Using limited information to enhance group stability

Sandip Sen, Neeraj Arora, & Shounak Roychowdhury
Department of Mathematical & Computer Sciences
University of Tulsa
600 South College Avenue
Tulsa, OK 74104-3189, USA
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 them about local and global goals and resources. In general it is assumed that the more accurate and comprehensive 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 kept ignorant about, or they deliberately ignore, any number of options, the group can converge faster to a stable and optimal configuration. We present a probabilistic analysis that sheds light on the observed phenomenon of quicker system convergence with less global information. This analysis suggests a desirable adaptive behavior on the part of individual agents. Experiments with agents following these adaptive behavior exhibits faster convergence. We also demonstrate how a couple of coalition formation schemes can improve the rate of convergence. A variable coalition formation mechanism is found to be more effective than a static one.

Running title: Information effects on group stability

1 Introduction

In a distributed, multiagent environment the behavior of a group of agents is evaluated in terms of the performance of agents and the utilization of resources. Researchers in the field of Distributed Artificial Intelligence (DAI) 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 and 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 at times 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 (Sen et al., 1996). 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 continually move between resources (Hogg and Huberman, 1991; Kephart et al., 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. Researchers have observed 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.

The rest of the paper is organized as follows: the Section 2 briefly states a related work in this area. In Section3 we present the multiple resource utilization problem and the decision procedures used by the agents. Some initial experimental results with this decision procedure is presented in Section4. In Section 5 we present a probabilistic analysis explaining the observed experimental results. In Section 6 we use insights from this probabilistic analysis to develop an adaptive decision mechanism that improves system performance. In Section 7 we discuss the effectiveness of two coalition formation schemes that can further improve system performance. A discussion on the salient features of the model and possible extensions are presented in Section 8.

2 Related Work

Durfee (Durfee, 1995) has argued for the use of sufficient but not complete information about its environment. But, in general, little work has been done to investigate the benefits of limiting information access by agents.

Hogg and Huberman (Hogg and Huberman, 1991) have analyzed a resource utilization problem similar to the one used here to study the effects of local decisions on group behavior (Hogg and Huberman, 1991; Kephart et al., 1989). Kephart *et al.* (Kephart et al., 1989) show how system parameters like decision rate can produce stable equilibria, damped oscillations, persistent oscillations, or can lead the system into a chaotic phase. 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 and 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

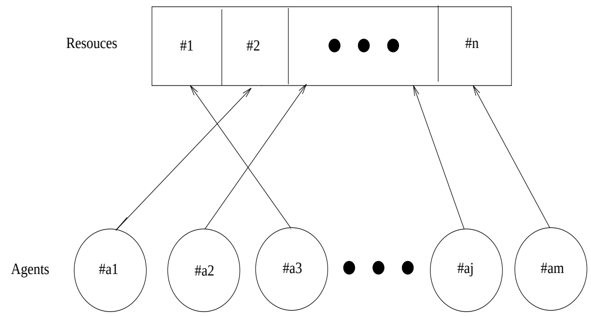

Figure 1: Agents sharing resources.

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.

3 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 (in general, $m > n$). 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 observe the load on some of its neighboring resources. 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 resource in its r-window with less utilization.

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 the agents. All agents remain active and they make their decisions synchronously.

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 probability that an agent will shift from the current resource to another resource is inversely proportional to the difference of the usage of these two resource. 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 or load on 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 within the r-window of an 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 as follows: the larger the r-window, the more will be the contention for the lesser used resources at any given point in time. This, in turn will lead to lesser in the system, which will take more time to reach an optimal equilibrium state. We now present results from some initial experiments we ran to verify this conjecture.

