Effects of local information on group behavior

Sandip Sen, Shounak Roychowdhury, & Neeraj Arora
Department of Mathematical & Computer Sciences
University of Tulsa
e-mail: {sandip,roy,arora}@euler.mcs.utulsa.edu

Abstract

The performance of individual agents in a group depends critically on the quality of information available to it about local and global goals and resources. In general it is assumed that the more accurate and up-to-date the available information, the better is the expected performance of the individual and the group. This conclusion can be challenged in a number of scenarios. We investigate the use of limited information by agents in choosing between one of several different options, and conclude that if agents are deliberately kept ignorant about any number of options, the entire group can converge faster to a stable and optimal configuration. We also demonstrate how a couple of coalition formation schemes improves the rate of convergence and conclude that a variable, rather than fixed, coalition formation mechanism is more effective.

Introduction

In a distributed multiagent environment the behavior of a group of agents is measured in terms of the performance of agents and the utilization of resources. Researchers in the field of Distributed Artificial Intelligence have studied the effects of local decision-making on overall system performance in groups of both cooperative as well as self-interested autonomous agents (Gasser & Huhns 1989; Huhns 1987). Ineffective system performance can be caused by several characteristics of distributed decision-making: conflicts of interests, contention for resources, asynchronicity in the decision process, lack of centralized control or information, incomplete or incorrect global information, etc.

In this paper, we focus on one particular aspect of distributed decision-making: the effect of limited local knowledge on group behavior. Whereas intuition suggests that agents are equipped to make better local decisions with more complete and correct information, self-interested choices can sometime lead to group instabilities with complete global information. We believe that reducing the amount of information available to such rational decision makers can be an effective mechanism for achieving system stability. The research question that we are asking is the following: Can limited local knowledge be a boon rather than a bane in a multiagent system?

To investigate this issue, we use a resource utilization problem where a number of agents are distributed between several identical resources. We assume that the cost of using any resource is directly proportional to its usage. This cost can be due to a delay in processing of the task in hand, or a reduction in the quality of the resource due to congestion. Hence, there is a justified need for agents to seek out and move to resources with lesser usage. Other researchers have shown that such systems can exhibit oscillatory or chaotic behavior where agents move back and forth between resources (Hogg & Huberman 1991; Kephart, Hogg, & Huberman 1989) resulting in lack of system stability and ineffective utilization of system resources. The case has also been made that the introduction of asynchronous decision making or heterogeneous decision-making schemes can improve system convergence. We see our current work as providing a natural, complimentary mechanism for enabling agents in similar situations to quickly converge to the optimal system state.

Not limited to artificial domains discussed here, we find an analogy of the resource utilization problem within the dynamics of human society. We often observe social trends in human societies where the populace tend to look for opportunities and search for better openings within a closed environment (Bartos 1967). For instance, it is obvious and practical under rational thinking to shift for greener pastures, move for better jobs with less competition, to search for resources with less utilization, etc. The self-interested nature of an individual leads to choices that are perceived to improve rewards from the environment. The theory of migration in social behavior and occupational mobility suggest a dynamic structure, the stability of which depends on how an individual chooses its action based on the prevailing circumstances. Similar to human societies, societies of agents also undergo changes and evolve with time. As agent designers, we are faced with the problem of developing decision mechanisms that allow agent societies to stabilize in states where system resources are effectively utilized. In this paper, we consider agent societies where agents decide on their social mobility based only on their perception of the current state of the world. This assumption of relying only on the current state and ignoring the effects of past history on decision making is also used in Markovian analysis (Howard 1971).

This study attempts to verify the following conjecture: *limited knowledge of the environment can be beneficial for an agent in comparison to complete global knowledge.* We present a decision mechanism to be used by individual agents to decide whether to continue using the same resource or to relinquish it in the above-mentioned resource utilization problem. We show that a spatially local view of an agent can be effectively used in a decision procedure that produces stable allocation of agents to an optimal global state in terms of effective resource utilization. Experimental results show that increasing the information available to an agent increases the time taken to reach the desired equilibrium state. We provide a probabilistic analysis explaining this phenomena. We further plan to study the effects of varying amounts of information on the convergence process of these agent groups.

