# Coalition Formation in Multi-Agent Systems under Real-World Conditions

**Michael van de Vijsel and John Anderson**
Autonomous Agents Laboratory
Department of Computer Science
University of Manitoba
Winnipeg, Manitoba, R3T 2N2, Canada
andersj@cs.umanitoba.ca

**Abstract**

The benefits of using multiple agents to tackle difficult problems are well-known. To date, however, most work in multi-agent systems has focused on making teams of agents work better under the assumption that those teams have already been hand-constructed. Comparatively little research has been performed on *coalition formation* — the process of forming groups of agents with possibly conflicting individual goals, in order to improve their collective capabilities. Current research in this area has produced a number of approaches for coalition formation, all of which work under strong assumptions of the nature of a coalition and the domain in which it operates. Unfortunately, these assumptions are generally not true of most real-world environments. This paper describes these assumptions in light of current work in coalition formation, and proposes a new approach.

## Introduction

Multi-agent systems research has illustrated again and again that groups of agents working together tend to achieve better results than individual agents working alone (Weiss 1999; Lerman & Shehory 2000; Fontan & Mataric 1996; Balch 1999). These successes manifest themselves in many different forms — achieving goals faster or at lower cost, delivering a higher quality of result, or producing reasonable performance with a greater number of more parsimonious agents.

Most work in multi-agent systems, however, involves starting with a team of agents and either providing a methodology allowing them to work together, or providing a mechanism for improving their group performance. Agents are in a group by default, and the work involved in deciding whether it is worth cooperating with other agents in the group is done by the researchers beforehand, generally using an assumption that every new body is a helpful one (e.g. (Fontan & Mataric 1996; Balch 1999)). While this may be fine for situations in which we are studying the effects of increasing the group size on a coordination mechanism (e.g. (Fontan & Mataric 1996; Arkin & Ali 1994)), or experimenting with the nature of the way in which individuals already in a group interact (e.g. (Veloso & Stone 1998)), there are many applications in the world where this is unreasonable. Most real world groups do not appear suddenly and fully staffed, but arise out of interactions and common ground between their members. Considering this brings up a host of issues that are ignored in most multi-agent systems work. In many real-world teams, for example, the likelihood of every potential member being of the same utility to the group (having the necessary level of ability or quality of work) is not high, such as an average human soccer team. We thus generally value the participation of some individuals over others, and even learn not to participate with less skilled individuals (Anderson, Tanner, & Baltes 2004). In some domains, agents one might have to decide whether to interact with might not only be less skilled, but malicious (e.g. e-commerce or other distance-based interaction, where software agents act on the motivations of a human user).

By presuming that a useful group formation process already exists, multi-agent systems research assumes away many important problems. Who should a particular agent attempt to form a group with? When should such a group be formed? Why should an agent want to form a group at all? How is a group formed? What set of constraints does being in a group impose on the agents involved?

In order to find answers to these and related questions, researchers have turned to the issue of *coalition formation* in multi-agent systems. The goal of coalition formation research is to form intelligent, cohesive groups that benefit all participating agents, to actively exclude agents that will hinder the group's performance, and to realize when it may not be necessary to work with other agents directly. In particular, the cost of forming and maintaining a coalition must be less than the gains achieved by the participating agents over the long term. Thus, the overall coalition formation process becomes evolutionary rather than static, with better and more refined groups being produced over time, as agents acquire more information about each other.

Two of the main measures used in evaluating the efficacy of a coalition are *throughput* and *stability*. Throughput measures the amount of work that a coalition can accomplish in a given amount of time, where the work being done varies from financial transactions (Breban & Vassileva 2001; 2002) to delivering packages (Sen & Dutta 2002). Stability measures the rate of change in coalition size and composition over time. Given that it takes some effort in the real world to evaluate joining a potential group, be accepted by that group, and similarly evaluate reasons for leaving and finding new groups to suit one's goals, stable coalitions are a desirable feature. There are similar benefits of stability to the group itself — in a human workplace, high turnover among employees in a company reduces the utility due to higher training costs and unfamiliarity among co-workers, not to mention lower employee morale in those who remain.

