Negotiation to Improve Role Adoption in Organizations

Asad Rahman and Henry Hexmoor
Computer Science and Computer Engineering
Engineering Hall, Room 340A
University of Arkansas
Fayetteville, AR 72701
{rahman, hexmoor}@uark.edu

Abstract

We present negotiation schemes for efficient role adoption that enhance utility in organizations. In one scheme, local utility computations determine role adoption. In the second scheme, utility of the entire system is considered by the negotiating agents. Our results show that when agents negotiate using the latter strategy, the system performance converges to its pareto optimal utility level.

1. Introduction

We consider organizations with fixed number of individuals in fixed roles and individuals at varying fitness levels to roles. Self-organization of individuals of an organization allows for them to change roles in the organization. Efficient role adoption can increase performance of the organization [10]. The adaptation depends on many factors, which include consideration of institutional norms and autonomy [5, 9]. Since agents are independent, centralized schemes for role assignment do not apply. Instead, an agent can freely choose to suggest a role swap with any other agent in the system. To ensure optimal system performance, agents use negotiation to adopt new roles. Negotiation is the process during which participants communicate with one another to come to a mutually acceptable agreement on any matter [1]. The motivation for negotiation is supported by the ideas presented in [10] where managers of a group of agents in a coalition use negotiation to efficiently allocate agents to other coalitions. This notion of negotiation can be further extended to allow all members of an organization to actively negotiate.

Negotiation between agents, though costly, can ensure an optimal system performance [7]. During this process, participant's perception and preferences may also change [7]. In our paper, we do not consider negotiation that alters participant's perception or their preferences. Negotiation between agents can be described in a variety of ways. We use the state diagram model in [9] and the conversational scheme in [10] to describe negotiation between agents.

The main premise behind the state diagram model in [9] is to describe negotiation in terms of state transitions. A negotiation state denotes whether or not there are any pending or new proposals in the system. A negotiation is initiated because a proposal to adopt a new role is suggested by an agent. This causes the system to transition into a new state. The negotiation is carried forward by subsequent counter-proposals or concessions leading to other states. The system may finally move back to the original state after, either a mutually acceptable agreement is reached between two agents, or no agreement is reached among any of the negotiating agents. This last state transition brings about a closure to the negotiation. The state transitions, however, do not capture the actual conversations that embody negotiation between agents.

To describe the conversations which take place between agents participating in a negotiation, we use a conversational scheme, as in [6, 10]. All conversations take place within a negotiation state. A state transition results from the culmination of the conversations within the negotiation state. An agent initiates a negotiation by announcing that it would like to adopt a new role, termed a *proposal*. Other agents in the system respond to this proposal and conversation between agents ensues. After repeated conversations, two agents may mutually agree to a particular proposal, and the negotiation ends. If none of the agents reach an agreement, the original proposal expires. The state diagram model shows how the system reacts to negotiation between agents. The conversational scheme, on the other hand, captures the way every agent reacts to negotiation.

We made the following assumptions while designing the simulation.

- The number of agents is equal to the sum of the number of agents required by each role.
- Fitness of an agent for a particular role is computed based on the match of its capabilities

and role requirements and synergy with other agents.
- No new agents are added to or removed from system.
- Every agent has a role at any given time, i.e., no idle agents exist.
- When an agent reaches a role where its contribution of utility to the organization is the highest, it will not consider adopting any other roles.

2. A Model of Roles

We introduce a number of terms.

Definition 1: A *capability* is basic agent ability which is denoted by a value between 0.0 and 1.0.

We assume that there is no decay in the capability of an agent and that in our simulation; an agent can not increase its capability. Furthermore, we assume capabilities are mutually exclusive. Let C denote the set of all capabilities required in the system to perform all the tasks i.e. C= $\{c_1, c_2, c_3, \ldots, c_n\}$. C is the set of all capabilities known by agents. This gives us an n-dimensional space of capabilities. We will call this C-space. Each agent will possess each capability c_i to a different level, denoted by $A(c_i)$.

Definition 2: A *role*, denoted with R_i, is a point in C-space that specifies a capability profile required by any agent to perform in that role.

