackle and extinguish fire and the victims escape from fire in an intelligent fashion. An agent responds to fire at different levels. At the lower level, the agent burns like any object, such as chair. At the higher level, the agent reacts to fire by quickly performing actions, generating goals and achieving goals through plan execution. This world contains all the significant features of a dynamic environment and thus serves as a suitable domain for collaborating agents. Agents operating in the domain face a high level of uncertainty caused by the fire. Agents in the fire domain do not face the real time constraints as in other domains, where certain tasks have to be finished within the certain time. However, because of the hostile nature of the fire, there is strong motivation for an agent to complete a given goal as soon as possible.

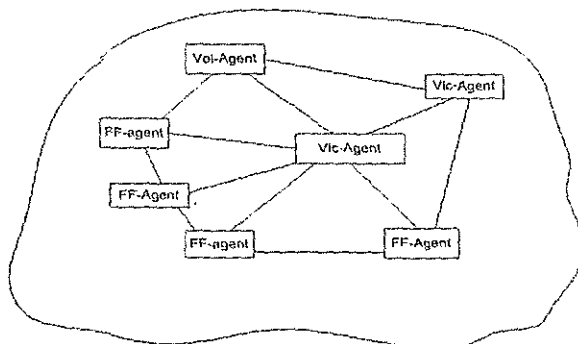


Figure 1: Coalition between fire-fighter, volunteer and victim agent.

There are three main objectives for intelligent agents in the world during the event of fire: self-survival, saving objects including lives of animals and other agents and put-off fire. Because of the hostile settings of the domain, there exist a lot of challenging situations where agents need to do the cooperative activities. Whenever there is fire, there is need of coalition between the fire fighters (FF-agent), volunteers (Vol-agent) and victim agents (Vic-agent) (figure 1). The fire fighters perform all the tasks necessary to control an emergency scene. The problem solving activities of the fire fighters are putting out fire, rescuing victims and saving property. Apart from these primary activities there are a number of sub tasks e.g. run towards the exit, move the objects out of the room, remove obstacles, and to prevent the spread of fire. The first and paramount objective of the victim agents is self-survival. The role of volunteer agents is to try to save objects from the fire and help out other victims who need assistance when they believe their lives are not under threat. To achieve these tasks there is need of coalitions between these agents is necessary. Thus the

fire world we consider is sufficiently complex to bring about the challenges involved in to study coalition formation in a typical multi-agent world.

3 Strategic Coalitions in an Agent Based Hostile World

The coalition facilitates the achievement of cooperation among different agents. The cooperation among agents succeeds only when participating agents are enthusiastically unified in pursuit of a common objective rather than individual agendas. We claim that cooperation among agents is achieved only if the agents have a *collective attitude* towards cooperative goal as well as towards cooperative plan. From collective attitudes, agents derive individual attitudes that are then used to guide their behaviours to achieve the coalition activity. The agents in a coalition can have different attitudes depending upon the type of the environment the agent occupies.

3.1 Definition of Attitude

Attitude is a *learned predisposition to respond in a consistently favourable or unfavourable manner with respect to a given object* [4]. In other words, the attitude is a preparation in advance of the actual response, constitutes an important determinant of the ensuing behaviour. However this definition seems too abstract for computational purposes. In AI, the fundamental notions to generate the desirable behaviours of the agents often include goals, beliefs, intentions, and commitments. Goal is a subset of states, and belief is a proposition that is held as true by an agent. Bratman [1] addresses the problem of defining the nature of intentions. Crucial to his argument is the subtle distinction between doing something intentionally and intending to do something. The former case might be phrased as deliberately doing an action, while intending to do something means one may not be performing the action in order to achieve it. Cohen and Levesque [3], on the other hand, developed a logic in which intention is defined. They define the notion of individual commitment as *persistent goal*, and an intention is defined to be a commitment to act in a certain mental state of *believing throughout what he is doing*. Thus to provide a definition of attitude that is concrete enough for computational purposes, we model attitude using goals, beliefs, intentions and commitments. From Fishbein's [4] definition it is clear that when an attitude is adopted, an agent has to exhibit an appropriate behaviour (predisposition means *behave in a particular way*). The exhibited behaviour is based on a number of factors. The most important factor is goal or several goals associated with the object. During problem solving, an agent in order to exhibit behaviour may have to select from one or several goals depending on the nature of the dynamic world.