There has been some research examining approaches to forming coalitions of individually motivated agents and studying these in terms of throughput, stability, and related concepts (e.g. (Breban & Vassileva 2001; 2002; Brooks, Durfee, & Armstrong 2000; Brooks & Durfee 2002; Tsvetovat & Sycara 2000)). However, work to date has been done on the basis of simplifying assumptions that make the work inapplicable to real-world coalition formation situations in three very important ways. First, individual members of a real-world coalition will all have different perspectives on the coalition as a whole. Each coalition member may have its own agenda, and each may be using the coalition for somewhat different purposes (for example, using a social group as a place to meet potential business contacts). Similarly, the internal perspective of those in the coalition may be very different from the public perspective the coalition itself presents (e.g. most members using a group advertised as social in nature as a potential place to meet business contacts). Non-members of the coalition will only know what the coalition members choose to advertise. Most current work generally assumes that membership in a coalition is globally available to everyone (e.g. (Breban & Vassileva 2001; 2002)). Much current work makes even more restrictive assumptions related to this — for example, that agents have only one goal, completely omitting even the potential for coalitions to be used for different purposes (e.g. (Breban & Vassileva 2002; Brooks, Durfee, & Armstrong 2000)), or that agents can belong to only one coalition at a time, eliminating most conflict (e.g. (Tsvetovat & Sycara 2000; Breban & Vassileva 2001)).

Second, in the real world not all coalition members are equal — some members are more valuable than others. This is not only true from the perspective of the coalition (some workers are better), but is also true from the point of view of individual members. In the real world, members of a coalition may see the value of a coalition in large part through the participation of one or a few members, either because they trust those members, know those members are particularly good, or have personal goals that happen to coincide with those members. It may be possible that if a particularly valuable coalition member were to exit the group, other members may no longer have the inclination to remain with the coalition. Current work such as Durfee's research (Brooks, Durfee, & Armstrong 2000; Brooks & Durfee 2002), as well as Breban and Vassileva's work (Breban & Vassileva 2001; 2002), make the simplifying assumption of equal attractiveness of agents.

Finally, in the real world, coalition members will not start out with a large body of knowledge on the former two points; knowledge of the group and other agents will be incomplete and will be gathered over time, as an agent's experience in the coalition grows. While current research has devoted some effort to the issue of acquiring knowledge over time (Breban & Vassileva 2002; Brooks, Durfee, & Armstrong 2000), this has not been adequately applied to the above two issues.

We are currently working on an approach to coalition formation and maintenance that allows for stable and high-performing coalitions while addressing the above deficiencies in existing algorithms. An evaluation is proceeding using a software simulation, and preliminary results using a software simulation to compare our approach to the others described here, from the standpoint of throughput and stability, will be presented at the workshop. This paper describes our approach and its implementation in a software simulation. Before doing so, we provide an overview of related work in light of the assumptions presented above.

## Related Work

Coalition formation and maintenance research has been performed from a number of different perspectives. Some recent work comes from the area of game theory. Tohme and Sandholm (Tohme & Sandholm 1999) examine coalition formation with the goal of developing an algorithm that derives a stable coalition structure among a group of self-interested agents, when communication and deliberation are actions that have an explicit cost. While showing that communication is useful in forming stable coalitions, their work assumes that coalitions are used for the same purposes for all agents, and that the domains involved are simple enough that any two different coalitions are better off from the point of utility to merge into one. More importantly, this work is purely theoretical even within these simiplifying assumptions, and involves no practical implementation or implementation considerations.

Axtell (Axtell 2002) takes a different perspective on game-theoretic models of coalition formation. He assumes that agents are *non*-cooperative by nature, and shows that agents will still form cooperative coalitions under the correct set of constraints. Axtell's research strives to obtain a Nash equilibrium for the coalition formation process. He notes, however, that "to limit the focus of one's analysis [of multi-agent systems] to equilibria, while certainly augmenting mathematical tractability, is both highly restrictive and unrealistic, and likely to render the resulting models empirically false and operationally brittle." Axtell's criticism of his own work applies to some degree to all game-theoretic coalition formation models — while mathematically elegant, they tend not to make pragmatic considerations necessary for implementation and deployment. In addition, Axtell's model makes the same qualifying assumptions outlined in Section 1 that invalidate its use in many real-world settings.

Researchers have also examined a number of different multi-agent systems solutions to coalition formation problems, many of which have used an e-commerce or information marketplace domain to test their results. Tsvetovat and Sycara (Tsvetovat & Sycara 2000) implemented an electronic marketplace domain that allowed selling agents to provide goods at lower, wholesale prices to groups of buying agents. In their research, agents would only join a coalition or buying group if it could be shown that their buying price would be lower in the coalition than it would be outside the coalition. This introduces the concept of an agent having incentive to join a coalition. However, many of their other assumptions make their model unrealistic: coalitions are only active for one transaction, for example, and agents are all homogeneous in ability and have the same goal.

