

Be Patient and Tolerate Imprecision: How Autonomous Agents can Coordinate Effectively*

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

A decentralized multiagent system comprises agents who act autonomously based on local knowledge. Achieving coordination in such a system is nontrivial, but is essential in most applications, where disjointed or incoherent behavior would be undesirable. Coordination in decentralized systems is a richer phenomenon than previously believed. In particular, five major attributes are crucial: the extent of the local knowledge and choices of the member agents, the extent of their shared knowledge, the level of their inertia, and the level of precision of the required coordination. Interestingly, precision and inertia turn out to control the coordination process. They define different regions within each of which the other attributes relate nicely with coordination, but among which their relationships are altered or even reversed. Based on our study, we propose simple design rules to obtain coordinated behavior in decentralized multiagent systems.

1 Introduction

Coordination is key to the design of multiagent systems. Often, the multiagent systems must be *decentralized*—whose member agents act autonomously based on local information. Such systems are essential in a number of applications where the agents may not wish to or be able to communicate or have a common plan.

Coordination has been studied before. In the context of distributed problem-solving and generalized partial planning, many good results have been obtained [Decker and Lesser, 1995; Durfee, 1999]. However, the key features of decentralized systems and their relationship to coordination have not yet been fully explored. Their study is the theme of this paper.

Let's begin with a brief historical overview. The early work on coordination considered knowledge as a key factor. Although decentralized systems of the kind we study

were not always considered, the community's folklore is that more knowledge leads to better coordination. It is also recognized that the locally best actions would not always lead to the best payoff for an individual agent much less for the system as a whole.

Schaerf *et al.* consider multiagent reinforcement learning in the context of load balancing in distributed systems [1995]. In their framework, the agents share a number of resources, which they autonomously select to use. When all agents are noncooperative, e.g., by always selecting their most preferred resources, they all stand to lose. However, when individuals sometimes select the less desirable resources, the entire population benefits. In this system, communication may not be useful in improving the performance of the population and may in fact be detrimental.

In a simpler framework, Sen *et al.* also study coordination among agents sharing resources [1996]. Coordination corresponds to achieving equilibrium. Sen *et al.* argue that, contrary to conventional wisdom, giving the interacting agents additional knowledge causes the coordination to slow down. Baray uses the same framework, but applies genetic algorithms to show how coordination can be speeded up [1998].

Rustogi & Singh study coordination in a similar framework [1999]. They show that in addition to knowledge, the choices available and the extent of the knowledge shared by the agents are also important. Rustogi & Singh show that coordination slows down when the available choices increase. When shared knowledge increases, then too coordination slows down. There is no direct contradiction with Sen *et al.*, because their results correspond to the case where the agents' knowledge also increases.

The present paper advances the above program of research by bringing in additional features of decentralized systems in order to better characterize the outcome of coordination. Our experiments indicate that perfect coordination is often inordinately more time-consuming than slightly imperfect coordination. Usually, if the agents exhibit higher patience or inertia in terms of not jumping to another resource, they can coordinate faster.

*Supported by the NCSU College of Engineering, the National Science Foundation under grant IIS-9624425 (Career Award), and IBM corporation.

Organization Section 2 describes our experimental setup. Section 3 describes the main experimental results we obtained. Section 5 discusses some relevant conceptual issues, mentions some related literature, and concludes with a description of some open problems.

2 Experimental Setup

Our setup (Figure 1) consists of an array of equivalent resources. Each agent uses exactly one resource, but a resource can support several agents. The agents prefer resources that support fewer other agents. The agents know the occupancy of a certain number (Kn) of other resources besides their own. They can elect to move to any of a certain number (Ch) of resources. The agents move only to resources that appear better. They gradually disperse from the more crowded resources toward the less crowded ones.

Equilibrium is achieved when the agents are uniformly distributed over all resources, and none move. Equilibrium corresponds to perfect coordination, because it means the agents have achieved a locally and globally optimal sharing of resources. Note that the present setting requires the same or complementary decisions. In general, complementary decisions are more interesting, because they cannot be hardwired in some trivial mechanism.




Figure 1: Knowledge (Kn), choice (Ch), and knowledge sharing of agents at resources i and j

2.1 Decision Protocol

Each agent stochastically decides whether to move and where. All agents use the same decision function and only move to better resources. 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, and r_i and r_j the number of agents at resources i and j , respectively. In our experiments, we set $\alpha = 5$, $\beta = 2$, and $\tau = 1$ unless otherwise specified.

The weights are normalized to yield probabilities. Thus, the probability of an agent moving from resource

i to resource j is given by

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

Intuitively, when the choices are limited (as when small problems are considered), 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 levels out, their p_{ii} values increase until each of them becomes 1 meaning that none of the agents can move.