$R_i = (c_1, c_2, c_3, \ldots, c_n)$. Each role requires a capability c_i at a different level, which is denoted by $R_i(c_i)$.

An agent can perform a role even if its capability is below the capability profile specified by the point for the role in the C-space. The agent's goal is to enact a role that contributes to the global aims of the society [2].

Definition 3: A *department*, denoted by D_i, is a fixed number of agents who are performing the same role R_i.

A department is a way to refer to a group of agents performing the same role. We will use this concept interchangeably in our paper. To adopt a new role indicates the willingness of an agent to leave its current department and move to a new department where all agents are performing the new role that the agent wants to adopt.

Definition 4: An *organization* is a fixed number of departments with an initial assignment of agents to departments.

At any given time, if there are k roles in the organization, than there will be k departments in that organization. An agent is allowed to negotiate for different roles in an organization. If an agent decides to leave its role, another agent from the system must be willing and able to swap its role with the agent. This ensures that all roles are staffed at all times within the organization.

Definition 5: The *productivity* of an agent A, performing in a role R_i, denoted by $P(A, R_i)$, is the sum of the ratios of its capabilities in the capability profile of the role.

$P(A, R_i) = A(c_1)/ R_i(c_1) + A(c_2)/ R_i(c_2) + \ldots + A(c_n)/ R_i(c_n)$

Definition 6: A *synergy network* among a group of agents is a graph among agents where the arcs represent a real value between -1.0 and 1.0 indicating negative or positive influence between pairs of agents. We will use *s(i,j)* as a symmetric function that returns the synergy value between agents i and j.

We will combine the effects of synergy with that of productivity to compute the utility of an agent.

Definition 7: The *utility* of an agent A, performing in a role R_i, denoted by $u(A, R_i)$, is a linear combination of its *productivity* and its *synergy* levels for that role., i.e., $u(A, R) = P(A, R_i) + [(1/\mathrm{sizeof}(R_i)) * \sum s(i, j)]$.

s(i,j) returns synergy value between agents i and j agents. *sizeof*(R_i) returns the number of agents who have adopted role R, and the value $[(1/\mathrm{sizeof}(R_i)) * \sum s(i, j)]$ computes the average synergy value of agent i with all other agents who are performing in role R_i at that time.

Definition 8: The *system utility* of the entire organization, denoted by U, is the sum of the individual utilities of all the agents in the organization.

At any given time, an agent is described in the conversational scheme by specifying its current role. E.g., an agent A with role R_i will be described as $A[R_i]$.

Two negotiation strategies are explored in this paper. The first strategy is called the *local strategy* and is based on the idea of an agents' local maximization of its utility as presented in [4]. The second strategy, called the *rational strategy*, is a cooperative strategy where the agent negotiates with all other agents in the system, and may choose to adopt a new role even though it might not maximize its own utility by doing so.

In the local strategy, an agent proposes to adopt a specific role based on its own comprehension of where it will maximize its contribution to the organization. It is a *greedy* approach where agents are trying to maximize their own utility and are therefore selfish. If we consider this negotiation state S_0 to be where no proposals have been made, then a new proposal moves the system to a new state S_1. The conversational scheme describes the conversation that takes place in this negotiation state. Agent A_1 makes the following proposal:

$A_1[R_i] \rightarrow A_1[R_j]$ "WOULD LIKE TO ADOPT"

This proposal is based on A_1's comprehension that adopting R_j (or moving to Department D_j) is going to bring about the most net positive utility gain for A_1. The arrow reflects the message in between the quotation marks, which in this state implies the willingness of A_1 to adopt role R_j.

This proposal is analyzed by all agents currently performing role R_j and if any agent's utility is going to increase due to a swap, a concession is sent back to the initiating agent, with the agent's corresponding change in utility. This takes the system to a new negotiation state S_2, as shown in Figure 1. At this new negotiation state, the initiator would analyze the concessions and choose to swap with the agent from R_j whose change in utility is the greatest.