Yamamoto and Sycara (Yamamoto & Sycara 2001) created another coalition formation algorithm that focused on splitting groups of buyers into smaller groups which they called coalitions in order to maximize the buying power of the group. Buyers in these coalitions have different goals, in that they are looking to purchase differing but related goods (e.g. different types of "cameras"). In this sense, agents that will not contribute positively to the group are excluded. However, their model makes several simplifying assumptions. It allows one agent to speak for the entire group, and assumes that accurate information from all agents in the group is available. As well, an agent's preferences as to the type of goods they want to purchase remain static throughout the simulation.

Lerman and Shehory (Lerman & Shehory 2000) have created another information marketplace coalition formation strategy that removes many of the complexities of previous methods. Their methods assume very simple, self-interested agents that only have local knowledge of the system, and that requires minimal communication between agents. However, many of their simplifying assumptions make their model unrealistic. Agents all have the same goal, and it is always beneficial for agents to be a part of a coalition. As well, agents leave coalitions randomly, in the hope that there is a better coalition available to join.

Brooks and Durfee (Brooks, Durfee, & Armstrong 2000; Brooks & Durfee 2002) have done several studies on groupings of agents called *congregations*. Congregations in their terminology are subdivisions of agents in a system that are intended to limit the search space for individual agents (i.e. they function as subgroups of a population). Agents execute transactions only with other agents in their congregation, and if they deem the transaction unsatisfactory, they change congregations. Once again, the same assumptions surface. Brooks and Durfee postulate that congregations are intrinsically defined by their membership. Later work introduces *labeller* agents that attempt to attract congregating (or buying/selling) agents by placing labels on marketplaces. They attempt to put the correct label on a marketplace in order to attract the most agents and win more business. Thus, labeller agents must learn the preferences of the congregating agents in order to have a successful marketplace. While this brings in the concept of adaptation over time, agents still only have one type of goal — congregating agents only prefer one type of good, and labellers are only focused on marketplace labels. As well, labeller agents do not attempt to learn strategies of other labellers, which would help them to "outwit" each other in trying to attract congregators.

Breban and Vassileva (Breban & Vassileva 2001; 2002) have proposed a coalition formation mechanism that focuses on long-term coalitions lasting over multiple transactions. These coalitions are based on trust relationships that form between agents over time. In their model, agents interact with one other by attempting to execute a buy/sell transaction. If the agents involved in the transaction happen to be in the same coalition, then a discount is applied to the transaction. Once the transaction is complete, both agents evaluate their experience. If the experience is evaluated positively, the agent's trust in the other agent increases, and vice versa. Once the transaction is complete, the agent then reasons about its current coalition state. The agent can leave their current coalition, join a coalition (if they are not currently part of a coalition) or form a new coalition with the agent with whom they just completed a transaction. This research moves coalition formation in a positive direction by allowing agents to reason about their coalition status and providing many options for the agent to consider. As well, long-lasting coalitions are a crucial part of any realistic model, as is learning about other agents over time by forming trust relationships. However, their model still makes many of the same simplifying assumptions detailed earlier. Agents can only be part of one coalition at a time, and there are no costs to joining or remaining in a coalition. Also, they consider coalition membership to be global knowledge.

Sen and Dutta (Sen & Dutta 2002) examine a cooperative model for agents to assist one another without explicitly forming coalitions. In their model, agents are assigned a series of tasks (represented in their domain as packages that must be delivered to a particular location). Each task is assigned a specified cost. When agents encounter one another, they determine via a series of calculations if one agent can perform the tasks of both agents for less total cost than the sum of the individual task costs. If so, then an opportunity for cost savings exists. Sen and Dutta experimented with various strategies for agents to use in this case — philanthropic agents that always honour requests for help, selfish agents that always refuse requests for help, or reciprocative agents that base their decisions on past experience with the other agent. While not truly forming coalitions, reciprocative agents show the beginnings of what is needed for realistic coalition formation (identifying potentially useful agents), and their experimental domain involving package delivery proved to be useful and interesting.