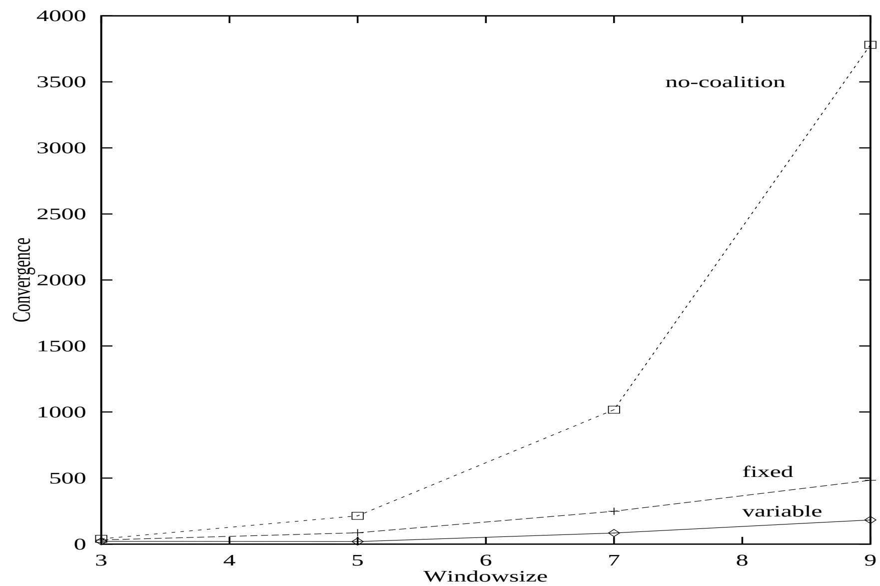

Figure 2: Number of steps to convergence for different r-window sizes and with or without the use of coalitions.

4 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 2 we present experimental results with 27 agents using 9 resources. For now, consider the plot labeled "no-coalition" (the other plots will be discussed in Section 6). 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 window to 9, we observe that the system takes much more time on the average to stabilize. The system can only converge to the optimal state with uniform distribution of agents to resources. This is the only state for which the probabilities for any agent moving out of any resource is zero. It is clear from the figure that increasing the window size leads to considerable increase in the time taken by the system to converge to a stable state.

