The Bases of Effective Coordination in Decentralized Multiagent Systems*

Sudhir K. Rustogi and Munindar P. Singh**

Department of Computer Science
North Carolina State University
Raleigh, NC 27695-7534, USA

skrustog@eos.ncsu.edu, singh@ncsu.edu

Abstract

Coordination is a recurring theme in multiagent systems design. We consider the problem of achieving coordination in a system where the agents make autonomous decisions based solely on local knowledge. An open theoretical issue is what goes into achieving effective coordination? There is some folklore about the importance of the knowledge held by the different agents, but the rest of the rich agent landscape has not been explored in depth. The present paper seeks to delineate the different components of an abstract architecture for agents that influence the effectiveness of coordination. Specifically, it proposes that the extent of the choices available to the agents as well as the extent of the knowledge shared by them are both important for understanding coordination in general. These lead to a richer view of coordination that supports a more intuitive set of claims. This paper supports its conceptual conclusions with experimental results based on simulation.

1 Introduction

The coordination of agents is a crucial problem in the study of multiagent systems. Consequently, the challenge of understanding the various bases of coordination is important. Although a number of strategies have been considered and applied on a variety of problems, there is little domain-independent agreement on the phenomena that affect coordination. A particularly interesting class of coordination problems arises in multiagent systems in which the decision processes are fully decentralized. Each agent decides its actions purely locally.

A number of interesting research questions arise in this context. In particular, we address the following questions, which emerge at the interface between agent theory and architecture.

---

* We are indebted to Sandip Sen for explaining his previous efforts on this topic, and to Jie Xing for useful discussions. We would also like to thank the anonymous reviewers for comments on a previous version.

** Munindar Singh is supported by the NCSU College of Engineering, the National Science Foundation under grants IRI-9529179 and IRI-9624425 (Career Award), and IBM corporation.

– What are the main concepts involved in achieving coordination in decentralized, i.e., locally autonomous, multiagent systems?
– What are the trade-offs involved in terms of these concepts from the standpoint of achieving coordination effectively?

The answers to the above questions are, inevitably, interleaved. Also, since multiagent systems are a new area of investigation, we follow Simon's advice to study carefully designed simulations to develop a clearer understanding of the theoretical concepts [11, p. 15].

Knowledge is a key component of several abstract agent architectures, e.g., the family of BDI architectures [5, 13]. The MAS folklore identifies the importance of the relationship of knowledge and coordination [3]. Previous studies indicate informally that knowledge helps, but the notion of knowledge is not formalized or quantified in an obvious manner. Sen *et al.* recently introduced a simple experimental setup in which coordination arises among agents (optimally) exploiting shared resources [10]. The agents decide locally, and coordination corresponds to their achieving equilibrium. Sen *et al.* argued that, contrary to what one might naively believe, giving the interacting agents additional knowledge can cause their coordination, i.e., the achievement of equilibrium, to slow down.

1.1 Key Concepts

Thus, it appears that the traditional answers to our two research questions are (a) only knowledge—howsoever formalized—matters for coordination, and (b) the trade-offs involving knowledge are not universally agreed upon.

We find both the above answers intuitively unsatisfactory. First, we believe that knowledge is not the only relevant concept influencing coordination. The following concepts are also potentially important. (We describe these terms technically below.)

– The inertia that the agents exhibit in updating their decisions in response to changes in the state of the world brought about by others' actions. A system whose agents has low inertia may exhibit chaotic behavior, and never achieve coordination.
– The choices that are available to the agents. Too many choices may also lead to chaotic behavior.
– The amount of shared knowledge among the agents. If the agents follow a homogeneous strategy, shared knowledge would tend to lead to similar decisions by all. Similar decisions could lead to more or less effective coordination depending on whether the setting requires the same or complementary decisions. In general, complementary decisions are more interesting, because they cannot be hardwired in some trivial mechanism.
– The extent of the precision in the coordination required. Potentially, the above factors may have a different kind of influence on effectiveness depending on whether we were considering coarse-grained coordination.

When these concepts are factored in, we obtain a richer understanding of the terrain of coordination.