In a dynamic multi-agent world, the behaviour is also based on appropriate commitment of the agent to all unexpected situations in the world including state changes, failures, and other agents' mental and physical behaviours. An agent intending to achieve a goal must first commit itself to the goal by assigning the necessary resources, and then carry out the commitment when the appropriate opportune comes. Second, if the agent is committed to executing its action, it needs to know how weak or strong the commitment is. If the commitment is weak, the agent may not want to expend too much of its resources in achieving the execution. The agent thus needs to know the degree of its commitment towards the action. This degree of commitment quantifies the agent's *attitude* towards the action execution. For example, if the agent considers the action execution to be *higher importance* (an attitude towards the action), then it may choose to execute the action with greater degree of commitment; otherwise, the agent may drop the action even when it had failed at the first time. Thus, in our formulation, an agent when it performs an activity, since the activity is more likely that it will not succeed in a dynamic world; agents will adopt a definite attitude towards every activity while performing that activity. The adopted attitude will guide the agent in responding to failure situations. Also the behaviour must be consistent over the period of time during which the agent is holding the attitude. Thus attitudes, once adopted, must persist for a reasonable period of time so that other agents can use it to predict the behaviour of the agent under consideration. An agent cannot thus afford to change its attitude towards a given object too often, because if it does, its behaviour will become somewhat like a reactive agent, and its attitude may not be useful to other agents. Once an agent chose to adopt an attitude, it strives to maintain this attitude, until it reaches a situation where the agent may choose to drop its current attitude towards the object and adopt a new attitude towards the same object. Thus we define attitude as: *An agent's attitude towards an object is its persistent degree of commitment to one or several goals associated with the object, which give rise to persistent favourable or unfavourable behaviour to do some physical or mental actions.*

3.2 Type of Attitudes

The attitudes of the agents in the world consist of attitudes towards the physical objects, mental objects, processes and other agents. When attitudes are attached to physical objects, the agents are able to evaluate the *liking*, *importance* or *location* etc. of these physical objects. When attitudes are attached to mental objects, agents are able to communicate and reason with those mental objects. For example, agents can actively monitor their plans so those plans can be re-organised or abandoned when the world state changes. If the object denotes a mental object such as a plan, *higher-priority* can be an attitude that the

agent may hold towards the plan. In that case, the agent will perform behaviour appropriate to this attitude, which may involve physical, communicative, and mental actions or a combination of these which may lead to behaviour where the agent gives higher preference to the plan compared to the other plans in all possible situations. Agents can also have attitudes towards processes such as execution of actions and plans, the process of achieving goals, etc. For example, if the execution of a plan goes on for too long, appropriate attitude is necessary to define how to handle the situation. Similarly, attitudes are also possible towards process states, such as the execution state of a plan after a crucial action has been executed. Similarly, attitudes towards other agents and their mental states define the agent's behaviour towards these agents.

Behaviours exhibited by an agent in a multi-agent environment can be either individualistic or collective. Accordingly, we can divide attitudes in two broad categories: *individual* attitudes and *collective* attitudes. The *individual* attitudes contribute towards the single agent's view towards an object or person. An agent's attitude toward an object is based on its salient beliefs about that object. The agent's individual attitude toward a fire world, for example, is a function of its beliefs about the fire world. The *collective* attitudes are those attitudes, which are held by multiple agents. The collective attitudes are individual attitudes so strongly interconditioned by collective contact that they become highly standardised and uniform within the group, team or society etc. The agents can collectively exist as *societies*, *groups*, *teams*, *friends*, *foes*, or just as *strangers*, and *collective* attitudes are possible in any one of these classifications. For example, the agents in the collection called *friends*, can all have a collective attitude called *friends*, which is mutually believed by all agents in the collection. A collective attitude can be viewed as an abstract attitude consisting of several component attitudes, and for an individual agent to perform an appropriate behaviour; it must hold its own attitude towards the collective attitude. Thus, for example, if A1 and A2 are *friends*, then they mutually believe they are friends, but also each A_i must have an attitude towards this infinite nesting of beliefs so that it can exhibit a corresponding behaviour. Thus, from A₁'s viewpoint, *friends* is an attitude that it is holding towards the collection {A₁, A₂} and can be denoted as *friends*_{A₁}(A₁, A₂). Similarly, from A₂'s view point, its attitude can be denoted as *friends*_{A₂}(A₁, A₂). However, in an extreme case, A1 may not be certain about A2's behaviour. That is why; it needs to have an attitude of its own, which generates a behaviour taking all the uncertainties introduced by the dynamic environment into account. Further note that A1 might implicitly expect A2 to perform its role, but it is only an expectation. This is a bottom up view of the friends-relation, where the relation is viewed in general

differently by each agent depending on the local situations the individual agents face in the world.