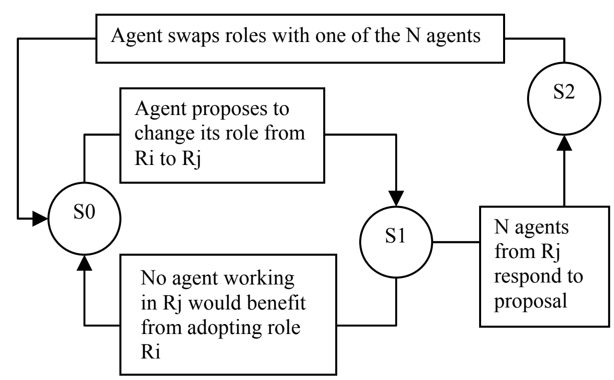

Figure 1: An example of local strategy using the Negotiation State diagram model

In state S_1, if no agents from department with role R_j are willing to swap with the initiating agent, then after a period of wait, the agent's proposal expires and the system returns to state S_0.

The conversational model in state S_1 looks like this:

$A_2[R_j] \rightarrow A_2[R_i]$ "WILL SWAP" (delta(utility))
$A_3[R_j] \rightarrow A_3[R_i]$ "WILL SWAP" (delta(utility))
…
…
$A_m[R_j] \rightarrow A_m[R_i]$ "WILL SWAP" (delta(utility))

The number of responses, m-1, for this proposal is fewer than the total number of agents currently performing in department D_j, performing role R_j, because only those agents respond from department D_j whose utility will increase by adopting R_i. The message here indicates a willingness to swap. Additional information is added which indicates the change in utility for the agent in role R_j if it adopts R_i. All these messages for a swap are sent to the initiating agent which ultimately sends a message back to the agent whose utility change is the highest amongst this group.

$A_1[R_i] \rightarrow A_5[R_j]$ "SWAP"

At this point, the initiating agent and the agent currently working in role R_j will leave their respective roles and adopt each other's role.

All the other agents that had sent a concession to swap with the initiating agent wait for a period and when no answer is received, they retract their concessions. The diagram in figure 2 suggests the messages that are exchanged during negotiation. Agents performing a role R_i are grouped together in a department D_i. Each arrow represents a message sent between agents during negotiation.

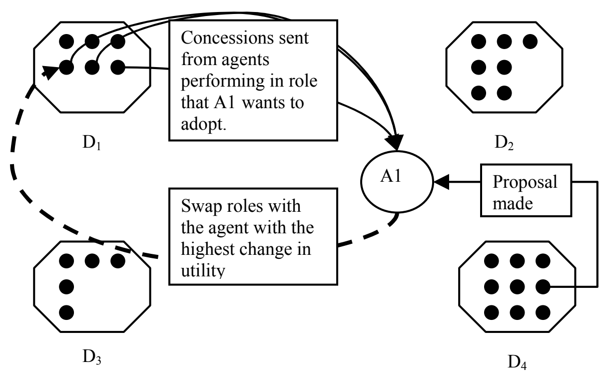

Figure 2: Conversational Scheme for Negotiation during State Transitions for local strategy

In the rational strategy, an agent A_1, does not broadcast a message to agents in just one department but to the entire system that it is willing to leave its current role R_i and move to any other department (thus adopting a new role). This changes the state of the system from S_0 to S_1, as shown in figure 3. The agent then waits for concessions from all the agents in the system whose utility would increase if they adopt role R_i. This makes the system enter state S_2. A_1 then filters out concessions from those departments that will not bring about a net positive gain in utility for A_1 if it chooses to move to that department and perform the role. This makes the system go to state S_3 where A_1 only analyzes concessions from the resulting subset of responses. If there are no agents left after the filtering out process, then the system does not go to state S_3 but instead the initial proposal by A_1 is retracted and the system moves to state S_0. This is different than the local strategy, where once negotiation is initiated by an agent and concessions are received back, the agent must swap role with one of the agents who sent the concessions.