2.2 Key Concepts

This simple framework provides enough structure to capture a variety of interesting concepts.

Choice. The number of actions an agent may choose from. 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. If resource j is not in the choice window, then r_j is not used, and $p_{ij} = 0$.

Knowledge. 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.

The variables r_i and r_j give the occupancy of resources i and j . They are accurate for resources within the agent's knowledge window. For other resources, they are estimated based on the total number of agents and the occupancy of the known part of the world.

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

where N is the total number of agents, K is the number of agents in the knowledge window, and u is the number of resources 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.

Inertia. This is the tendency of an agent to stay in its resource even if better alternatives are known. Greater inertia means that the probability p_{ii} is higher. In our setup, inertia is controlled by α . As remarked above, if all p_{ii} values increase to 1, coordination is achieved. Thus inertia can facilitate coordination. A system whose agents have low inertia may exhibit chaotic behavior, and never achieve coordination. On the other extreme, very high inertia would lead to an inactive system, with a similar result.

Sharing. Shared knowledge corresponds to overlapping knowledge windows. Rustogi & Singh estimate the total amount of sharing in the system as roughly proportional to the cube of the size of the knowledge window.

Under a homogeneous strategy (as here), shared knowledge would tend to lead to similar decisions, which could influence coordination.

Precision. Imprecision is the distance from a perfectly coordinated state, i.e., the minimum number of agent relocations required to coordinate. Given an acceptable level of imprecision, we control the simulations to halt when that level is reached. Introducing precision into the experimental framework had important consequences. First, because coordination is achieved much faster when imprecision is allowed, we could simulate much larger configurations than otherwise possible. Second, allowing some imprecision made the trends more robust by reducing the likelihood of pathological situations in which the system may get stuck. Third, imprecision helps us study the above pathological situations, which are interesting in their own right. This is the basis for some technical results presented later.

3 Results

The following figures indicate our results. The tuple in each caption indicates, respectively, the number of resources, the number of agents, the initial deviation (distance of the agent distribution from a coordinated state), and the imprecision tolerated.

| Knowledge | Choice | | | | | | | |
|-----------|--------|----|----|----|-----|------|------|-------|
| | 2 | 4 | 6 | 8 | 10 | 12 | 14 | 16 |
| 2 | 64 | 21 | 22 | 31 | 44 | 110 | 127 | 74 |
| 4 | | 17 | 20 | 20 | 59 | 115 | 139 | 174 |
| 6 | | | 17 | 31 | 103 | 180 | 254 | 366 |
| 8 | | | | 51 | 142 | 194 | 518 | 931 |
| 10 | | | | | 339 | 710 | 1007 | 1464 |
| 12 | | | | | | 1096 | 2464 | 7270 |
| 14 | | | | | | | 5257 | 10550 |
| 16 | | | | | | | | 21069 |

Table 1: Number of steps to coordination (16, 48, ± 24 , ± 4)

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 decision 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 simulations, the randomization can cause minor variations.)

3.1 Sharing of Knowledge

Figure 2 is based on the last column of Table 1. Figure 2 shows that the time to achieve coordination has the same order as the sharing metric. To reduce clutter, we only show the graphs for a cubic polynomial that was fit to the data, and data corresponding to the last




Figure 2: Effect of sharing of knowledge (16, 48, ± 24 , ± 4)

column (constant, maximal choice) of Table 1. This figure indicates 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.

3.2 Precision

Reducing the required precision enhances the scalability of coordination. In other words, as the quality of the coordination increases, the cost in terms of time becomes extremely high. We studied this observation further by delineating the effect of the deviation from coordination at the start of each simulation run. In our setup, the deviation ranges from 1 (almost coordinated) to 15 (maximally uncoordinated).




Figure 3: Effect of initial deviation (10, 30, ± 0)

Figure 3 demonstrates that it takes far fewer steps to progress from maximal uncoordination to almost perfect coordination than to go from almost perfect coordination to perfect coordination. The last little bit of precision consumes almost all of the effort.

The previous result suggests that the time to coordinate increases exponentially as the allowed imprecision is reduced to zero. Figure 4 supports this claim. The exponential variation occurs as inertia drops significantly, resulting in increasing instability. The exponential variation described above, however, does not manifest itself when the agents have little choice, because in such scenarios, the agents cannot move around much anyway. Instead, as Figures 5 and 6 demonstrate, for small choice




Figure 4: Effect of imprecision $\langle 10, 30, \pm 15 \rangle$

windows, the time for coordination increases only polynomially with reducing allowed imprecision. In these figures, to help visualize the trends better, each curve is normalized to 1 with respect to its maximum value. It should be obvious, however, that for low values of imprecision (including 0), the actual time to coordinate increases with choice. For higher values, the time to coordinate is practically independent of the choices available to the agents or the knowledge possessed by them.