3.3 Attitude Based Agents

We adopt a BDA (figure 2 modified BDI architecture) based approach in which agent is comprised of: *beliefs* about itself, others and the environment; set of *desires* representing the states it wants to achieve; and *attitudes* corresponding to the plans adopted in pursuit of the desires.

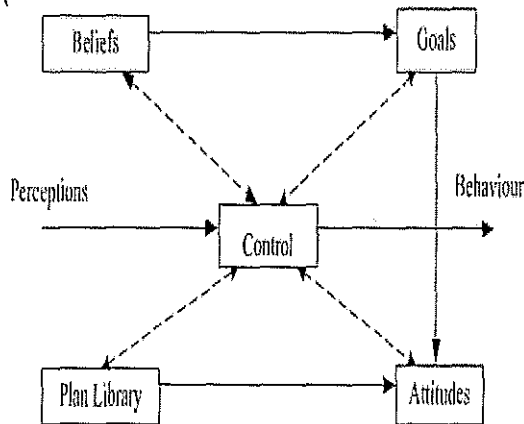


Figure 2: BDA agent architecture.

In comparison to traditional BDI [2] model, we have replaced *intentions* with *attitudes*. We say that *intentions* are primitive forms of *attitudes* without degree of commitment in them. An agent has a set of attitudes, each with a degree of commitment which persists according to the current situation. The attitudes are represented by following attributes:

Name of Attitude: This attribute describes the name of the attitude e.g. *like, hate, cautious* etc.

Description of Object: The description of the object contains the name of the object and a description of the internal organisation in terms of the components of the object.

Basic agent behaviour towards x: This attribute specifies the behaviour that will be performed by the agent with respect to the object *x*.

Evaluation: This attribute specifies whether the attitude is favourable or not.

Concurrent attitudes: This attribute specifies any other attitudes that can coexist with this attitude.

Persistence of Attitude: *This attribute specifies how long the attitude will persist under various situations. For example, it may specify how the attitude itself will change over time; that is, when to drop it and change it to another attitude, when to pick it up and how long to maintain it.*

Type of Attitude: This attribute specifies whether the attitude is individual or collective.

3.4 Attitude Based Coalition Model

We claim that successful coalition is achieved only if the agents have coalition as a collective abstract attitude. From this collective attitude, agents derive individual attitudes that are then used to guide their behaviours to achieve the coalition. Suppose there *n* agents in a coalition i.e. A_1, \dots, A_n . So the collective attitude of the agent A_1, \dots, A_n towards the coalition is represented as $Coal_{A_1, \dots, A_n}(A_1, \dots, A_n)$. But from A_1 's viewpoint, team is an attitude that it is holding towards the collection (A_1, \dots, A_n) and can be denoted as $Coal_{A_1}(A_1, A_2)$. Similarly from A_n 's viewpoint, its attitude can be denoted as $Coal_{A_n}(A_1, \dots, A_n)$. But the collective attitude $Coal_{A_1, \dots, A_n}(A_1, \dots, A_n)$ is decomposed into the individual attitudes only when all the agents mutually believe that they are in the coalition. The coalition attitude can be represented in the form of individual attitudes towards the various attributes of the coalition i.e. coalition methods, coalition rule base, and coalition responsibility.

The attitudes of an agent existing in a coalition consist of attitude towards coalition as well as attitude towards coalition activity. At any time, an agent may be engaged in one of the basic coalition activities i.e. coalition formation, coalition maintenance, and coalition dissolution. Instead of modelling these basic activities as tasks to be achieved, we have chosen to model them as attitudes.

Coalition (A_1, \dots, A_2)

This attitude is invoked when the agents are in a team state. This attitude guides the agents to perform the appropriate coalition behaviours.

Name of attitude: Coalition

Description of object: (1) Name of object: set of agents (2) Model of object: $\{A_1, A_n \mid A_i \text{ is an agent}\}$

Basic agent behaviour: coalition behaviour specified by agent's rule base

Evaluation: favourable

Persistence: This attitude persists as long as the agents are able to maintain it.

Concurrent attitudes: all attitudes towards physical and mental objects in the domain.

Type of attitude: collective.