1.2 Main Results

We developed an experimental framework that generalized over the one used by Sen *et al*. Whereas they considered only knowledge (which we find is coupled in their setup with choice), we considered the other important concepts mentioned above. When the enhancements were eliminated, we did indeed achieve results similar to those of Sen *et al.*, but in light of our more extensive exploration, were forced to different conclusions.

When we increased the choices available to an agent independently of its knowledge, we found as we had suspected that it took longer and longer to converge. More choices lead the agents to coordinate slowly.

However, we found that holding the extent of the choices constant and increasing the knowledge also led to increased times for convergence. This was a big surprise! But it was still good news, because surprises are what make empirical research, especially simulations, worthwhile [11, p. 14]! We conjectured that the inherent symmetry in our problem might be causing this. When we tried to break it by offsetting the agents' choices and knowledge, however, it had no substantial effect on the above behavior. So we discarded that conjecture.

We made another interesting observation. When the local knowledge of the agents is increased, another hidden effect is obtained. This is the amount of *sharing* that the agent has with other agents. Intuitively, as the agents share more and more knowledge, their decisions can become more and more similar, resulting in greater instability. We attempted to characterize the sharing of knowledge among the agents. When the sharing was factored in, we found that it appears to explain the decreased effectiveness of coordination when the extent of choices are held constant.

1.3 Organization

Section 2 describes our experimental setup. Section 3 describes the main experimental results we obtained. Section 4 discusses some relevant conceptual issues, mentions some related literature, and concludes with a description of some open problems.

2 Experimental Setup

The setup consists of an array, each of whose elements is thought of as a resource. Figure 1 shows the array—accessed as a ring—that captures the resources available in the experiment. A number of agents are given. The agents use a given resource by being in the array index corresponding to that resource. There can be multiple agents using a resource; each agent uses exactly one resource. It is tacitly assumed that the quality of a resource received by an agent varies inversely in some way with the number of agents using that resource. Thus—although the utility accruing to an agent is not explicitly modeled in the present version of the experiments—each agent would like to be using a resource that is used by as few agents as possible.

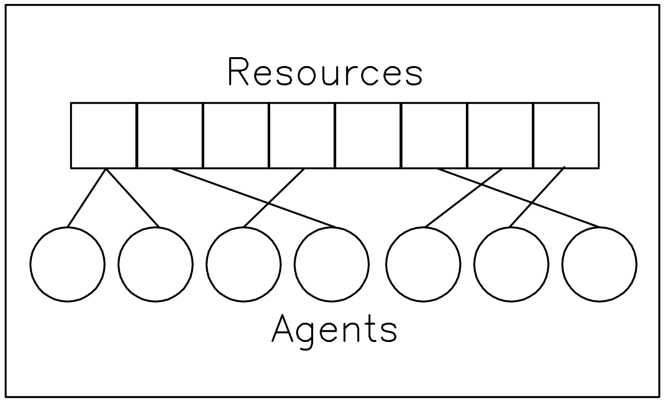

Fig. 1. Resources and Agents

2.1 Knowledge and Choice

As motivated above, it is crucial to distinguish between *knowledge* and *choice*. Knowledge refers to a reduction in uncertainty perceived by the agent. The amount of knowledge available to an agent is given by the number of resources whose occupancy is known to the agent. Thus, the knowledge of an agent increases as the agent is given information about an increasing number of resources.

Choice has to do with the number of actions that an agent is allowed to choose among. In other words, by choice, we mean raw physical choice. Note that a rational agent may find it has fewer realistic choices when it comes to know more facts, but that aspect is not directly measured here.