Figure 3 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

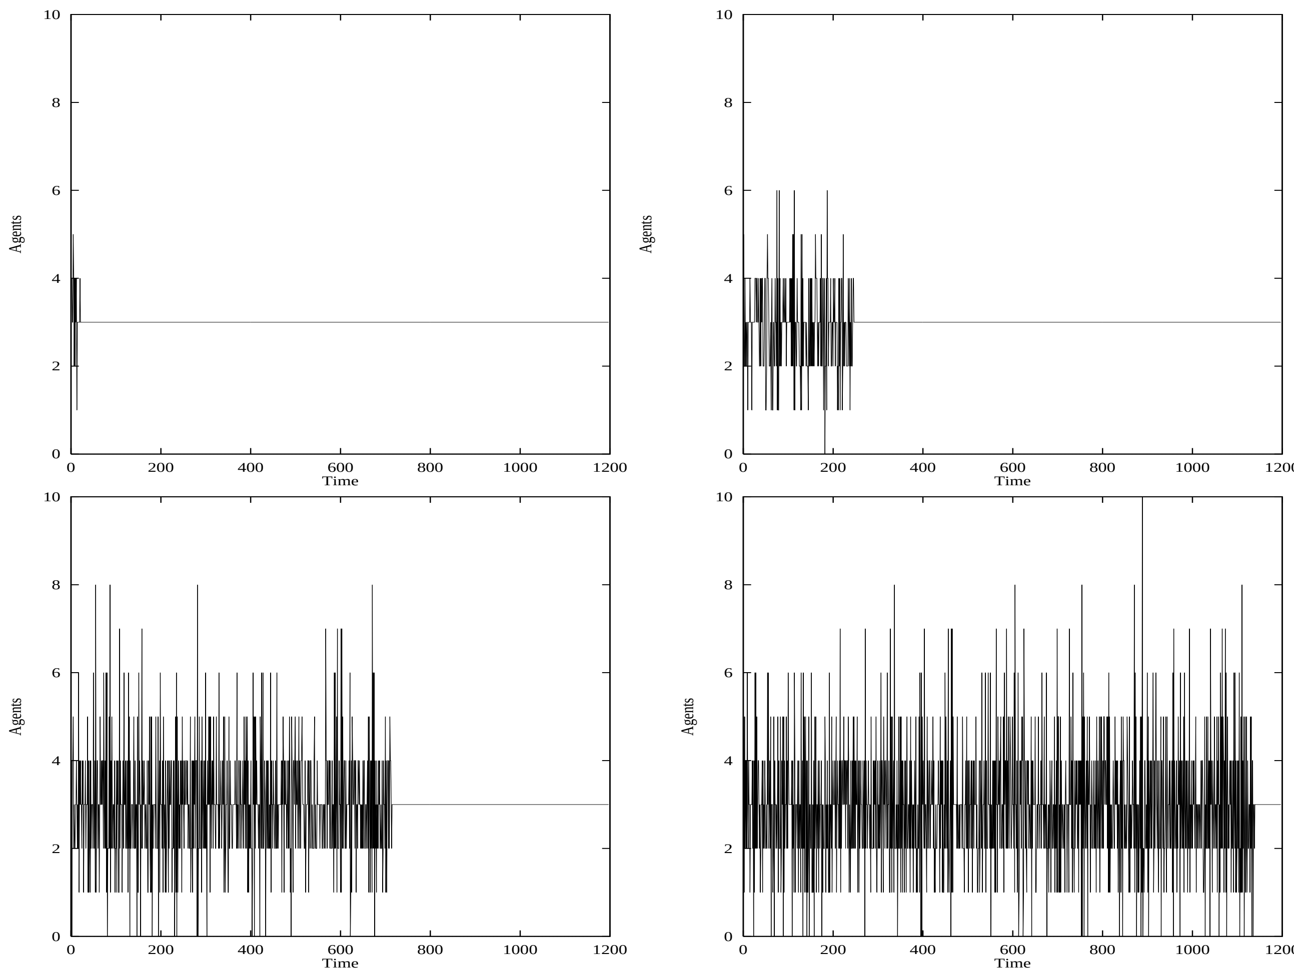

Figure 3: System convergence with 27 Agents and 9 resources. From top left in clockwise order we have R-windows of 3, 5, 9, and 7 respectively.

a given resource is higher with larger window sizes. Our initial experiments, therefore, suggest that agents converge to a stable, optimal state in less number of time steps when they have relatively less global information.

5 Probabilistic analysis

We now present a probabilistic analysis of the behavior of the system to explain slower convergence of the system when agents use larger r-windows, i.e., have more information available locally to base their decisions on.

Consider a resource i which has higher load than the surrounding resources (as shown in

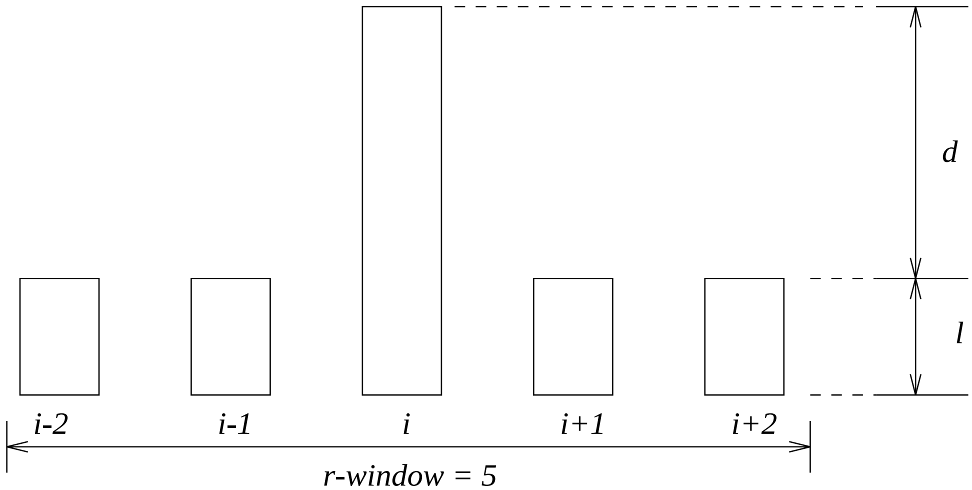

Figure 4: Resource i has d agents more than every other resource in its r-window.

the Figure 4). 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 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 5 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 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_{ji} 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:

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

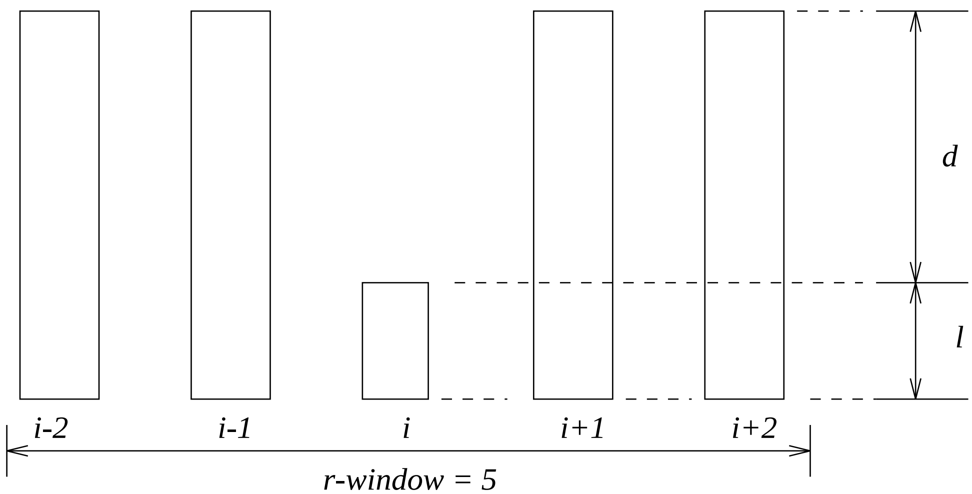

Figure 5: Resource i has d agents less than every other resource in the r-window.

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 6 plots the expressions in Equations 4 and 5 for different d values and different r-window sizes. The plot for expectations of X confirms the intuition that with larger window sizes and larger d values (difference between current resource load and the loads on the neighboring resources), a larger number of agents are expected to move out of the current resource. But figure 6 also reveals a very interesting phenomena. For large window sizes, the variance of the number of agents staying in the resource increases as the difference between the loads of the current and neighboring resource decreases. The two plots together can be used to draw the following conclusion: initially the agents will quickly spread out from a highly utilized resource to neighboring, less utilized resources. But when all resources have approximately the same load, high variance will probably cause some imbalance in the resource usages leading the system back towards the initial situation. This kind of behavior can go on in a cyclical manner for some time.

The situation is precisely the opposite for small window sizes: here, the variance decreases with the decreasing difference between the current and the neighboring resource loads. This means that even though there is a relatively slower convergence towards a near-uniform distribution of agents to resources (as inferred from the expectations plot), there is a continuing

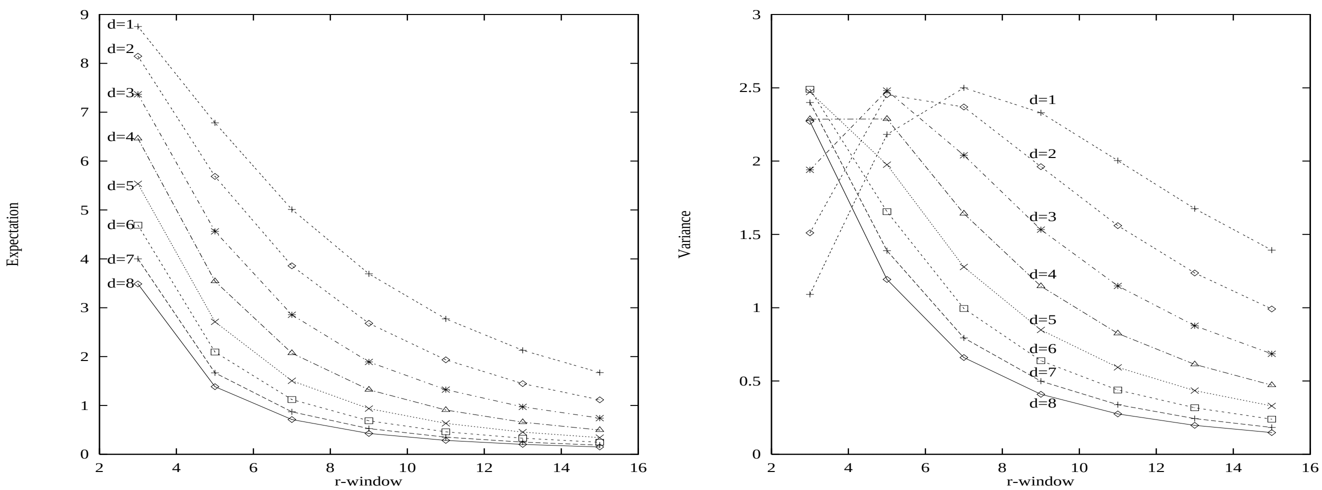

Figure 6: Expectation and variance of an agent staying in the current resource (corresponding to Figure 4), and $l + d = 10$.

pressure towards a completely uniform distribution of agents to resources. This process is further helped by a greater inertia of moving out of the current resource at smaller r-window sizes as seen from expectation plot in Figure 7 (this plot corresponds to agents moving in to a lesser loaded resources from surrounding highly loaded resources).