Related Work

Hogg and Huberman (Hogg & Huberman 1991) have analyzed a resource utilization problem similar to the one mentioned in the previous section to study effects of local decisions on group behavior (Hogg & Huberman 1991; Kephart, Hogg, & Huberman 1989). Kephart *et al.* (Kephart, Hogg, & Huberman 1989) show how system parameters like decision rate can produce stable equilibria, damped oscillations, persistent oscillations, or can burst into a chaotic regime. They also provide an analysis of how agents that monitor system behavior and accordingly adjust their performance can bring the system closer to a stable behavior. Hogg and Huberman (Hogg & Huberman 1991) present a robust procedure for suppressing system oscillations using a reward mechanism based on performance.

We share their motivation of achieving stability in a multiagent system when individual agents are making decisions based on self-interest. However, whereas they are interested in investigating decision procedures that lead to heterogeneity in agent types, we focus our efforts on identifying a simple decision procedure that can be used by all agents but would still lead to stable systems. On another note, we are particularly interested in evaluating the effects of agent decisions based on limited system knowledge on the stability of the system. Thus we have chosen to investigate systems with relatively larger number of resources as compared to others.

We should also clarify that various other forms of heterogeneity including asynchronicity of decision making, different communication delays, different decision algorithms, etc. will help speed up convergence and attain group stability. Our purpose in this paper, is to investigate the conjecture that access to less global information can help agents achieve stability under certain situations. It should be noted that because local information is different for different physically distributed agents, limiting agent decisions to the use of local information only provides another source of heterogeneity in the system.

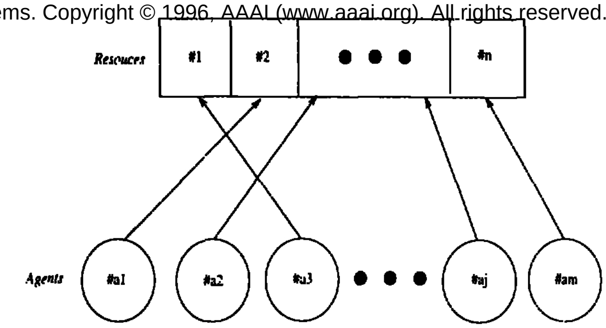

Figure 1: Agents sharing resources.

The Model

In this section we present a simple model of agents sharing a set of identical resources as shown the Figure 1. There are m agents and n identical resources. At any time instant, an agent use only one resource, and over time tries to move to a resource that is less used by other agents. In this study, we show that when an agent has less knowledge about the utilization of each resource in the resource set, the contention for resources decreases and results in quicker convergence to stable resource usage.

At present we model the knowledge of an agent about the resources by using an *r-window*. An r-window is a window through which an agent can see which of the resources among the resource set it should look for before making a decision. At each time step each agent has to make the following decision: whether it should continue to use the present resource or should it move to another one with less utilization. If agent a_k is currently using resource i, then it will consider a move to one of the resources in it's r-window (resource in the vicinity of the current resource).

The model makes a few basic assumptions. We assume that that all resources are equivalent. Moreover, resources are neither introduced nor eliminated during the life time of agents. Similarly all agents remain active and they make their decisions synchronously. All agents retain the same r-window size during the process of decision making. The probability of an agent to shift from the current resource to another resource is inversely proportional to the difference of the usage of these two resource.