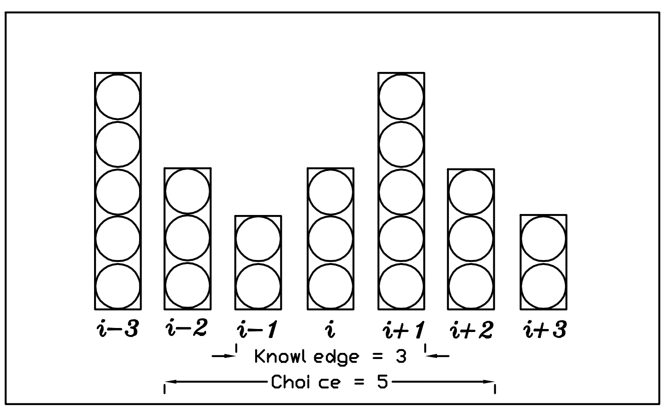

Fig. 2. Knowledge and Choice Windows

Intuitively, knowledge and choice are orthogonal properties. Figure 2 illustrates the knowledge and choice windows for an agent at location i. In the initial experiment, the knowledge and choice windows of an agent were symmetrically distributed around its current location. In later experiments, we allowed for the knowledge and choice windows to be skewed with respect to each other and the agent's current location. This had no significant bearing on the trends observed. For this reason, we report results from the simplest case, where the knowledge and choice windows are placed symmetrically around the agent, just as shown in Figure 2.

2.2 The Protocol

It is postulated that each agent has knowledge of a limited number of the available resources. This knowledge is in terms of the occupancy at a given resource. Using this knowledge, each agent fires a simple rule (the same for all agents) to stochastically decide whether to move to a new (less occupied) location, and if so, which one. In this scheme, the agents gradually disperse from the more crowded locations toward the less crowded ones. The system as a whole stabilizes when all of the resources are equally occupied. This convergent situation represents coordination, because it corresponds to the agents having achieved a sharing of resources that maximizes the performance or utility for each of them. Typically, to facilitate convergence, we set an integral ratio of agents to resources. However, when the convergence condition is liberalized, so that the systems stops even when an exact match is not obtained, the integral ratio requirement can also be safely relaxed.

The expressions used by an agent to compute the probability of moving from current resource i to another resource j in its choice window are given as follows. The f_{ij} are treated as weights.

$$f_{ij} = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{if } i \neq j \text{ and } r_i \leq r_j \\ 1 - \frac{1}{1+\tau \exp(\frac{r_i - r_j - \alpha}{\beta})} & \text{otherwise} \end{cases}$$

where α , β and τ are control parameters, r_i the number of agents at resource i, and r_j the number of agents at resource j. In our experiments, we set $\alpha = 5$, $\beta = 2$, and $\tau = 1$.

The weights are normalized so they are guaranteed to add to 1, and are then treated as probabilities. Thus, the probability of moving from resource i to resource j is then given by

$$p_{ij} = \frac{f_{ij}}{\sum_j f_{ij}}$$

The variables r_i and r_j, which refer to the number of agents at a resource are based on what the given agent knows about the environment. If location j is within the agent's knowledge window (thus it is in the intersection of its knowledge and choice windows), then r_j is the actual value of resource occupancy. If j is not in the agent's knowledge window, then we use an estimated value for it based on the total number of agents and the occupancy of the known part of the world. If location j is not in the choice window, then r_j is not used.

$$r_j = \begin{cases} \text{occupancy of } j & \text{if } j \text{ is in knowledge window} \\ (N-K)/u & \text{otherwise} \end{cases}$$

where N is the total number of agents in the system, K is the number of agents in the knowledge window, u is the number of locations that are not known about. Thus, N and u are a form of global knowledge in the system. Since eliminating them would complicate the present experiment considerably, that aspect is deferred to future work.

2.3 Inertia

Inertia refers to the tendency of an agent to stay in its location even if preferable alternatives are available. This is reflected in the probability p_{ii}. It turns out that the above protocol used by the agents in deciding their actions maximizes the agents' inertia for problems of small dimensions. With small dimensions, especially when the choices are limited, the agent typically has only a few good alternatives. Each good alternative gets a small positive weight; each undesirable alternative gets a weight of 0. Thus, the value of p_{ii} comes out fairly high. As the distribution of the agents becomes more uniform, the inertia of each of them goes up, resulting in an inertia of 1 at equilibrium. An inertia of 1 for all agents denotes convergence, because then none of them move.

From our experiments, we can safely state that inertia is crucial to coordination. Without substantial inertia, the system can become highly unstable leading to situations in which convergence never occurs. We revisit inertia below, but suffice it to state here that we used high inertia just for our simulations to terminate.

3 Analysis

The hypotheses we wished to test were based on the idea that in the original formulation the knowledge and choice are tied together as a single variable, whereas they could be orthogonal in principle. Our hypotheses were

- H1. Increasing the choice and the knowledge simultaneously would increase the time taken to coordinate—in our opinion, this is essentially Sen *et al.*'s main result.
- H2. Increasing the choice while holding the knowledge constant would increase the time taken to coordinate.
- H3. Increasing the knowledge while holding the choice constant would not increase, and may even decrease, the time taken to coordinate.