3.4.1 Coalition Formation

Impetus for attitude and coalition formation may arise from the world and a particular domain or from agent's themselves. Having identified the potential for coalition action with respect to one of its goals, a leader agent will solicit assistance from some group of agents that it believes can achieve the goal. Our agents form coalition, because the inherent nature of the world requires agents to exist and act together. However, a particular situation may force the agents to dissolve the coalition for some time.

In the fire world, the event triggering the coalition formation process is a fire. Whenever there is fire, the security officers call the fire-fighting company to put out the fire. Then the fire fighters arrive at the scene of fire and get the information about when, how and where the fire had started. Suppose there is a medium fire in the campus, which results in the attitudes *medium-fire* and *dangerous-fire* towards the object fire. The attitude *Coal-form* is also generated, which initiates the team formation process. We propose a dynamic team formation model, in which we consider initially the mental state i.e. the beliefs of all the agents is same. The fire-fighting agents recognise appropriateness of the team model for the task at hand; set up the requirements in terms of other fellow agents, role designation, and structure; and develop attitudes towards the team as well as towards the

ain.

In order to select a member of the team, our agent will select the fellow agent who has following capabilities:

- Has knowledge about the state of other agents.
- Has attitude towards the coalition formation.
- Can derive roles for other agents based on skills and capabilities.
- Can derive a complete joint plan.
- Can maintain a coalition state.

Our method of forming a coalition is like this; the agents start broadcasting message to other agents "Let us form a coalition". The agents will form a coalition if two or more than two agents agree by saying, "Yes". If the agent does not receive the "Yes" message, it will again iterate through the same steps until the coalition is formed. The *coal-form* is maintained as long as the agents are forming the team. Once the team is formed, agents will drop the *coal-* attitude and form the *coal* attitude, which will guide the agents to produce various team behaviours.

Coal-form (A_1, \dots, A_2)

This attitude is invoked when the agents have to form a coalition to solve a complex problem.

Name of Attitude: Coal-form

Description of object: (1) Name of object: set of agents. (2) Model of object: $\{A_1, A_n \mid A_i \text{ is an agent}\}$

Basic agent behaviour: invokes coalition formation rules.

Evaluation: favourable

Persistence: The agent holds this attitude as long as it believes that a coalition formation is possible.

Concurrent attitudes: All attitudes towards physical and mental objects in the domain.

Type of attitude: individual

3.4.2 Coalition Maintenance and Dissolution

While solving a problem (during fire fighting activity) the coalition agents have also to maintain the coalition. During the coalition activity the agents implement the coalition plan to achieve the desired coalition action and sustain the desired consequences. The coalition maintenance behaviour requires what the agent should do so that coalition does not disintegrate. In order to maintain the coalition each agent should ask the other agent periodically or whenever there is a change in the world state, whether he is in the coalition. So the attitudes like *periodic-coalition-maintenance* and *situation-coalition-maintenance* are produced periodically or whenever there is a change in the situation. These attitudes help the agent to exhibit the maintenance behaviours.

When the team task is achieved or team activity has to be stopped due to unavoidable circumstances, the attitude *coal-unform* is generated. This attitude results in the dissolution of the team and further generates attitude *escape*. For example, when the fire becomes very large, the agents have to abandon the team activity and escape. The attitude *coal-unform* is maintained as long as the agents are escaping to a safe place. Once the agents are in the safe place, the attitudes *team-unform* and *escape* are relinquished. In case the fire comes under control, the agents again form a team by going through the steps of team formation.

4 Experimental Results

The key evaluation criterion for the overall performance of the coalition in any environment is that agents must successfully accomplish their tasks within their given environments, both efficiently and accurately. We have done several experiments for fire fighting in *FFWorld* domain to verify our ideas about the coalition model. The motivation behind doing these experiments is to determine the advantages and disadvantages of our attitude based coalition model. The agents in a coalition react to the changes in the world states by generating and achieving new goals. Meanwhile, old goals and plans are constantly being monitored and re-structured if necessary. The attitudes are mainly concerned with how to re-organise plans and goals due to situational changes. Whenever there is a new goal because of the changes in the world, one or more attitudes are usually created along with the goal. The experiments have concentrated on evaluating the performance of attitude based coalition in case of unexpected events. When a problem occurs, the attitude model stipulates a new set of behaviours for the agents. The agents with attitude respond to changes in the world by adopting a set of attitudes towards these changes.