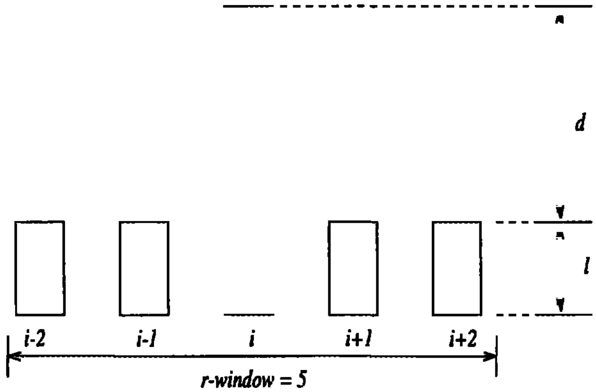

Figure 2: Resource i has d agents more than every other resource in its r-window.

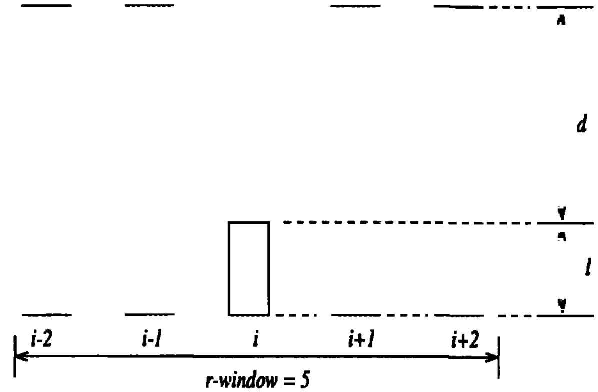

Figure 3: Resource i has d agents less than every other resource in the r-window.

We now discuss the decision procedure we use to determine the resource to be used by an agent in the next time step. It can be shown that a deterministic and greedy decision procedure of choosing the resource with the lowest utilization in the r-window will lead to system oscillations. Hence, we are motivated to use a probabilistic decision procedure. The particular procedure that we use first calculates the probability of moving to each of the resources in the r-window, and then normalizes theses values by the corresponding sum. The probability of an agent that decides to continue to use the same resource i is given by:

$$f_{ii} = \frac{1}{1 + \tau \exp^{\frac{r_i - \alpha}{\beta}}}, \tag{1}$$

where r_i is the number of agents currently using resource i (this is also the utilization of that resource), and τ, α, and β are control parameters. On the other hand, the probability of moving to another resource $j \neq i$ is given by:

$$f_{ij} = \begin{cases} 1 - \frac{1}{1+\tau \exp^{\frac{r_i - r_j - \alpha}{\beta}}} & \text{if } j \in W_i \ \&\ r_i > r_j, \\ 0 & \text{otherwise}, \end{cases} \tag{2}$$

where W_i are the resources accessible to agent using resource i. Now, the probability that an agent a_k occupying a resource i will occupy a resource j in the next time step is given by normalizing the above terms:

$$Pr(i,j) = \frac{f_{ij}}{\sum_j f_{ij}}. \tag{3}$$

Our conjecture for the behavior of the group is: the larger the r-window, the lesser is the stability of the system, and it takes more time to reach an optimal equilibrium state. This slower convergence can be explained by a probabilistic analysis. Consider a resource i which has higher load than the surrounding resources (as shown in the Figure 2). We further assume that n agents are using that resource at a given instance of time. Let X be a random variable corresponding the number of agents who will not leave the resource in the next time step. Therefore, values for X follow a binomial distribution i with probability $Pr(i,i)$. The expected value of X is therefore given by:

$$E[X] = nPr(i,i), \tag{4}$$

and the variance of X is given by:

$$Var[X] = nPr(i,i)(1 - Pr(i,i)). \tag{5}$$

Similarly, as the Figure 3 shows, the resource i is being less utilized when compared with its neighbors. Obviously there will be a tendency of an agent who is currently not using i to move to resource i. Let Y be the random variable corresponding to the number of agents who will move into resource i in the next time step. Therefore values for that Y follow a binomial distribution with the probability $\sum_{j \neq i} Pr(j,i)$. We can also think of Y as a sum of several independent binomially distributed random variables, Y_{ji}, where Y corresponds to the number of agents who will move into resource i from resource j in the next time step. Y_{ji} has an expected value of $nPr(j,i)$ and a variance of $nPr(j,i)(1 - Pr(j,i))$. Therefore, the expected values of Y is given by:

$$E[Y] = \sum_{j \neq i} nPr(j,i). \tag{6}$$

And the corresponding variance is:

$$Var[Y] = \sum_{j \neq i} nPr(j,i)(1 - Pr(j,i)). \tag{7}$$

Let us now analyze the implications of these analysis. Figure 4 plots the expressions in (4) and (5) for different d values and different r-window sizes. Figure 4 reveals a very interesting phenomena. For large window sizes, the variance of the number of agents staying in the resource decreases as the difference between the utilization of the current resource usage and

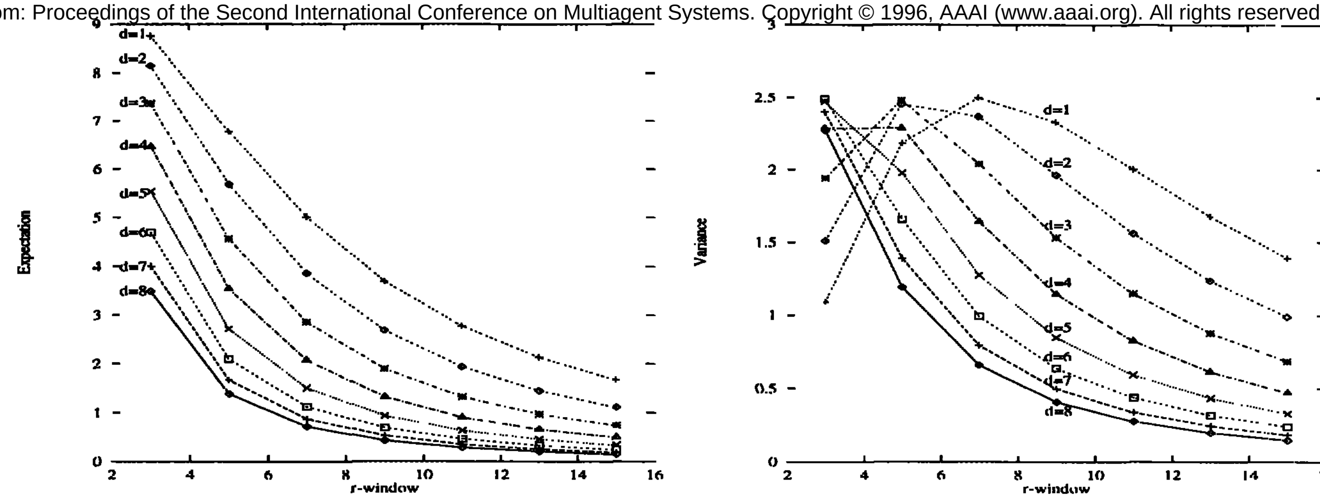

Figure 4: Expectation and variance of an agent staying in the current resource (corresponding to Figure 2), and $l + d = 10$.