Figure 5: Effect of choice relative to imprecision $\langle 10, 30, \pm 15 \rangle$ (each curve is normalized to 1)




Figure 6: Effect of choice relative to imprecision $\langle 16, 48, \pm 24 \rangle$ (each curve is normalized to 1)

3.3 Inertia

Recall that inertia refers to the tendency of an agent to stay in its present resource even if it knows of better resources. 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 an imprecision of 0, can prevent coordination for moderately large dimensions.




Figure 7: Effect of inertia (α) $\langle 10, 30, \pm 15, \pm 0 \rangle$

In our setup, inertia is characterized by the parameter α . The preceding results were based on $\alpha = 5$; now we vary α above and below this value. Figure 7 shows that increasing the inertia facilitates coordination. This is because when the agents are less likely to move, a low occupancy resource will not suddenly be occupied by several agents. Conversely, decreasing the inertia to a low value can make coordination extremely slow. The agents appear to jump about too much and system takes longer and longer to converge. For such cases, the detrimental effect of shared knowledge still applies; thus adding knowledge slows coordination.

Interestingly, for high inertia, an increase in knowledge or choice further improves the coordination. This relationship is a reversal from when the inertia is low. It appears that the trend changes, because higher inertia limits agent movement to such an extent that the benefits of additional local knowledge in decision-making overshadow the usual ill effects of increased sharing of knowledge.

The improvement of coordination due to increasing inertia is observed only if the inertia is not too high. Increasing the inertia to a very high value results in slow coordination. This is because very high inertia causes the agents to freeze in whatever resources they occupy.

3.4 Other Variants Considered

Our interest is in understanding the phenomenon of coordination in general, not analyzing the specific setup used in our experiments. Thus we emphasize the trends observed in the simulations, and the qualitative relationships among the trends, such as whether the number of steps is increasing or decreasing and if so at what poly-

mial order. Our experiments included complex scenarios, but which also yield the same trends as the simple scenarios on which the above results are directly based.

- Our results hold for several decision functions, but we present only the simple decision function used by Sen *et al.*
- Like Rustogi & Singh, we observed that keeping the knowledge and choice windows of an agent symmetrically distributed around its current resource yield the same trends as when the windows are skewed with respect to each other; therefore, we focus on the simpler situation.
- To enable convergence, we set an integral ratio of agents to resources. This is not strictly necessary when imprecise coordination is allowed, but changing the ratio has no effect on the trends, so we report only the integral situations here.
- Except when precision itself is a variable, we can make do with lower precision, because it yields faster convergence without affecting the qualitative nature of the trends.
- We studied the role of inertia and its interplay with knowledge and choice, by altering the control parameters (α , β , and τ) in our protocol. The results highlighted an interesting interplay among the various bases of coordination. Varying α alone, however, provides representative results.

4 Mapping the Terrain

Our experimental study of decentralized multiagent systems brought out a number of important factors that affect coordination. Some of these factors—*inertia* and *precision*—have not been empirically studied in such systems. Others—*knowledge* and *choice*—have been studied but, as our analysis showed, the trends relating to these are richer than believed. Trends due to inertia and precision can dominate and sometimes reverse the simpler trends.

The following simple rules summarize our qualitative results.

- R1.** Low inertia & low imprecision \implies knowledge sharing governs \implies local knowledge & limited choice performs better
- R2a.** Moderately high inertia \implies extent of knowledge or choice is less important
- R2b.** High imprecision \implies extent of knowledge or choice is less important
- R3.** Very high inertia \implies system inactivity

The above rules demarcate the most important regions of our terrain. Figure 8 illustrates the corresponding regions. Rule *R1* supported by Figure 2, is mapped to Region *I* in Figure 8. To achieve effective coordination in this region, agents must limit their knowledge as well as choice. The results of Sen *et al.* and Rustogi & Singh lie within this region. Figures 5-6 and 7 support the

rules *R2b* and *R2a*, respectively. The results of Baray lie within this region—this is the reason he obtains much faster coordination than Sen *et al.* These rules, mapped to region *II* of Figure 8, imply that knowledge and choice are less relevant for coordination. Rule *R3* is intuitively obvious and is represented by region *III* in Figure 8.




Figure 8: Mapping the terrain of decentralized systems

The study of coordination is interesting from a practical engineering standpoint. The above rules yield heuristics to aid in the engineering of a multiagent system. Our first conclusion is that for maximal scalability, we should allow some imperfection in coordination. Even a slight imperfection improves performance considerably. A moderately high value of inertia is desirable. Selecting the right value is nontrivial, especially because it will change in a dynamic system. An open problem is to devise online learning techniques to adapt to the right inertia during execution.