3.1 Error Tolerance

Instead of defining convergence as precise convergence, we found it convenient to allow a small band of tolerance of error. Thus, a state would be deemed acceptable (and the simulation would halt) if the resources assigned to each agent were within a certain range of the optimal. By reducing the time taken to converge, this enabled us to test configurations involving a larger number of agents and resources than otherwise possible.

We discovered that including some tolerance for error made the system more robust in that the trends were more reliable than otherwise. Intuitively, this is because it reduces the chance that the system would be stuck in a suboptimal state that was several moves away from the optimal, e.g., if almost all of the resources were being used optimally, but one of the resources was under-used and another, faraway resource was over-used.

3.2 Results

		Choice						
		3	5	7	9	11	13	15
Knowledge	3	7	9	20	45	59	100	180
	5		9	20	50	102	117	246
	7			37	65	144	237	335
	9				151	205	519	831
	11					790	901	1563
	13						2974	5023
	15							7520

Table 1. Summary of Number of Iterations to Convergence $\langle 15, 45, \pm 1, 50 \rangle$

		Choice			
		3	5	7	9
Knowledge	3	40	136	749	2336
	5		194	837	4622
	7			977	4759
	9				6766

Table 2. Summary of Number of Iterations to Convergence $\langle 9, 27, \pm 0, 50 \rangle$

Some of our experimental results are displayed in Tables 1 and 2. The tuple in each caption indicates, respectively, the number of agents, the number of resources, the error tolerance, and the number of simulation runs over which the results are averaged.

We always average the results over several runs, but it takes more runs for the results to be reliably duplicated if the tolerance is set low. However, the interesting aspect of the trends is not the exact number of steps taken to converge, but the qualitative relationships among them, such as whether the number of steps is increasing or decreasing and if so at what polynomial order. For this reason, Table 1, which has more data points and a larger error tolerance, is taken as the more important one. Table 2, in which the error is limited to 0, should be treated mostly as a corroboration of Table 1.

We compute the tables only for the upper triangular submatrix, because the lower triangular submatrix is readily determined from it. The lower triangular submatrix corresponds to the knowledge window being a superset of the choice window. In our reasoning protocol, this extra knowledge is useless and harmless, because it does not affect the agent's decisions. Thus, the values are essentially constant along each column below the principal diagonal. (In an actual simulation, they would not be exactly constant because of randomization, but they are reliably approximately equal.)

3.3 Initial Hypotheses

Tables 1 and 2 show that we had mixed success in establishing our initial hypotheses. Hypothesis H1 corresponds to the principal diagonals of Tables 1 and 2. This hypothesis is clearly supported by the evidence. In this respect, by restricting our system, we were able to reconstruct the numerical trends exhibited in [10]. However, because of the above case corresponds to increasing knowledge and choice simultaneously, we do *not* support the conclusion that increasing knowledge *alone* causes a loss of the effectiveness of coordination.

Hypothesis H2 corresponds to rows of Tables 1 and 2. Reading to the right, the time to convergence increases as the choices increases, if the knowledge is held constant. Thus this hypothesis is supported.

Hypothesis H3 corresponds to the columns of Tables 1 and 2. Reading downwards, the time to convergence increases as knowledge increases, even as choices are held constant. Thus this hypothesis is not supported! We explain why next.

3.4 Sharing of Knowledge

We define a metric to estimate the extent of sharing of knowledge among the agents. This metric estimates the "amount" of knowledge of a given agent that is also available to others. This metric obviously depends on the size of the knowledge window. As the windows for the agents increase, the windows overlap to a greater degree with more agents, resulting in higher effective sharing.

To define our metric, let the window size available to all agents be k. The given agent's window overlaps to the extent of $(k-1)$ with agents one slot to the right or left of it, $(k-2)$ with those two slots away, and so on. Thus each agent has a sharing of $\Theta(k^2)$. The sharing in the entire system is $\Theta(Nk^2)$, for a total of N agents. When k is large, we can treat this as $\Theta(k^3)$. In fact, the interesting results are the rightmost column (where choice equals N, and k increases toward N). Now we have the following hypothesis.