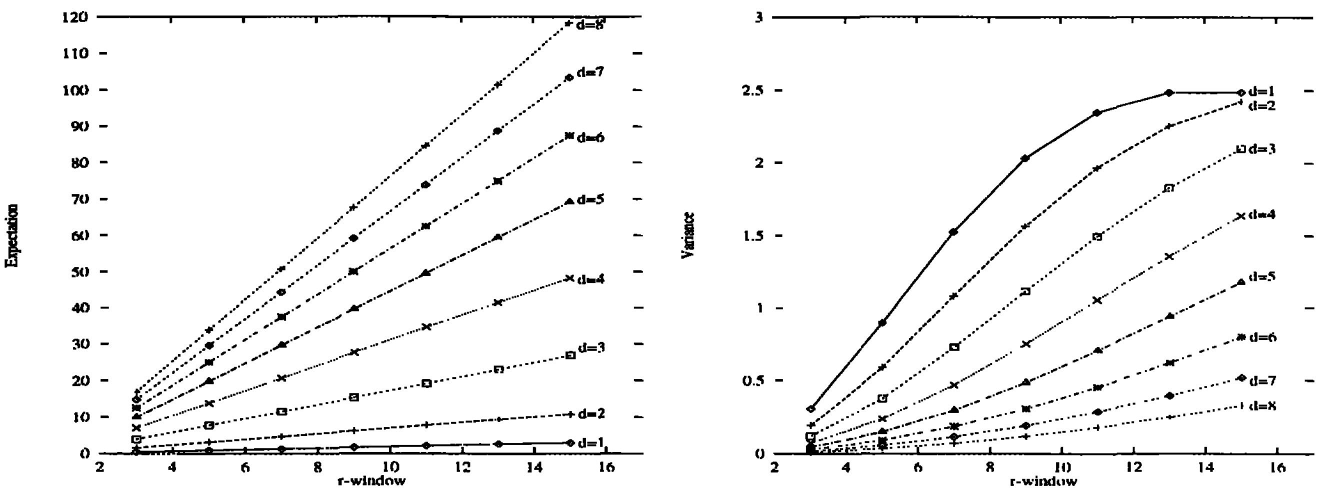

Figure 5: Expectation and variance of an agent moving to the less used resource (corresponding to Figure 3).

the neighboring resource usages decreases. This means initially the agents will quickly spread out, but later it will have difficulty to converge when all resources have roughly the same utilization. At this point high variance can again cause some imbalance in the resource usage. The situation is precisely the opposite for small window sizes: here, the variance decreases with the decreasing difference between the current and the neighboring resources. This means that there will be a relatively slower convergence towards a balanced distribution of agents to resources, but there is a continuing pressure towards more uniform distribution of agents to resources. This process is further helped by a greater inertia of moving of the current resource at smaller r-window sizes as seen from expected number of agent plot in Figure 5. A similar phenomena is observed in Figure 5 where we consider the variance in the number of agents coming to a resource which is less utilized than the neighboring resources. These two figures give a more formal explanation of the faster convergence with smaller windows. We are currently performing a more detailed analysis of this phenomena.

Results

We assume that the resources are arranged in a ring and each agent knows the number of agents using the resource it is using and the neighboring resources within the r-window to the left and right. Each time step consists of all agents making a decision regarding which resource to use next. In Figure 6 we present experimental results with 27 agents using 9 resources. The data for these plots are averaged over 10 random initial assignments of agents to resources. Starting from r-window size of 3, as we increase the size of the

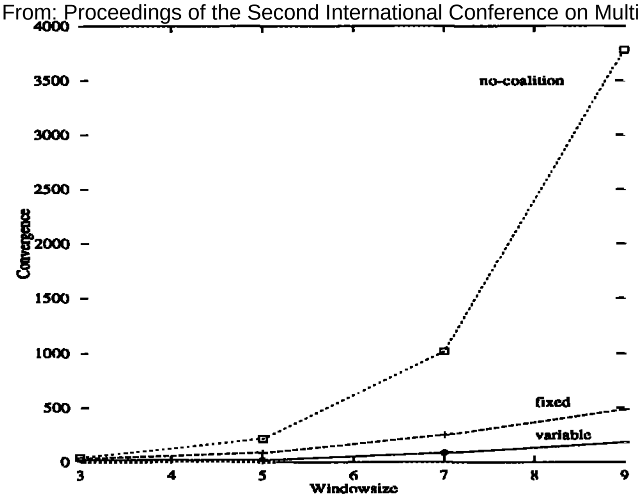

Figure 6: Number of steps to convergence for different R-window sizes.

window to 9, we observe that the system takes much more time on the average to stabilize. Figure 7 presents the number of agents occupying resource 1 at different time steps with r-window sizes of 3, 5, 7, and 9 respectively. These figures confirm our experimentation that together with taking more time to converge, the variation in the number of agents occupying a given resource is higher with the larger window size.