Interestingly, for most of the situations in our setup, local information performs better than global information. Even when local information gives suboptimal results, for many applications, it can provide a reasonable tradeoff with the cost incurred in acquiring the global information.

5 Discussion

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 [1996]. The buyer agent can use more than one resource concurrently, and seeks to minimize the total price. As in our approach, the buyer’s decision-making is probabilistic. Although Kuwabara *et al.*’s model is similar to ours, they do not study the reasons for achieving effective coordination.

Rachlin *et al.* show how agents, using the A-Team architecture, can achieve coordination without explicit communication [1999]. An A-Team is an asynchronous team of agents that shares a population of solutions that evolve over time into an optimal set of solutions.

Through sharing of the solution population, cooperative behavior between agents may emerge leading to better solutions than any one agent could produce. Often, however, a human agent may be necessary to help achieve coordination by imparting domain-specific knowledge.

Results by Hogg & Huberman indicate the potential benefits of introducing heterogeneity of different forms [1991]. 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 resources. This problem is closely related to the emergence of conventions for resource sharing [Lewis, 1969].

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.

Although we introduced some interesting considerations, a lot remains to be done. Choice and inertia bear 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 appears especially relevant. Kinny & Georgeff empirically investigate how the agents' commitment to their current plan contributes to their effective behavior [1991]. The agents in their work are characterized as bold, normal, or cautious based on the extent of their commitment (akin to inertia here), ranging from high to low, in that order. The cautious agents continually reconsider their plan at every step, in the face of a dynamic environment, and therefore exhibit the least commitment. For the most part, bold agents, despite their higher degree of blind commitment, perform better than normal and cautious agents except when the rate of change is very high. Kinny & Georgeff, however, do not study the effectiveness of behavior when the degree of commitment is very high.

We identified several of the key attributes affecting coordination in a way that agrees with but subsumes previous results. We also give some heuristics to develop decentralized multiagent systems. To cover additional applications, we need to consider communication among agents and to determine the circumstances under which it helps or disrupts coordination. We need to consider systems whose membership involves agents being removed and added back.

References

- [Baray, 1998] Cristobal Baray. Effects of individual decision schemes on group behavior. In *Proceedings of the 3rd International Conference on Multiagent Systems (ICMAS)*, pages 389–390. July 1998.
- [Decker and Lesser, 1995] Keith Decker and Victor Lesser. Designing a family of coordination algorithms. In *Proceedings of the International Conference on Multiagent Systems*, pages 73–80, 1995.
- [Durfee, 1999] Edmund Durfee. Distributed problem solving and planning. In Gerhard Weiß, editor, *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*. MIT Press, Cambridge, MA, chapter 3, pages 121–164. 1999.
- [Hogg and Huberman, 1991] Tad Hogg and Bernardo Huberman. Controlling chaos in distributed systems. *IEEE Transactions on Systems, Man, and Cybernetics*, 21(6):1325–1332, 1991.
- [Kinny and Georgeff, 1991] David Kinny and Michael Georgeff. Commitment and effectiveness of situated agents. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 82–88, 1991.
- [Kuwabara *et al.*, 1996] Kazuhiro Kuwabara, Toru Ishida, Yoshiyasu Nishibe, and Tatsuya Suda. An equilibrium market-based approach for distributed resource allocation and its applications to communication network control. In Scott Clearwater, editor, *Market Based Control: A Paradigm for Distributed Resource Allocation*. World Scientific, 1996.
- [Lewis, 1969] David Lewis. *Convention*. Harvard University Press, Cambridge, MA, 1969.
- [Rachlin *et al.*, 1999] John Rachlin, Richard Goodwin, Sesh Murthy, Rama Akkiraju, Fred Wu, Santhosh Kumaran, and Raja Das. A-Teams: An agent architecture for optimization and decision-support. In *Intelligent Agents V: Proceedings of the 5th International Workshop on Agent Theories, Architectures, and Languages (ATAL-98)*, pages 261–276. Springer-Verlag, 1999.
- [Rustogi and Singh, 1999] Sudhir Rustogi and Munindar Singh. The bases of effective coordination in decentralized multiagent systems. In *Intelligent Agents V: Proceedings of the 5th International Workshop on Agent Theories, Architectures, and Languages (ATAL-98)*, pages 149–162. Springer-Verlag, 1999.
- [Schaerf *et al.*, 1995] Andrea Schaerf, Yoav Shoham, and Moshe Tennenholtz. Adaptive load balancing: A study in multi-agent learning. *Journal of Artificial Intelligence Research*, 2:475–500, 1995.
- [Sen *et al.*, 1996] 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.