Forming Stable, Overlapping Coalitions in an Open Multi-agent System (Position Paper)

Brandon Wilson and **Marie desJardins**
Department of Computer Science and Electrical Engineering
University of Maryland, Baltimore County
Baltimore, MD 21250

Abstract

Coalition formation is an essential component for multi-agent systems in environments where tasks can be too complicated or resource intensive to be completed efficiently by a single agent. By forming coalitions, agents are able to work together and combine resources to complete a task that would not otherwise be possible. We propose a mechanism for forming stable, overlapping coalitions of self-interested agents. Current coalition formation algorithms tend to assume that coalitions cannot overlap, meaning that an agent can only be a member of one coalition at any given time; their restrictions prevent the system from efficiently allocating all of the agents' resources.

Keywords: Coalition, coordination, cooperation

Introduction

There is an obvious potential for agents in a multi-agent system to work together as a group, or coalition, to solve large problems. As tasks are introduced to a the system, the resources required to complete each task can vary. For some tasks, the resources of any single agent may not be sufficient. The formation of coalitions of agents can alleviate this problem. A *coalition* is defined as a group of self-interested agents that work together and share resources to complete a task with the expectation of receiving a benefit upon completion, either in the form of future cooperation on its tasks, or an immediate payoff from completing the task. Therefore a coalition allows groups of agents to complete tasks that may not be possible to accomplish alone or to complete them more efficiently. Figure 1(a) depicts a multi-agent system with a set of agents, $A = A_0, ..., A_8$, in one potential coalition configuration.

An example of realistic situation that is simplified via coalition formation is the process of moving residences. While it is usually possible to pack and relocate all of one's belongings alone, a person will typically enlist the help of friends. Many of these friends are self-interested and will decide to help because they assume that their friend will help the next time they are moving, or perhaps because they are being paid immediately for their services as is the case with moving companies. In both situations, the people work together as a coalition to achieve the higher-level task of moving all of the belongings from one residence to another.

Current work typically restricts coalition formation to disjoint coalitions; overlapping coalitions are often assumed not to exist. A set of N overlapping coalitions, $C = C_0, ..., C_N$, is defined as coalitions that share at least one agent in common (Shehory & Kraus 1996). Figure 1(b) illustrates a configuration of agents in which some coalitions overlap. Overlapping coalitions are most beneficial in large-scale domains with thousands of agents where a select few agents have a rare, but highly demanded, resource. By allowing the agents that are in high demand to join multiple coalitions, they can maximize their utility by committing to multiple coalitions. Consider an example where agents have only a single resource: computing power. An agent that is implemented on a supercomputer should not be restricted to joining a single coalition because it has more than enough processing power to share with several coalitions, as long as tasks are of moderate size. The coalitions are also profiting from this relationship because they do not have to wait for an agent to leave one coalition, then rush to acquire them before they join another coalition. We are interested in developing protocol will allow for overlapping coalitions and attempt to maximize the allocation of agent resources.

As a solution to the problems discussed thus far, we propose methods of forming stable, overlapping coalitions in large-scale, open multi-agent systems. The remainder of this paper is structured as follows. Section 2 contains a brief background and review of related work in the field of coalition formation. In Section 3, we present a formal definition of the problem domain; Section 4 provides a detailed explanation of the proposed approach that we plan to take in solving the problem. Section 5 concludes the position statement.

Background and Related Work

We define the concept of a *multi-agent system* (MAS) as a system that contains two or more agents. We can formally define a multi-agent system as a tuple, {A, E, R}. *A* is a set of *n* agents, $A = A_1, ..., A_n$, where $n > 2$, *E* is the environment, and $R = R_{11}, ..., R_{nj}$, are the resources available to each agent in the system; there are *j* total resources.

The problem we are trying to solve can be divided into two areas: disjoint coalition formation, and overlapping coalition formation. A great deal of work has been performed on disjoint coalition formation, but forming coalitions that may overlap is relatively unstudied. The following sections discuss the work that has been performed in each of the individual areas.