There have also been a small number of implemented approaches to coalition formation in industry (Pechoucek, Marik, & Stepankova 2000; Pechoucek, Marik, & Barta 2002; Contreras & Wu 1999). In all these cases, however, the same simplifying assumptions are made — agents are all working towards a single goal, agents are only part of a single coalition, and there is no acquiring of information over time regarding agent behaviours or strategies. In these cases, the application domains for these implementations allow for these simplifications without any limitation in functionality.

## Coalition Formation under Real-World Conditions

We are currently implementing a new approach to coalition formation that increases the realism in the overall process while maintaining a scalable approach suitable for systems with a large number of agents. The implementation of this algorithm and its evaluation both require a domain in which to operate. We begin with a description of this domain.

## Package Delivery Domain

The domain in which we are exploring coalition formation is an adaptation and extension of Sen's package delivery system (Sen & Dutta 2002). The domain involves a set of agents that receive packages to deliver to specified addresses. Sen employs a domain in which agents receive their initial packages at a central hub, and the delivery address is on one of a number of radial fins extending away from the hub. Agents cannot move between radial fins, only along them, like the spokes on a wheel. Once the delivery location is reached, agents must to travel back to the central depot to receive their next package to be delivered. They can enlist the help of another agent at the depot, but not of any agents encountered while travelling. An agent can only provide help if it is already intending to travel along the desired radial fin (Sen & Dutta 2002).

Our domain extends this work by adding several key elements. First, agents are permitted to have multiple goals at the same time. Without multiple goals, an agent never has any conflict over what it should do next (and thus no choice of action that can affect its performance) — it can focus completely on getting its current package delivered. In our model, agents have multiple packages (i.e. multiple goals) at any point in time, and their delivery locations can be in opposing directions. Agents must decide which package to deliver first, and must also determine for which package it would be most advantageous to solicit help from another agent. In order to successfully implement this element, agents must be given a greater degree of freedom of movement through their environment. This desire led us to move from a radial to a grid-based implementation.

Second, agents can solicit help from other agents in a coalition whenever they encounter another agent belonging to the same coalition, and can also form or extend coalitions upon encountering another agent. This is intended to increase the number of agent interactions throughout the model — agents are not limited to interacting only at the central depot. More interactions between agents means more opportunities to form coalitions, establish relationships, solicit help delivering packages, and learn about other agents. Our hypothesis is that this should translate to the earlier arrival of a stable set of coalitions, as well as increased throughput, since agents will be helping one another as often as possible.

A third change to Sen's model is that there are several package depots scattered throughout the environment instead of at a single central depot. These depots are placed randomly and located at a minimum distance from one another. By specifying a minimum distance between depots, we allow for more variation in travel patterns, requiring agents to branch out to all points on the grid from each package depot.

Finally, we are extending Sen's model by providing the agent with a payoff for each package that is successfully delivered. The initial payoff value is derived based on the distance from the location where the agent received the package to the package's delivery address – the further away the delivery, the higher payoff on the package. Payoff for packages diminishes over time, eventually becoming a penalty if the package is not delivered (with a maximum penalty for non-delivery of a package). This payoff mechanism allows agents to make decisions based on expected payoff of delivering a package. It also gives agents a "currency" that can be used to implement costs of particular actions (e.g. vacating a coalition, paying coalition membership dues, etc.)

Together, these changes transform a traditional experimental environment for coalition formation into one that is reflective of many of the characteristics we observe in real-world multi-agent problem-solving environments. This environment is being implemented using Java, in order allow for a variety of derived agent types to interact with this environment in a modular fashion. The environment is modular in that it accepts an agent that fits a base pattern (or interface) and agent models can be extended to implement more complex reasoning as necessary.

## Agent Model

The agent model is the heart of the entire system – with the agent resides the knowledge and skill to form coalitions intelligently. Several things are required to allow agents to make reasonable decisions about coalition formation.

First, agents must be heterogeneous. If all agents were the same except for their private goals, then there would be little inclination to form coalitions with any one agent as opposed to another. To support agent heterogeneity, agents have specific, unchanging values for each of the following attributes:

- Honesty: Agents with a high honesty value will present accurate pictures of themselves to each other, while dishonest agents may exaggerate their attributes to other agents.
- Memory: Agents with a high memory value will have a greater probability of remembering to deliver packages for other agents as well as being able to remember significant locations in the environment (e.g. where to pick up packages). Agents with a low memory value will tend to forget that they agreed to help another agent deliver their package, thus inducing a penalty for non-delivery to the original agent.
- Speed: Agents with a high speed value can cover more ground while moving around the grid than agents with a slow speed value. This attribute is very valuable, since agents with high speed values will deliver packages faster, generating higher payoffs for the originating agent.
- Trust: Agents with a high trust value will have a higher inclination to give their packages to other agents for delivery. Agents with a low trust value will prefer to deliver their packages themselves.