Our initial experiments, therefore, suggest that agents converge to a stable state (which is also optimal because agents are equally distributed among the resources) in less number of time steps when they have relatively less global information.

Forming coalitions

In the previous section, we observed that agents with a limited view of global scenario converged faster to optimal states. However, this work was based on the assumption that every agent made an individual decision which was based on the current resource utilizations. The results showed that in cases where the window-size is large the system took significantly longer to converge (in some cases the system did not converge even after a large number of time steps). One reason which we attributed to this delayed convergence was that the individual agent had no information about the decision of the other agents. As a result, all the agents tried to move towards the least utilized resource within their view thus letting the previously under-utilized resource to become over-utilized in the next time step and vice versa.

We conjectured that some of the convergence problems mentioned above can be alleviated by forming coalitions of agents, where agents belonging to a given coalition will cooperatively decide on their next move. For example, within any such coalition, agents may take turns in selecting which resource they are going to occupy in the next time step and then inform other agents in the coalition about that decision. Thus, agents will have more up-to date and accurate information about the likely resource usages in the next time step, and hence are in a position to make a more effective movement decision. In the extreme case, if all agents form one coalition and the R-window included all resources, each agent will have a complete and correct global information at all times, and the system will immediately converge if each agent moves to the least used resource at the time it makes its movement decision.

We studied two modes of forming coalitions: in the first mode agents were randomly partitioned into equal-sized coalitions before the start of the simulation and no agents ever changed coalitions (we use a coalition size of 5); in the second form, agents occupying the same resource at any given time formed a coalition and hence coalitions changed from one time step to the next. In both the groups, an individual agent's decision of moving to a resource is not only based on the current utilization of the resources within its view window but is also guided by the actual status of that resource after all the other agents in its group have decided to move to a particular resource.

We ran experiments for both these coalition types by varying the window size and keeping the number of agents and resources constant. The results of these experiment averaged over 10 runs are shown in the Figure 6. The convergence patterns with two types of coalitions were very similar to the convergence pattern with no coalitions, i.e., increasing the window still resulted in slower convergence. Runs with coalitions, however, converged faster than runs with no coalitions. This was particularly true for larger window sizes where runs without coalition often took an extremely long time to converge. In fact, we believe for larger window sizes (will require more resources too) and number of agents, the system may not converge if some form of coalitions are not used.

When comparing the performance of two coalition types, we find the variable coalitions converge faster than fixed coalitions. This observation can be explained by two reasons:

- Agents belonging to a static coalition may be dispersed over all the resources at any given point in time. So, the movement decision of any one such agent may not impact all the other agents in the coalition (the agent may be moving from and to resources both of which may be outside the window of some of the other agents). Hence, only some of the information that is shared among the coalition members is useful. On the other hand, in the variable coalition case, movement decisions of any one agent impacts every other agent in the coalition. Thus, for same sized coalitions, agents in variable coalitions take more informed decisions compared to agents in fixed coalitions.