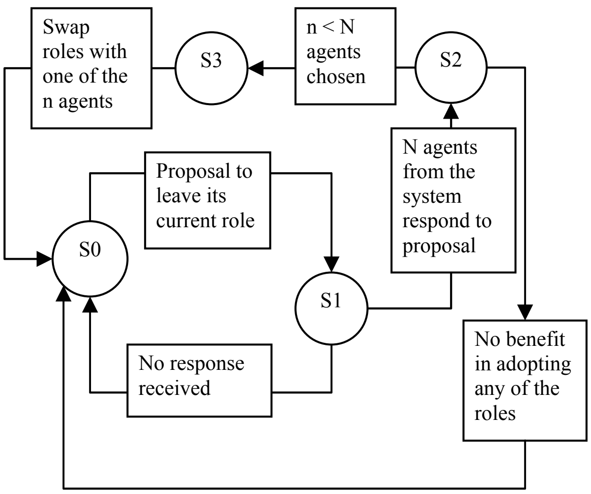

Figure 3: Negotiation State Diagram for the rational strategy

In state S_1, the initiating agent A_1, is expecting to hear from everyone in the system. In this case, the social utility is being maximized [4]. All the agents who will increase their utility by adopting the current role of A_1 respond to this proposal. However, A_1 will only choose to consider responses coming from those departments which will positively affect its utility if it adopts the role being performed in that department. In this case, an immediate message is transmitted to all the agents who will not be considered by A_1.

$A_1[R_i] \rightarrow$ allRespondingAgents(D_j) [SORRY].

Here, *allRespondingAgents(D_j)* is the group of agents from department D_j, performing role R_j, who sent a concession to agent A_1, showing their willingness to swap with A_1. This message no longer binds those agents to wait for a swap with A_1, and instead they can participate in any other negotiation. Another extension to this strategy could be for the initial proposal to only list roles that the initiating agent would be willing to adopt (because of a positive increase in its utility). This will, however, not change the results.

At the end, the initiating agent A_1, chooses to send another message to the agent who has the highest change in utility among the remaining agents just like in the local strategy.

Outcomes of this strategy are *Pareto Optimal*. An outcome is Pareto optimal if and only if there are no other outcomes available which guarantee a higher utility for the system without making any single agent have lower utility [3]. In the global strategy, an agent begins by not being selfish and after hearing all the concessions, continues to negotiate only with those agents that will increase its utility if it chooses to swap roles with them. The initiating agent, however, is not trying to maximize its utility and considers adopting any role that will increase its utility, however small that increment might be.

A third strategy is a global social strategy that will always yield an optimal result. In this strategy, however, agents would have to move to a role that may decrease their utility. Since agents are selfish and only choose to increase their utility, this strategy is not acceptable. All things being equal, pareto efficient solutions are preferred over those that are not [8]. The results of the global strategy, however, help us in determining the deviation of our results from the pareto strategy.

3. Experimentation and Results

Experiments were carried out in a simulated environment. The JAVA programming language was used to code the algorithm. Figure 4 shows the pseudo-code we used for our simulation.

1. Initialize all agents within the organization (with their capabilities) and assign them roles randomly.
2. Initialize the synergy network for all agents.
3. An agent from each role is chosen during each cycle. In the local strategy, it chooses to initiate negotiation if it will maximize its utility by adopting a different role. In the global strategy, the agent initiates negotiation with the entire system.
4. In the local strategy, agents from the role which the initiating agent wishes to adopt send concessions to the agent if their utility will increase due to the swap. In the global strategy, all agents who will increase their utility by adopting the initiating agent's role send concessions.
5. The future and current values for the utility of all the agents involved in the negotiation is calculated.
6. After receiving all responses, the initiating agent either chooses to swap with one of the agents (local and global strategy) or retracts its proposal (global strategy). If no concessions are received, the proposal is retracted also.
7. The most recent change in system utility and the change in entire system utility plotted for each cycle.

Figure 4: Pseudo-code for algorithm used in our simulation

The random association of roles guaranteed a system which had not reached its optimal utility value. Since, we consider the state where no negotiations are placed as S_0, it implies that the system is in state S_0 after all the roles and agents have been initialized in the system.

We first examined the system's behavior when no agents were negotiating, as depicted in Figure 5.