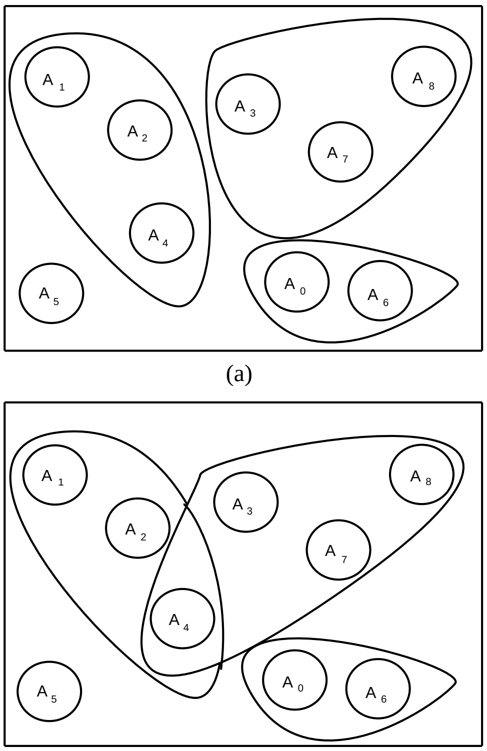

(b)

Figure 1: *Figure (a) represents a set of nine agents, $A = A_0, ..., A_8$, in a multi-agent system are partitioned into a set of three coalitions, $C = C_0, C_1, C_2$. In this instance, $C_0 = A_1, A_2, A_4$, $C_1 = A_3, A_7, A_8$ and $C_2 = A_0, A_6$. Agent A_5 does not belong to a coalition at the time represented by this snapshot. (b) The same set of agents as in a different coalition structure, with two overlapping coalitions that share agent A_4.*

Disjoint Coalition Formation

The study of coalition formation has generally focused on task-based multi-agent systems in which a set of agents are posed with a series of tasks that must be completed. The formation of coalitions allows the agents to complete tasks that are impossible to complete alone within a given time bound.

Caillou et al. (2002) describe a coalition formation protocol that is capable of determining a Pareto-optimal set of coalitions, based on the assumption that each agent has a set of preferences over the possible coalitions. This set is considered Pareto-optimal because it has the characteristic that in order to make any agent in the set more satisfied with the assignment, at least one other agent will become less satisfied. Here satisfaction is measured by a utility function that may vary from agent to agent. The protocol determines the set of coalitions without disclosing the individual agent preferences. This is distinct from other coalition formation protocols in that it does not aggregate individual agent utility functions to determine the coalitions. The secrecy of the agent preferences makes this algorithm applicable to an open system where information disclosure can come at a high cost. This method, however, does not consider coalition overlap.

Rathod and desJardins (2004) studied the effect of coalition duration on the total tasks completed and total revenue of a coalition. The domain of interest was a task oriented domain and the system was populated with self-interested agents that each possess one resource. A subset of the agent population was selected to be team leaders. The team leaders bid on tasks and formed coalitions to complete the tasks that they were awarded. The results of the study showed that in such a domain, the stable coalitions, coalitions that remained together for longer than one task, accrued a larger overall total revenue than dynamic, one-shot coalitions. This study assumed that the system was closed and agents belonged to only a single coalition. From this study, we draw the hypothesis that stable coalitions will also perform well when overlap is permitted between coalitions and agents possess multiple skills.

One of the major difficulties in the coalition formation problem is that computing the optimal set of coalitions is NP-complete (Gaston 2005). Sandholm et al. (1998) present an coalition formation algorithm that is designed to form coalitions for real-time applications. The algorithm models the problem as a search problem through the graph of all possible coalition sets. The final solution may not be optimal, but it is guaranteed to be within a specific bound of the optimal solution. The authors also derive an upper bound on the running time of the algorithm that is acceptable for a real time application.

Sandholm et al. (1998) decrease the running time further by distributing calculations among the agents. To address the possibility of malicious agents reporting incorrect results of their portion of the calculation, checks are performed at random by duplicating calculations to catch lying agents. A significant punishment is applied to an agent that is caught reporting false results and thus providing motivation to behave honestly.