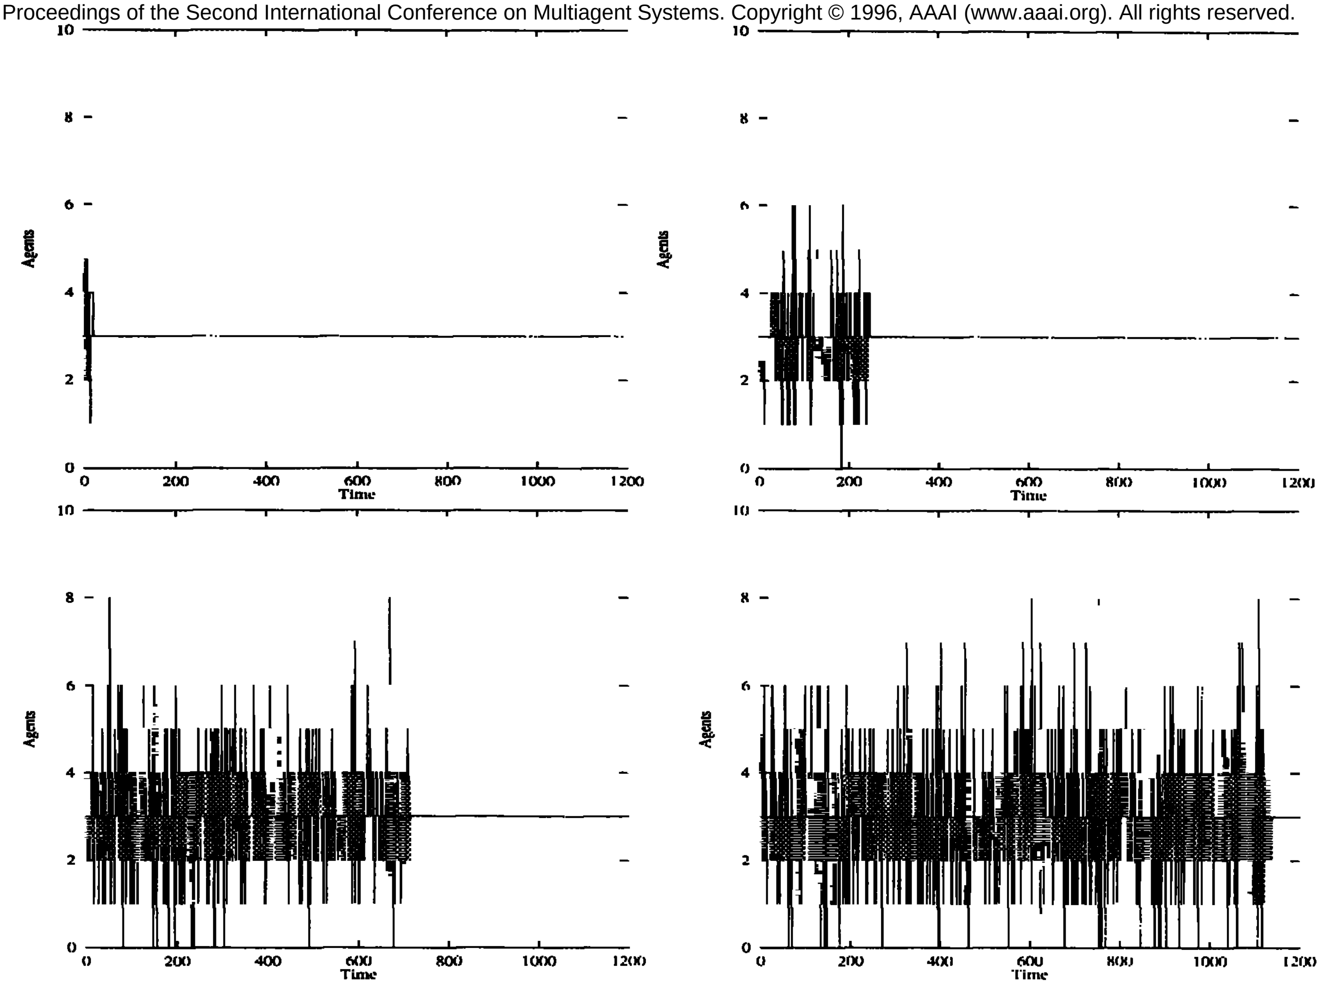

Figure 7: System convergence with 27 Agents and 9 resources. From top left in clockwise order we have R-windows of 3, 5, 9, and 7 respectively.

- The size of fixed coalitions is determined *a priori*, whereas the size of variable coalition dynamically changes. The larger the load on a resource, the larger is the size of the corresponding variable coalition, and the more informed is the decisions made by corresponding coalition members. Hence, our proposed variable coalition formation scheme allows agents access to more information precisely when it is critical. This allows variable coalitions to converge faster.