Values for these attributes are randomly assigned when agents are created, resulting in a range of agents with variable likelihoods of being successful in a coalition for this domain (and from the point of view of the coalition and the agents in it, a variable degree of value to the coalition).

Agents are initially placed at random locations in the world, with no information about the world other than the coordinate system used for addresses and its own current location (expressed as a coordinate). Agents wander until they encounter another agent or a package depot.

If the agent encounters a package depot, it remembers the depot's location (to the degree the agent's memory is limited, allowing it to potentially forget and have to find package depots again, resulting in less productive work for agents with limited memory) and it will receive multiple packages to deliver to allow for a range of competing goals. The agent determines which package to deliver first by calculating which will have the highest payoff based on the package's payoff value, the distance to the delivery address and its own speed value. It then begins journeying towards that package's delivery address (a specific coordinate in the world).

If the agent encounters another agent, they have the opportunity to learn useful information about one another, as well as an opportunity to form or extend coalitions.

### Coalition Formation Mechanism

The realistic elements of the environment described above allow the inclusion of more features in a coalition formation approach than have been used in the prior studies cited in Section 2. The mechanism employed for coalition formation, described abstractly, is as follows. When two agents meet, there are a series of questions that each agent asks of itself:

1. Have I met this agent before? If so, do I have a favourable opinion of this agent?
2. Is this agent in a coalition that I am already in?
3. Do I think this agent is worth asking to join any coalitions I belong to?
4. Am I interested in joining any of the coalitions this agent is a member of?
5. Do I want to form a *new* coalition with this agent?
6. Should I ask this agent to assist me with any of my goals? (this is only asked if the agents now share a coalition)

The answers to these questions determine whether each agent will attempt to work with the other in a new or existing coalition. By considering the newly-encountered agent's suitability for any existing coalitions, as well as considering one's own suitability for any of the encountered agent's coalitions, each agent can determine which coalition would be most appropriate to invite the other agent to join. In particular, to answer questions 3-5 above, the inviting agent considers:

- The new agent's attributes (as advertised by that agent — these may be exaggerated based on that agent's honesty);
- The average attribute values of agents already in in the agent's current set of coalitions, and any obvious deficiencies among coalition members;
- The current size of the coalition — smaller coalitions will want to grow by inviting members, while larger coalitions may have enough functional members to be self-supportive and not be viewed as having a strong need for further members.

Once an invitation to join a coalition is issued, the agent being invited must also consider several factors before accepting or declining the invitation, such as:

- Public-knowledge features about the coalition, such as cost to join, cost to maintain membership, cost to leave the coalition, number of members, identities of members it knows (membership may or may not be public information) and their perceived value;
- The agent's perception of the inviting agent — if the inviting agent is viewed unfavourably, the agent may not want to be part of a coalition of such unfavourable agents;
- The agent's current coalition situation — if it is not a member of any coalitions, it will be more likely to join an unfavourable coalition in order to gain colleagues.

If there is no invitation issued, or the invitation is rejected, the two agents may still decide to form a *new* coalition (question 5). This follows similar reasoning – if each agent feels the other is trustworthy and possesses complementary skills, but there is no reasonable fit to any existing coalition, they can decide to start a coalition on their own. In this case, the agents must determine the "ground rules" for the coalition, including:

- What are the costs associated with the coalition?
- Are there rules for enlisting the help of other members?
- How much coalition information is available to the public?

For the purposes of this work, these ground rules will be controlled externally, so that the costs, available information, etc. is similar to other coalitions for evaluative purposes. Negotiation of the ground rules of a coalition is beyond the scope of this work, but provision for it is being left in the general approach in order that this can be examined in future.