Figure 7 plots the expressions in Equations 6 and 7 for different d values and different r-window sizes. The variance plot in this figure also supports a similar trend. Here, for all window sizes, the variance of the number of agents moving to the less loaded resource increases with decreasing difference between the loads of the less loaded resource and the surrounding resources. However, the increase in variance for small window sizes is much less compared to when large window sizes are used. This means that when the system comes to a near uniform agent distribution to resources, larger instabilities can be introduced when using larger window sizes. Figures 6 and 7, therefore, help us develop a more more formal explanation of the faster system convergence with smaller windows.

6 Adaptive agents

The probabilistic analysis of agent movements in Section 3 suggests a possible improvement in agent behaviors over the use of fixed r-window sizes that we have seen so far. We will

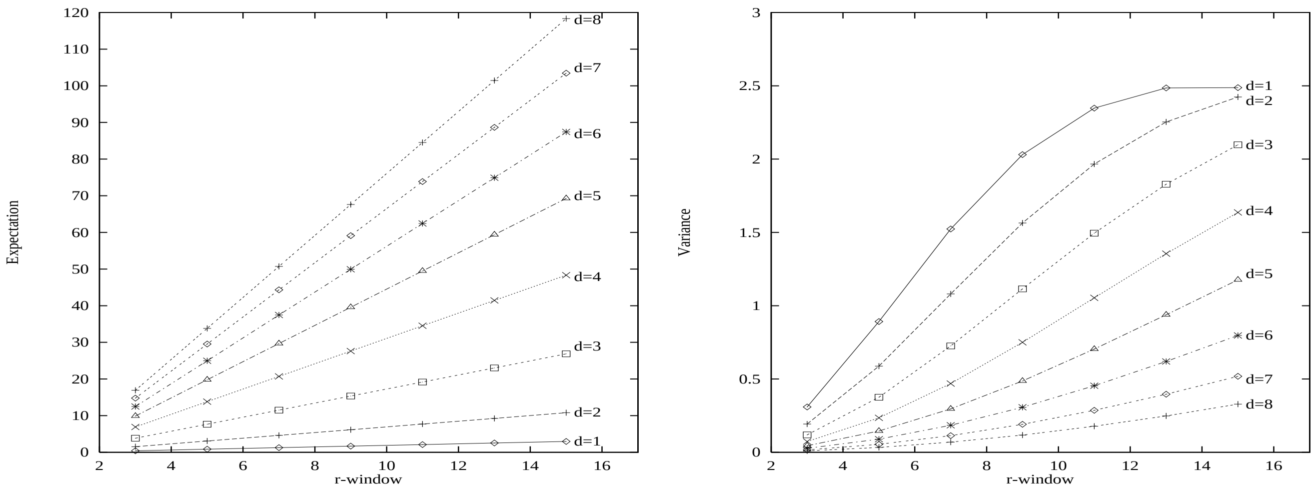

Figure 7: Expectation and variance of an agent moving to the less used resource (corresponding to Figure 5), and $l + d = 10$.

briefly revisit the analysis here to precisely identify the scope for improvement. From Figure 6 we observed that for large r-windows, the variance of the number of agents staying in the resource increases as the difference between the utilization of the current resource usage and the neighboring resource usages decreases. The opposite is the case for small r-windows. Also, the expected value of the number of agents leaving the resource more loaded than its neighbors is higher for large rather than small r-windows. Hence, if the initial distribution of agents to resources had a marked imbalance and if the agents were using a large r-window, agents will quickly spread out to a near uniform distribution. At this point, the discrepancy of usage between neighboring resources will be much less than it was at the initial state. Now, from Figure 6, agents will be better served to use a small, rather than large, r-window to reduce the variability in their movements in the next time steps (which is equivalent to a reduction in the variance in the occupancy of the associated resources). Therefore, a likely effective adaptive agent strategy would be to initially use a large r-window size, but quickly reduce this size after some initial movements.