Overlapping Coalition Formation

Most of the work on coalition formation constrains each agent to belong to only one coalition at a time. Allowing agents to join multiple coalitions, or allowing coalitions to overlap, proves to be useful in situations where some agents are in possession of resources that are rare, yet in high demand. Agents with rare resources can maximize their profit without hindering the other agents. It would hinder the system greatly if an agent had a common resource as well as a rare resource and it belonged to a coalition that utilized its common resource, but neglected the rare one. Figure 1(b) represents a set of agents that have been divided into coalitions. Two of the coalitions share agent A_4 as a common member, presumably because it has resources that both coalitions need, but are not possessed by any of the "free" agents (in this case, A_5).

Shehory and Kraus (1996) adapted the idea of overlapping coalitions to a distributed problem solving domain. Their algorithm successfully utilizes resources in an efficient manner and maintains low computational complexity. This algorithm assumes that the agents in the system are all interested in the success of the overall system. The algorithm will not work in a system of self-interested agents because it does not consider the possibility that agents may have incompatible goals, or an agent may be unreliable. Therefore, any agent may decide to be uncooperative, destroying the whole system.

Problem Domain

The domain we are modeling is task oriented and populated with self-interested agents. The agents are assigned tasks randomly; each task t is specified by $R = r_0, ..., r_n$, which is the set of resources required to complete the task. Each task has an associated timeout and a payoff $P(t)$ that designates the payment that will be received if the task is completed prior to expiring after a timeout. Maximizing the total payments received is assumed to be a general goal of all agents in the system.

There is a task generator that randomly introduces tasks and assigns them to an agent in the system. The agent is then required to find a coalition that can accomplish the task. If the combined resources of the agents in a coalition are at least the same as those required by the task, then the task can be completed. On the other hand, if the combined set of resources of all agents in a coalition are more than a task requires, then the surplus resources are being wasted. For this reason that we allow agents to join multiple coalitions to maximize their resource usage.

Agents in the system are motivated to achieve the highest profit possible; therefore, we model the agents as service providers. The services that the agents offer are the resources that they possess. This provides two sources of income in the system: a direct payoff for completing a task and a payment for assisting another agent in completing a task. As is the case in the real world, an agent is able to set its own price for its services, allowing agents with rare resources to maximize their profit. This also limits an agent from overcharging for a common resource because there are so many other providers.

Proposed Solution Approach

We propose a solution to the problem of forming overlapping, stable coalitions of agents based on the *membership fees and benefits* business model. Many wholesale retailers utilize this model to maintain their customers' loyalty. These businesses sell products in bulk at a reduced rate as a benefit to those who purchase a membership. Those who do not wish to become members can still shop at other stores that do not offer the same benefits, but also do not require a paid membership.

Membership Fees

For our model, we define a coalition as a group of self-interested agents that are trying to maximize their profit. To join a coalition, C_i, an agent must pay a membership fee that is a function of the size of that coalition, $f(|C_i|)$. The optimal function for this fee will be derived through experimentation. The nature of the function results in large coalitions being costly to join, deterring these large coalitions from ever existing. The membership fee is charged periodically to ensure that agents do not remain in a coalition if it is no longer useful. The fee also discourages agents from joining many coalitions since this can become non-profitable. Agents are then encouraged to start coalitions so they can avoid this costly fee when the coalitions are small. If at any time an agent is not able to pay its membership fee, it is forcefully removed from the coalition.

For the model to be robust against malicious agents, it must reward agents that contribute rare resources to the coalition and deter agents that contribute resources that are already in surplus. This is accomplished by discounting the membership fee by an amount proportional to the *rarity* of the resource. For a particular coalition, the rarity of a resource R_i is measured as the percentage of coalition resources accounted for by R_i.