After going through all the above deliberations, the two agents may or may not find themselves sharing a coalition. If they do not share a coalition, they move on. If they do share a coalition, they each must now decide whether to enlist the help of the other agent in delivering any current packages (question 6). The agent asking for help examines the expected gain in payoff that would be received if it enlists the help of the other agent in delivering the package, taking under consideration the speed and memory attributes of the other agent. If it decides to ask for help, the other agent will take into consideration its current direction and the expected loss of payoff from existing packages to be delivered if it decides to help. If the expected loss is acceptable, the request for help is accepted.

This approach improves on prior work in coalition formation by removing the three main assumptions that previous work entails. Since each agent has differing attributes, agents' perspectives on coalitions will be different as well. For example, an agent with a low speed attribute may join a coalition of agents that all have high speed attributes in order to compensate for its lack of strength in this area. Other agents might be in the same coalition to compensate for other deficiencies, or even for reasons unrelated to their own attributes — low cost of coalition participation, for example. Each member may have a different perspective on why the coalition is valuable, including attributes of other agents that belong to the coalition. This, along with the fact that coalitions have both public information and information to be gleaned only as a member of the coalition, eliminates two of the three major limiting assumptions of previous work.

After an agent encounter, agents update their own memory with their impressions of the encountered agent so that they may recall the other agent if they meet again. Agents record attribute values (if those were exchanged between the agents) and continue to update their memory of the encountered agent on subsequent meetings. Thus, agents acquire information regarding the tendencies and attributes of other agents over time, removing the third limiting assumption of previous work. The recollections of other agents are not subject to memory limitations here – the memory limitations are designed to have an impact on the agent's ability to perform well in the environment and thus be of greater or lesser value to other agents. Limiting the ability to recall others is a level above this, limiting the agent's ability to participate in coalition formation. Here all agents have the equal ability to form coalitions; future work will involve varying elements at the coalition formation level as well.

At any point, an agent may decide to leave any of their existing coalitions. Agents track the amount of help that a coalition has provided to them in terms of payoff resulting from help by agents in that coalition. This and other information the agent acquires serve to augment the initial assessment of value based on publicly available information. Agents thus want to remain in coalitions that provide help to them, and will not want to remain in coalitions that are not useful.

This model provides the needed flexibility to realistically model coalitions between many agents, while still limiting communication between agents to direct encounters. This limit on communication is important in terms of the utility of this approach, since significant ongoing communication – another common problem among complicated coalition formation algorithms (Tsvetovat & Sycara 2000) — is costly in practice and results in a bottleneck. The only ongoing information an agent receives in this approach is from the environment itself, and consists of its current location and the payoff provided when a package is delivered (either by the agent or another enlisted agent).

## Evaluation and Discussion

The implementation of the approach described in the previous section is still in progress and initial results of its performance will be presented at the workshop. The evaluation of this approach will examine issues of system throughput and coalition stability, using the mechanisms of Breban and Vassileva, described in Section 2, as a baseline. This particular approach was selected for comparison over other alternatives because it is the most advanced of the existing approaches cited here in terms of dealing with real-world issues, and also translates well to the domain we are working with.

As an evaluation we intend to run a series of trials for each of Breban and Vassileva's approach and the approach described above, using a 100 x 100 grid for the testbed with a set of 500 agents. This will allow agents to take up five percent of the environment area, which should allow for reasonable freedom of movement while still providing for relatively frequent agent interactions. The number of packages picked up at each location will be a random value between three and five. The alternative approaches on which that described here is based do not consider a spread of agent honesty, as well as many of the other potentially variable elements in the version package delivery domain used here. In order to allow valid comparison between our approach and prior work, we will control these potential random variables, ensuring that an even spread of "good" and "bad" agents (both in terms of honesty and memory) are created, as well as ensuring that such items as coalition costs (joining, leaving, maintaining membership), making membership information public or private, and costs for enlisting help of other members are all evenly distributed throughout the simulations. We will track the number of coalitions in the system, the number of agents in each, and the number of coalition changes over time in order to examine coalition stability, while the number of packages delivered and the average delivery time for a package will serve as a measure of system throughput.

The richness of the environment also allows many other interesting factors to be examined, and we also hope to gather some initial data on the effects of varying rates of agent honesty and cost of leaving and joining coalitions. This will serve as a preliminary study using this domain to examine issues of trust and deception in coalition formation, and the additional data should also allow a fine-tuning of the system in order to provide better overall performance. While there are many additional avenues of experimentation, these will be left to future work.