H4. Increasing the knowledge while holding the choice constant increases the time proportional to the sharing metric defined above.

Figure 3 is based on the last column of Table 1. (We do not pursue Table 2 further, because it has too few data points. Suffice it to state here that the results are essential alike.) Figure 3 shows that the time to convergence has the same order as the sharing metric. To reduce clutter, we only show the graphs for a cubic polynomial that was

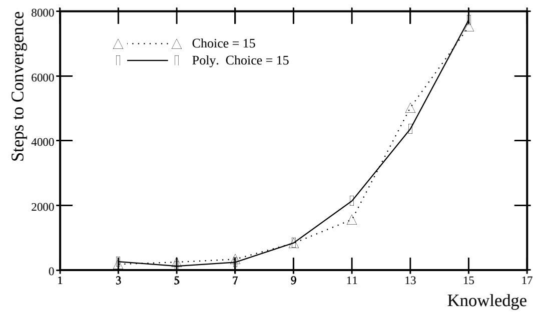

Fig. 3. Relating Sharing of Knowledge and Time to Coordinate $\langle 15, 45, \pm 1, 50\rangle$

fit to the data, and data corresponding to the last column (constant, maximal choice) of Table 1. This does not prove that the sharing of knowledge is the real reason for the delay in convergence. It does, however, give an indication that sharing may have a significant role to play in the final understanding of coordination in decentralized systems where the agents are homogeneous and coordination calls for complementary decisions, as here.

3.5 Inertia Revisited

Recall that inertia refers to the tendency (or probability) of an agent to stay in its present location even in the face of available alternative locations. From the probability calculations of section 2, it should be clear that, in general, as the number of choices increase, $\sum_j f_{ij}$ increases, and consequently the inertia (i.e., p_{ii}) decreases. This reason, especially when coupled with a band of 0 tolerance, can prevent convergence for moderately large dimensions.

We considered an alternative formulation of inertia, in which inertia is given directly in terms of a constant probability. An agent decides among its choices to move by normalizing the probabilities as before. The probability for moving to an undesirable alternative is still set to 0; however, the sum of the probabilities of moving to good alternatives are limited to $(1 - p_{ii})$. We discovered in preliminary experiments that going from a high inertia (0.9) to a medium inertia (0.7 or 0.5) can cause significant variations in the trends observed. Those results are not yet suitable for reporting.

4 Discussion

This paper developed some experimental results about coordination in a simple setting involving multiple, potentially conflicting, autonomous agents. Despite its simplicity, it led to nontrivial and surprising results. By using an experimental framework more general than that of Sen *et al.*, we were able to reproduce their numeric results as a special case, yet also show how their conclusions were not supported.

There are some limitations of the present experimental setup. It focuses on cases where the resource conflicts are direct and immediately perceived, the resources are homogeneous, the agents all use the same decision-making protocol, and the agents do not communicate directly. Further, there are well-known limitations of reinforcement learning in terms of time taken to learn even simple concepts. The present experiments leave open the possibility that more sophisticated agents in more flexible environments, where their learning is supervised in certain ways might discover better ways of coordination, which may turn out to have different characteristics in terms of the influence of knowledge and choice.

Our contribution, however, is not only in developing the results we presented, but in identifying some of the several factors that play a role in determining the coordination of autonomous agents. We also made some progress in delineating the trade-offs among these factors. In general, in making claims about an intuitively interesting concept, we must avoid the risk that other factors may intrude into our representation, processing, or measurement and collation. This is a difficult task where theoretical development must be interleaved with controlled experimentation or simulation. We have only taken the initial steps of such a systematic study.

Although the present results should not be taken as final, it is essential to report and discuss them. This is because of two major reasons. One, the problem of learning to coordinate and its relationship to other concepts is crucial to theories and architectures of agents and multiagent systems. Two, the present kinds of studies are of the category of *exploratory research*, which Cohen [2] eloquently argues is key to empirical research and must occur prior to the formulation of more precise questions and experimental protocols that are ultimately the core of experimental science.