Admittedly, our analysis is based on idealized scenarios in Figures 4 and 5, where all but one resource have the same occupancy. Our conjecture is that the analysis will still apply, with some loss of accuracy, to more realistic situations where resource occupancies are more graded. To verify our conjecture, we ran some experiments with the following adaptive agent decision mechanism for dynamically selecting an r-window size: *agents initially start with an r-window which includes all resources; any time an agent changes resources consecutively for k time steps, it narrows its r-window size down to 3 (i.e., from thereon the agent considers only the load of the current and neighboring resource when making a movement decision).* The motivation for this strategy is that if an agent continually jumps from one resource to another, the system is unstable. Assuming that using too much global information is the cause of this instability, the agent decides to focus on its current neighborhood.

In a set of exploratory experiments we observed best results when we used $k = 1$ in the above-mentioned decision mechanism, i.e., each agent reduced the window size the first time it moves from the resource that it initially occupied. Subsequently, we ran experiments using both adaptive and static window sizes and with 27 agents and 9 resources. In our prior experiments the initial allocation of agents to resources was obtained by using a uniformly distributed random number generator. This resulted in almost uniform resource usage to begin with. However, we are more interested in evaluating agent strategies when the initial distribution is particularly skewed. We generated 10 random scenarios each for skewed and uniform initial distributions of agents to resources. For each initial distribution, experiments were run with 10 random seeds of the random number generator (the agent algorithms use random numbers to make probabilistic choices). Table 1 presents results from these set of experiments. The table clearly demonstrates the effectiveness of our proposed adaptive strategy for choosing r-window size over the best static-window size choice (always using a window size of 3).

Initial distribution	Window size = 3	Adaptive window sizing	Improvement
Skewed	47.62	37.8	21%
Uniform	45.74	44.42	3%

Table 1: Average time steps taken to convergence with adaptive and static window sizes. Experiments involved 27 agents and 9 resources.

It is instructive to note that while the improvement of the adaptive scheme over the static choice is remarkable when the initial agent distribution is skewed, the corresponding improvement for uniform initial agent distribution is minimal. But this observation lends further credence to the accuracy of our probabilistic analysis, which was developed with skewed distributions.

7 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 assumes that agents independently make decisions based on observed 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 for this delayed convergence is that the individual agents 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 mode, 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 movement decision is not only based on the current utilization of the resources within its r-window but is also guided by the actual status of those resources after some of the other agents in its group have decided on their moves.

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 2. 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 extremely long times to converge. In fact, we believe that for even larger window sizes (this 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.

- 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 are the decisions made by corresponding coalition members. Hence, our proposed variable coalition formation scheme allows agents information about decisions made by other agents 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 likely to be myopic to pre-assign the coalition to which an agent should belong over its lifetime.

8 Conclusions and Future Work

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. It appears that strategic, limited usage of local knowledge may be an effective way to address stability issues in multiagent systems.

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:

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:

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

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

Choice versus information: In this paper, we have assumed that an agent moves only to a resource within its r-window. However, r-window is only used as a filter to restrict the information available to agents. One can imagine "adventurous" agents deciding to move to resources for which they have no current load information. This separation of choice from available information will be an interesting point to study and we plan to investigate scenarios where agents perceive information only from their r-windows, but are free to choose any resource to move to.

Acknowledgments:

This work has been supported, in part, by the National Science Foundation under a Research Initiation Award IRI-9410180 and a CAREER award IRI-9702672.

References

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

Durfee, E. (1995). Blissful ignorance: Knowing just enough to coordinate well. In *First International Conference on Multiagent Systems*, pages 406–413, Menlo Park, CA. AAAI Press/MIT Press.

Gasser, L. and Huhns, M. N., editors (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*. John Wiley & Sons, Inc., New York, NY.

Huhns, M., editor (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., editors, *Distributed Artificial Intelligence*, volume 2 of *Research Notes in Artificial Intelligence*. Pitman.

Sen, S., Arora, N., and Roychowdhury, S. (1996). Effects of local information on group behavior. In *Proceedings of the Second International Conference on Multiagent Systems*, pages 322–329, Menlo Park, CA. AAAI Press.