In addition to the membership fee, there is also an exit fee, which is where our model differs from the conventional business model. For an agent to leave a coalition, it must pay a fee that is proportional to how reciprocative the agent has been to the rest of the coalition. For example, an agent that enters a coalition and uses its members to complete five of its own tasks, but never assists the members with their tasks, will face a large fee upon exiting the coalition; whereas another agent that has also used the coalition to complete five tasks, but has also assisted the coalition members in completing five of their own tasks will be permitted to exit for free. The reciprocative nature of an agent is also incorporated into the membership fee to prevent an agent from abusing a coalition forever.

Joining a Coalition

An agent that has been assigned a task that it can complete alone does not have any motivation to join a coalition. Agents that have been assigned a task and need help to complete it are faced with deciding which coalition to join.

An agent will examine all of the coalitions that it currently belongs to and if any of them have the resources to complete the task, then the agent works with them. If the agent does not belong to a coalition that can help or it does not belong to any coalitions, then it examines the other coalitions in the system. An agent that is interested in a coalition C can request a list of agents and resources of each agent from any member of C. The coalition that possesses the resources to complete the task and offers the lowest membership fee is selected. In the case that the agent still does not find a valid coalition, it starts a coalition of its own. Once an agent is part of a coalition or starts a coalition, it can recruit members by advertising open tasks to "free" agents. Despite starting a coalition, the agent periodically performs a check of the coalitions to see if any have grown to offer a valid set of resources. If the agent finds a coalition during one of these periodic checks then it can dismantle the coalition it formed without penalty.

Membership Benefits

Members of a coalition gain several benefits, some of which are implicit to our mechanism. When an agent enlists the resources of another agent in the same coalition, it is supplied the resources at a discounted cost, where the discount factor, α, is a percentage (in the range [0.0 - 1.0]).

If an agent is assigned a task that it cannot complete with agents in its coalition, then it can obtain information about agents that could be helpful in the task from other agents in the same coalition. To encourage agents to cooperate with these queries, their responses contribute to reciprocity among the coalition and thus reduce their exit fee.

The membership fees are another potential benefit of joining a coalition. The membership fees that are paid for a coalition membership are distributed among the members of the coalition evenly. This provides a motivation to join small coalitions instead of creating oversized coalitions.

Disbanding and Reforming of Coalitions

The proposed mechanism will form stable coalitions in the presence of similar tasks. Tasks are considered similar if they have approximately the same resource requirements. Agents form coalitions that have a resource pool that is capable of accomplishing a given type of task. If the task distribution alters significantly, then the agents in the coalitions will stop receiving benefits from their coalitions and will ultimately exit their current coalitions – which are no longer making profitable returns – reforming in coalitions that can perform the new tasks. As a result, we believe that the mechanism we propose will gracefully disband coalitions that no longer offer significant benefits and reform coalitions that are a more efficient allocation of the agents' resources.

Conclusion

Coalition formation is a well-studied area of multi-agent systems, but most of the work has focused on disjoint coalition formation that can result in an efficient allocation of resources. Overlapping coalitions have been overlooked in much of the existing research, but they are crucial in environments where agents possess multiple resources and wasting those resources is very costly.

We have presented our ideas on forming coalitions that result in maximizing the allocation of resources. Our ideas also address the problem of self-interested agents unreliability of reciprocating actions..

References

Caillou, P.; Aknine, S.; and Pinson, S. 2002. A Multi-Agent Method for Forming and Dynamic Restructuring of Pareto Optimal Coalitions. *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems: Part 3* 1074–1081.

Gaston, M. 2005. *Organizational Learning and Network Adaptation in Multi-Agent Systems*. Ph.D. Dissertation, University of Maryland, Baltimore County.

Rathod, P., and desJardins, M. 2004. Stable Team Formation Among Self-Interested Agents. In *In Working Notes of the AAAI Workshop on Forming and Maintaining Coalitions in Adaptive Multiagent Systems*.

Sandholm, T.; Larson, K.; Andersson, M.; Shehory, O.; and Tohme, F. 1998. Anytime Coalition Structure Generation with Worst Case Guarantees. *Arxiv preprint cs.MA/9810005*.

Shehory, O., and Kraus, S. 1996. Formation of overlapping coalitions for precedence-ordered task-execution among autonomous agents.