An agent initiated a swap based on its own comprehension that moving to the new role would increase its utility. The agent swapped with a randomly chosen agent from the role it wanted to adopt. The initiating agent tries to maximize its utility but since there is no negotiation among agents, the change in system utility is unpredictable, and can both be positive and negative.

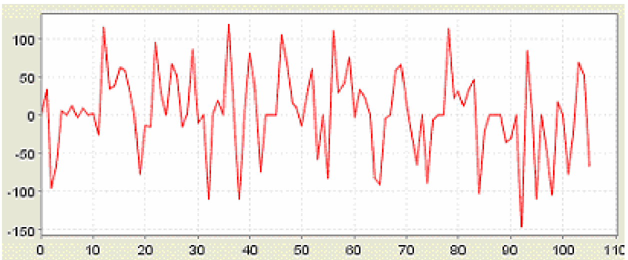

Figure 5: Change in System Utility over time with no negotiation

The results for the local negotiation strategy are depicted in Figure 6 and Figure 7 (by the dotted line). Figure 6 depicts a positive change in utility between each time cycle, and the change is higher in the beginning. The figure also shows that the graph is not continuous implying that sometimes none of the agents are willing to adopt a new role. This does not, however, mean that no agent in the system is initiating negotiation. It just implies that an agent's desire to go to a particular role is not fulfilled because no one from that particular role wants to swap with the initiator.

This negotiation strategy brings about a zero change in system utility fairly quickly. As we can see from figure 7, this is not an optimal level of system utility. In this strategy, negotiation is carried on locally based on the local perspective of the initiating agent. The agent tries to maximize its own utility by choosing to adopt a new role, but since that role might not have agents that will positively benefit from the swap, no agreement is reached and the negotiation is unsuccessful.

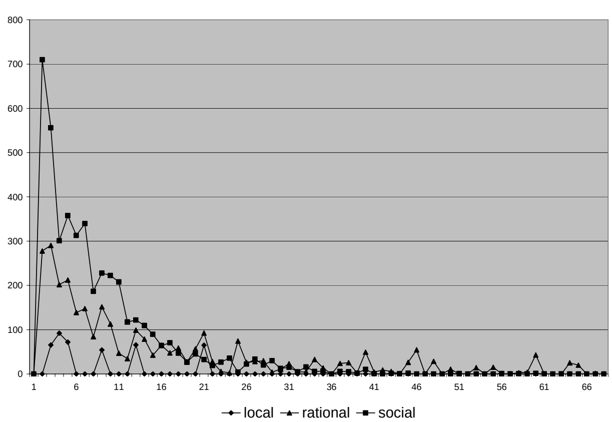

Figure 6: Change in Utility between each time cycle for local, rational and global social negotiation strategies.

In the rational strategy, negotiation is carried among all the agents in the system. The results (figure 7) show a monotonically increasing graph for system utility. An agent from a particular role initiates negotiation. All other agents from the system respond to this proposal. The initiating agent then considers responses from agents currently in those roles that will bring about a net positive gain if it were to swap with those agents. The agent with the highest change in utility amongst this group is chosen by the initiating agent. If a swap is made, it does not guarantee that the new role for either of the two agents is the most optimal for that agent. However, since the entire organization is considered, the change is greater than the one noticed in the local strategy.

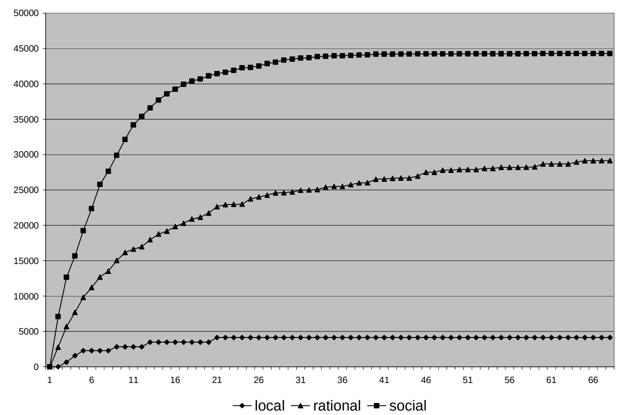

Figure 7: Change in Total System Utility over Time for local, rational and global social negotiation strategies.