The main contribution of this work is a coalition formation mechanism that allows for differing perspectives on coalitions between agents, a variation in the value of each agent to a coalition, and the gathering of information regarding these two items by agents over time. By introducing agents with different attribute values, we hope to promote heterogeneity among agents and among the coalitions that they form. Agents model the domain they reside in over time — thus forming their own perspectives about the agents they encounter and the coalitions they join. Agents are constantly updating their information as they learn more about the attributes and subsequent behaviours of the agents around them, thus adding an element of learning to the overall domain. This allows agents to form coalitions when it is advantageous to all concerned, allowing for stability in the overall system and producing a high-functioning society.

## References


Anderson, J.; Tanner, B.; and Baltes, J. 2004. Exploiting opportunities through dynamic coalitions in robotic soccer. In *Forming and Maintaining Coalitions and Teams in Adaptive Multiagent Systems*. (to appear).

Arkin, R. C., and Ali, K. S. 1994. Integration of reactive and telerobotic control in multi-agent robotic systems. In *Proceedings of the Third International Conference on the Simulation of Adaptive Behaviour*.

Axtell, R. L. 2002. Non-cooperative dynamics of multi-agent teams. In *Proceedings of the first international joint conference on Autonomous agents and multiagent systems*. ACM Press.

Balch, T. 1999. The impact of diversity on performance in multi-robot foraging. In *Proceedings of the Third Annual Conference on Autonomous Agents*, 92–99.

Breban, S., and Vassileva, J. 2001. Long-term coalitions for the electronic marketplace. In Spencer, B., ed., *Proceedings of the E-commerce Application Workshop*.

Breban, S., and Vassileva, J. 2002. Using inter-agent trust relationships for efficient coalition formation. In Cohen, R., and Spencer, B., eds., *Proceedings of the 13th Canadian Conference on AI*.

Brooks, C. H., and Durfee, E. H. 2002. Congregating and market formation. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems*, 1–100.

Brooks, C. H.; Durfee, E. H.; and Armstrong, A. 2000. An introduction to congregating in multiagent systems. In *Proceedings of the Fourth International Conference on Multiagent Systems (ICMAS-2000)*, 79–86.

Contreras, J., and Wu, F. F. 1999. Coalition formation in transmission expansion planning. *IEEE Transactions on Power Systems* 14(3):1144–1152.

Fontan, M. S., and Mataric, M. J. 1996. A study of territoriality: The role of critical mass in adaptive task division. In Maes, P.; Mataric, M.; Meyer, J.-A.; Pollack, J.; and Wilson, S. W., eds., *Proceedings, From Animals to Animats IV, Fourth International Conference on Simulation of Adaptive Behaviour*, 553–561. MIT Press/Bradford Books.

Lerman, K., and Shehory, O. 2000. Coalition formation for large-scale electronic markets. In *Proceedings of the Fourth International Conference on Multiagent Systems (ICMAS-2000)*, 216–222.

Pechoucek, M.; Marik, V.; and Barta, J. 2002. A knowledge-based approach to coalition formation. *IEEE Intelligent Systems* 17(3):17–25.

Pechoucek, M.; Marik, V.; and Stepankova, O. 2000. Coalition formation in manufacturing multi-agent systems. In Pechoucek, M., and Marik, V., eds., *Workshop on Industrial Applications of Holonic and Multi-Agent Systems, 11th International Conference on Database and Expert Systems Applications*, 241–246. IEEE Computer Society.

Sen, S., and Dutta, P. S. 2002. The evolution and stability of cooperative traits. In *Proceedings of the first international joint conference on Autonomous agents and multiagent systems*. ACM Press.

Tohme, F., and Sandholm, T. 1999. Coalition formation processes with belief revision among bounded-rational self-interested agents. *Journal of Logic and Computation* 9:1–23.

Tsvetovat, M., and Sycara, K. 2000. Customer coalitions in the electronic marketplace. In *Proceedings of Autonomous Agents 2000*, 263–264.

Veloso, M., and Stone, P. 1998. Individual and collaborative behaviours in a team of homogeneous robotic soccer agents. In *Proceedings of the Third International Conference on Multiagent Systems*, 309–316.

Weiss, G. 1999. *Multiagent Systems*. Cambridge, MA: MIT Press.

Yamamoto, J., and Sycara, K. 2001. A stable and efficient buyer coalition formation scheme for e-marketplaces. In *Proceedings of the International Conference on Autonomous Agents*, 576–583.