The more general lesson from this set of experiments is that in order for agents to be flexible to changing environmental demands, it is more appropriate to provide a coalition formation and dissolution mechanism that utilizes current problem loads and inter-relationships between agents. As these critical factors change over time, it is often myopic to pre-assign the coalition to which an agent should belong over its lifetime.

Discussions

To "bury the head in the sand" and ignore most of the information (in this case of using a small *r-window*) does not appear to be a sound principle in general. However, to observe what neighbors are doing may be good precept, but to base our decisions closely on what is happening anywhere in the whole wide world can be misleading at times, and can be detrimental in specific circumstances. One can easily find the effectiveness of such principles in daily chores of our lives. To name a few: a visit to a ticket counter, which highway to take to work, computational jobs waiting in various queues for their turn to get processed, etc. Similarly, we believe that a homogeneous agent society utilizing a set of limited resources might be able to utilize their resources efficiently by avoiding complete knowledge about the entire set of resource. Analyzing the data from these experiments suggests some further investigations on the interplay between limited global knowledge and group stability. We discuss some of our planned experiments below:

Adaptive agents: Counter to our normal expectations we have shown that it may be detrimental to search far and wide for the best opportunity if everyone is doing the same. In retrospective, an intelligent agent may adopt an adaptive policy of increasing its information window until it senses an instability in the system (finds itself jumping continuously from one option to another). At that point it may be prudent to reduce the information window. Adaptive policies may use specific dynamic learning strategies to handle specific problems. For instance, resource utilization problem might use a learning strategy which might not be suitable in a problem where communication is a critical.

Graded movements: We can also model agents with graded inertia of rest. These agents prefer to shift to a nearer resource with less utilization rather than to a more distant resource with negligible utilization. A more uniform treatment of this approach would be to add a notion of *stability* to the probability calculation, i.e., the further off a resource is located from the current resource, the less will be the probability of making the move given the same difference in resource utilizations. Agents may have large window size, but is more and more reluctant to move further away from its current choice. This mechanism assumes a distance metric between choices. A simple extension to equation (1) can be shown as follows:

$$f_{ij} = \begin{cases} 1 - \frac{1}{1+\tau \exp^{(r_i - r_j) * \delta_{ij}}} & \text{if } j \in W_i \ \& \ r_i > r_j, \\ 0 & \text{otherwise,} \end{cases} \tag{8}$$

where δ_{ij} is the distance between resource i and resource j.

Conclusions

In this study we investigated the problem of resource utilization and global performance based on limited local information. The agents with a limited view of global scenario converged faster to optimal states. We provide a probabilistic analysis that sheds some light on this interesting phenomenon. We argued in favor of dynamic, rather than static, coalition formation mechanism to improve system performance. We also identified future avenues of work that will produce adaptive agents which perform more effectively than agents with static strategies.

Acknowledgments: This work has been supported, in part, by an NSF grant IRI-9410180.

References

Bartos, O. J. 1967. *Simple Models of Group Behavior.* New York, NY: Columbia University Press.

Gasser, L., and Huhns, M. N., eds. 1989. *Distributed Artificial Intelligence*, volume 2 of *Research Notes in Artificial Intelligence.* Pitman.

Hogg, T., and Huberman, B. A. 1991. Controlling chaos in distributed systems. *IEEE Transactions on Systems, Man, and Cybernetics* 21(6). (Special Issue on Distributed AI).

Howard, R. A. 1971. *Dynamic Probabilistic Systems, Volume I: Markov Models.* New York, NY: John Wiley & Sons, Inc.

Huhns, M., ed. 1987. *Distributed Artificial Intelligence.* Morgan Kaufmann.

Kephart, J. O.; Hogg, T.; and Huberman, B. A. 1989. Dynamics of computational ecosystems: Implications for DAI. In Huhns, M. N., and Gasser, L., eds., *Distributed Artificial Intelligence*, volume 2 of *Research Notes in Artificial Intelligence.* Pitman.