The results (figure 7) indicate a sharp increase in system utility in the beginning which flattens out as more and more agents successfully negotiate to adopt roles. The negotiation makes the different roles stable as time passes. This happens because, with time, more and more agents end up in their pareto optimal role. The change in utility graph (figure 6), for the same simulation, indicate large positive changes in the beginning which tone down as time passes. The graph is fairly continuous in the beginning indicating many successful negotiations. As time passes, the graph is less continuous indicating a decrease in the number of successful negotiations because more agents are performing in their pareto optimal role.

The results for the social strategy are also depicted in figure 7. Although this strategy brings us to an optimal level of system utility, it is not a desirable strategy, because it may decrease an agents' utility from its original level during a negotiation cycle.

4. Conclusion

In this paper, we have shown that efficient role adoption by agents in an organization is achieved through negotiation, and that during this process the system performance converges to its pareto optimal level. The agents in the system selflessly adopt roles which maximize change in system utility for that time period. Eventually, the system converges to its pareto optimal level. The paper uses the negotiation state diagram model [4] and the conversational scheme [2] to describe the negotiations.

This paper suggests a desirable negotiation scheme. We will further investigate the rational negotiation scheme to include changes in the organizational structure like, for instance, layoffs, turnovers etc, or changes to agents' perceptions or their capabilities. The main contribution of this paper is to outline a strategy simulating an automated negotiation between agents for role adoption. This, in turn, can decrease cost and time for efficiently distributing talent within an organization. Resource allocation is another place where a strategy like this can be used to allocate resources to a task.

5. References

[1] Bussmann, S., and Muller, H. (1992). A negotiation framework for cooperating agents. In *Proceedings of the CKBS-SIG Conference*, pages 1-17, 1992.

[2] Dignum, V. (2004). The OperA model for Agent Societies. *A Model for Organizational Interaction.* Doctoral Thesis, Utrecht University, Netherlands.

[3] Endriss, U., Maudet, N., Sadri, F., and Toni, F. (2003). On optimal outcomes of negotiations over resources. In *Proceedings Second International Joint Conference on Autonomous Agents and Multiagent Systems, (AAMAS-2003).*

[4] Guttman, R., and Maes, P. (1998). Cooperative vs Competitive Multi-Agent Negotiations in Retail Electronic Commerce. *Proceedings of the Second International Workshop on Cooperative Information Agents (CIA '98)*, July 3-8, 1998, Paris, France.

[5] Hexmoor, H., and Pasupuleti, S. (2003) Towards Institutional versus Interpersonal Influences on Role Adoption. *In Workshop: Representations and Approaches for Time-Critical Decentralized Resource/Role/Task Allocation, Second International Joint Conference on Autonomous Agents and Multiagent Systems. (AAMAS 2003),* July 14-18, 2003, Melbourne, Australia.

[6] Hexmoor, H. and Vaughn, J. (2002). Computational Adjustable Autonomy for NASA Personal Satellite Assistants. *Symposium on Applied Computing (SAC '02)*, March 10-14, 2002, Madrid, Spain.

[7] Kersten, G., and Lo, G. (2001) Negotiation Support Systems and Software Agents in E-Business Negotiations. *The First International Conference on Electronic Business,* December 19-21, 2001, Hong Kong.

[8] Lomuscio, A., Wooldridge, M., and Jennings, N. (2000). A classification scheme for negotiation in electronic commerce. *Agent-Mediated Electronic Commerce: A European Perspective (eds. F. Dignum and C. Sierra), Springer Verlag, 19-33*.

[9] Parsons, S., Sierra, C. and Jennings, N. R. (1998). Agents that reason and negotiate by arguing, Journal of Logic and computation, 8(3):261-292.

[10] Sims, M., Goldman, C., and Lesser, V. (2003) Self-Organization through Bottom-up Coalition Formation. *AAMAS 2003*, July 14-18, 2003, Melbourne, Australia.