4.1 Literature

In addition to the works mentioned above, some interesting relevant approaches are known in the literature. For instance, Kuwabara *et al.* present a market-based approach in which agents controlling different resources set their prices based on previous usage, and buyer agents choose which resources to use [7]. The buyer agent can use more than one resource concurrently, and seeks to minimize its total price it has to pay. As in our approach, the buyer's decision-making is probabilistic. Although their model is similar to ours, they do not study the reasons for achieving effective coordination.

Results by Hogg & Huberman indicate the potential benefits of introducing heterogeneity of different forms [4]. These agree with the intuition that in homogeneous settings, the sharing of knowledge may have an undesirable effect on coordination. This is especially so when the agents must make complementary decisions so as to coordinate, i.e., move to different locations. This problem is closely related to the emergence of conventions for resource sharing [8].

4.2 Future Work

Although we introduced some interesting considerations, a lot remains to be done.

Choice bears an interesting relationship to the notion of commitments. It appears that the two are complementary in that the greater the agent's choice the lower its commitment to a particular decision. Previous experimental work by Kinny & Georgeff [6] and Pollack *et al.* [9] appears especially relevant. However, there is more structure to commitments that the present setup does not capture; some of this is discussed in [12].

We mention some high-level open issues that would extend the experiments described above. Although not as detailed as hypotheses, they can be studied in variations of the present experiments.

- I1. The improvement in speed with a nonzero band suggests a natural trade-off between the time taken and the quality of the solution. We conjecture that the time required increases exponentially as the tolerance is reduced to zero.
- I2. In settings where the agents coordinate by making the similar, but noncomplementary, decisions, increasing the sharing of knowledge will improve coordination.
- I3. A large class of strategies leading to adaptive behavior may be approximated by varying the inertia of the agents dynamically.

Our programs and data are available (for educational and research purposes) from `http://www4.ncsu.edu/eos/info/dblab/agents/skrustog/data/`.

References

1. Scott Clearwater, editor. *Market-Based Control: A Paradigm for Distributed Resource Allocation*. World Scientific, 1996.
2. Paul R. Cohen. *Empirical Methods for Artificial Intelligence*. MIT Press, Cambridge, MA, 1995.
3. Edmund H. Durfee. *Coordination of Distributed Problem Solvers*. Kluwer, 1988.
4. Tad Hogg and Bernardo A. Huberman. Controlling chaos in distributed systems. *IEEE Transactions on Systems, Man, and Cybernetics*, 21(6):1325–1332, 1991.
5. F. Felix Ingrand, Michael P. Georgeff, and Anand S. Rao. An architecture for real-time reasoning and system control. *IEEE Expert*, 7(6), 1992.
6. David N. Kinny and Michael P. Georgeff. Commitment and effectiveness of situated agents. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 82–88, 1991.
7. Kazuhiro Kuwabara, Toru Ishida, Yoshiyasu Nishibe, and Tatsuya Suda. An equilibratory market-based approach for distributed resource allocation and its applications to communication network control. In *[1]*. 1996.
8. David K. Lewis. *Convention: A Philosophical Study*. Harvard University Press, Cambridge, MA, 1969.
9. Martha E. Pollack, David Joslin, Arthur Nunes, Sigalit Ur, and Eithan Ephrati. Experimental investigation of an agent commitment strategy. TR 94-13, Computer Science, University of Pittsburgh, Pittsburgh, June 1994.
10. Sandip Sen, Shounak Roychoudhury, and Neeraj Arora. Effect of local information on group behavior. In *Proceedings of the International Conference on Multiagent Systems*, pages 315–321, 1996.
11. Herbert Simon. *The Sciences of the Artificial*. MIT Press, Cambridge, MA, third edition, 1996.

12. Munindar P. Singh. Commitments in the architecture of a limited, rational agent. In *Proceedings of the Workshop on Theoretical and Practical Foundations of Intelligent Agents*, pages 72–87. Springer-Verlag, 1997.
13. Munindar P. Singh, Anand S. Rao, and Michael P. Georgeff. Formal methods in DAI: Logic-based representation and reasoning. In *[14]*, chapter 8. 1998. www.csc.ncsu.edu/ faculty/ mpsingh/ papers/ mas/ formal-DAI.ps.
14. Gerhard Weiß, editor. *Introduction to Distributed Artificial Intelligence*. MIT Press, Cambridge, MA, 1998.