We carried out different types of experiments to investigate the problem solving behaviour of the

agents during team action and to analyse their coherence in hostile dynamic environments. These sets of experiments demonstrate the significance of employing attitudes when agents have to deal with individual, group, social and team goals in a changing world. In our experiments, we have tried to measure the average payoff of the coalition under the varying conditions. To provide comparative, as well as quantitative results, three distinctive types of coalition were considered (i) *high* coalition (Coal (high)) (ii) *medium* coalition (Coal (medium)) (iii) *low* coalition (Coal (low)). In the *high* coalition, the commitment of agents is very large (> 80%). In the *medium* coalition, the commitment of agents towards coalition is medium (60%-80%). In the *low* teams (T_{low}), the nature of the task is such that they are not required to work closely together all the time and percentage of attitudes is very less (<60%).

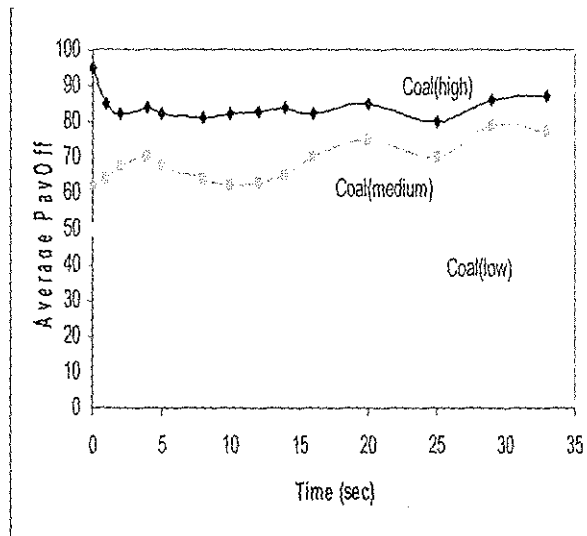


Figure 3: Average payoff of the agents with C_{high} , C_{medium} and C_{low} .

5 Conclusions

This paper has developed a novel framework for managing coalitions in a hostile dynamic world. Coalition is guided by the agent's dynamic assessment of agent's attitudes given the current scenario conditions, with the aim of facilitating the agents in coalitions to complete their tasks as quickly as possible. In particular, it is outlined in this paper that how agents can form and maintain a coalition, and how it can offer certain benefits to cooperation. Our solution provides a means of maximising the utility and predictability of the agents as a whole. Its richness presents numerous possibilities for studying different patterns of collaborative behaviour.

6 References

[1] M.E. Bratman, *Intentions, Plans and Practical Reason*, Harvard University Press, Cambridge, MA (1987).

[2] M.E. Bratman, D. Israel and M. Pollack, "Plans and resource bounded practical reasoning", *Computational Intelligence*, 4, pp 349-355 (1988).

[3] P.R. Cohen and H.J. Levesque, "Teamwork", *Special Issue on Cognitive Science and Artificial Intelligence*, 25(4) (1991).

[4] M. Fishbein and I. Ajzen, *Belief, Attitude, Intention and Behaviour: An Introduction to Theory and Research*, Addison-Wesley, Reading, MA (1975).

[5] M. Goyal, *An Attitude Based Teamwork Model for a Hostile Dynamic Multiagent World*, PhD Thesis, University of New South Wales (2001).

[6] M. Goyal, "Collaborative negotiations in a hostile dynamic world", *Proceedings International Conference on Artificial Intelligence (ICAI '04)*, Las Vegas (2004).

[7] N. Griffiths and M. Luck, "Coalition formation through motivation and trust", *Proceedings 2nd International Conference on Autonomous Agents and Multi-agent Systems (AAMAS 2003)*, Melbourne, ACM Press, pp 17-24 (2003).

[8] S. Kraus, O. Shehory and G. Taase, "Coalition formation with uncertain heterogeneous information", *Proceedings 2nd International Conference on Autonomous Agents and Multi-agent Systems (AAMAS 2003)*, Melbourne (2003).

[9] T. Sandholm, K. Larson, M. Andersson, O. Shehory and F. Tohmé, "Coalition structure generation with worst case guarantees", *Artificial Intelligence*, 111(1-2), pp 209-238 (1999).

[10] O. Shehory and S. Kraus, "Task allocation via coalition formation among autonomous agents", *Proceedings 14th International Joint Conference on Artificial Intelligence*, Montreal, pp 655-661 (1995).

[11] C. Li, S. Chawla, U. Rajan and K. Sycara, *Mechanisms for Coalition Formation and Cost Sharing in an Electronic Marketplace*, Technical Report CMU-RI-TR-03-10, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA (2003).