Abstract

RUSTOGI, SUDHIR KUMAR. Empirical Studies of Coordination in Decentralized Multiagent Systems (Under the direction of Munindar P. Singh).

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. Inertia together with imprecision, are also crucial for the scalability of coordination. Based on our study, we propose simple design rules to obtain coordinated behavior in decentralized multiagent systems.

EMPIRICAL STUDIES OF COORDINATION IN DECENTRALIZED MULTIAGENT SYSTEMS

BY

SUDHIR KUMAR RUSTOGI

A THESIS SUBMITTED TO THE GRADUATE FACULTY OF

NORTH CAROLINA STATE UNIVERSITY

IN PARTIAL FULFILLMENT OF THE

REQUIREMENTS FOR THE DEGREE OF

MASTER OF SCIENCE

DEPARTMENT OF COMPUTER SCIENCE

RALEIGH

AUGUST 1999

APPROVED BY:

CHAIR OF ADVISORY COMMITTEE

Biography

Sudhir Kumar Rustogi was born on March 11, 1966 in Pataudi, Haryana, India. He obtained his Bachelor of Technology degree in Civil Engineering from Indian Institute of Technology, Delhi, India in 1987. From 1987 to 1988, he worked as a bridge design consultant with Gammon India Ltd., New Delhi.

Before coming to North Carolina State University in 1994, he worked as a design engineer with Nuclear Power Corporation India Ltd. (NPCIL), Bombay. At NPCIL, from 1988 to 1989, he took advanced training which included coursework in civil engineering, computing, and nuclear engineering. He was awarded the prestigious Homi J. Bhabha prize in 1989 for securing the first rank among all trainee engineers.

At North Carolina State University, in 1998, he obtained his doctoral degree in Civil Engineering with a specialization in Structural Engineering. He was elected to the Honor Society of Phi Kappa Phi in 1996.

Acknowledgments

I am grateful to Professor Munindar P. Singh for his constant guidance and support throughout my graduate studies at North Carolina State University. I also wish to thank the other members of my advisory committee, Professors James C. Lester and Robert St.Amant for their many valuable comments.

This research was partially supported by the National Science Foundation under grants IRI-9529179 and IRI-9624425, and by IBM corporation.

I also appreciate the help and support of all my colleagues working at the database laboratory, especially, Mahadevan Venkatraman, Anuj Jain, Jie Xing, and Feng Wan.

I am indebted most of all to my parents for their unconditional support, to my wife Neerja for encouraging me to pursue my research and to my daughter Shreya for her love and understanding.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

Coordination is the act of choosing one's action based on the expectation of others' actions. We need to coordinate our actions whenever we are sharing goals, resources, or problem solving capabilities. Lack of coordination can result in difficulty or even inability in achieving personal as well as social goals. Therefore, coordination is the key to individual as well as group success.

The study of coordination is a multidisciplinary activity [Malone and Crowston, 1994]. In computer science, the coordination of multiple agents is crucial to the design of multiagent systems. A particularly interesting class of coordination problems arises in decentralized multiagent systems. Agents in such systems often act autonomously based on incomplete (local) information that is often affected by uncertainties and time delay. Further, agents in such systems may not wish to or be able to communicate or have a common plan.

Several researchers have studied coordination in decentralized multiagent systems.

Distributed artificial intelligence (DAI) researchers have primarily focussed on cooperative distributed problem solving [Durfee, 1988, Gasser and Jr., 1990, Decker and Lesser, 1995, Durfee, 1999] by sophisticated agents that work together to solve problems that are beyond their individual capabilities. The Partial Global Planning (PGP) approach [Durfee, 1988] requires agents to coordinate by scheduling the timely generation of partial results, avoiding redundant activities, shifting tasks to idle nodes, and indicating compatibility between goals. Recognition of coordination relationships among tasks in PGP, however, is dependent on details of a particular task environment. Generalized Partial Global Planning (GPGP) [Decker and Lesser, 1995] generalizes the coordination relationships used in PGP, thus forming an extendable family of coordination mechanisms, any subset or all of which can be used in response to a particular task environment. The TAEMS framework represents the coordination problem in a way that abstracts out the details of a particular task environment [Decker and Lesser, 1995].

The above cooperative distributed problem solving approaches incorporate cooperation as an integral part of the design of the system. The agents in the system are all working towards a single goal. In a slightly different approach, Liu and Sycara [1994] present a collective problem solving framework, where problem solving is viewed as an emergent functionality from the evolving process of a society of diverse, interacting, well-coordinated reactive agents.

Despite many useful results, the above studies provide little domain-independent agreement on the phenomena that affect coordination in decentralized multiagent systems. They do not particularly study the key features of decentralized systems and their relationship to coordination. Further, above approaches often rely on communication for achieving coordination. Communication, though an invaluable tool, does not guarantee coordinated behavior [Halpern and Moses, 1990], is time-consuming, costly and can sometimes be undesirable.

A direction of research explores the key features of decentralized systems and their relationship to coordination. Advancing this program of research is the primary focus of this thesis. But first a brief historical review of this research direction is warranted.

The early work on coordination considered knowledge as a key factor [Durfee, 1988]. Knowledge is also a key component of several abstract agent architectures, e.g., the family of belief-desire-intention (BDI) architectures [Ingrand et al., 1992, Singh et al., 1999]. Although decentralized systems of the kind studied here were not always considered, the folklore in the research community 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 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. This is analogous to the well-known prisoner's dilemma [Axelrod, 1985]. In Schaerf *et al.*'s system, agents rely on limited information to achieve coordination without explicit communication. Here 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 [1994, 1996, 1998]. The agents decide locally, and coordination corresponds to their 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 as Sen, but applies genetic algorithms and shows that coordination can, in fact, be speeded up by giving additional knowledge [1998]. The study, however, does not examine the reasons behind such a reversal of trends from Sen's results.

Huberman and Hogg [1988] study the behavior of computational ecologies which explicitly incorporates the features of incomplete knowledge and delayed information. Their theory gives rise to a panoply of interesting behaviors that include asymptotic regimes characterized by fixed points, oscillations, and chaos. This study, although rigorous, does not consider many other key attributes of decentralized systems, which are described later in this thesis. Kephart *et al.* further study the asymptotic dynamics predicted by Huberman and Hogg and show a mechanism to damp out the oscillatory and even chaotic behavior [1989].

This thesis advances the above program of research by bringing in additional features of decentralized systems in order to better characterize the outcome of coordination [Rustogi and Singh, 1999a,b]. In particular, the following questions that emerge at the interface of agent theory and architecture are addressed.

- 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 is a new area of investigation, we follow Simon's advice to study carefully designed simulations to develop a clearer understanding of the theoretical concepts [Simon, 1996, p. 15]. The experimental results indicate that in particular, five major attributes of decentralized systems are crucial:

- the extent of local knowledge and choices of member agents.
- the extent of their shared knowledge.
- the level of their inertia or patience in terms of not jumping to another resource.
- and the level of precision of the required coordination.

Our experiments show that for a low value of inertia, coordination slows down when the choices available to agents increase [Rustogi and Singh, 1999a]. When shared knowledge increases, then too coordination slows down. It is shown that perfect coordination is often inordinately more time-consuming than slightly imperfect coordination [Rustogi and Singh, 1999b]. However, if the agents exhibit higher patience or inertia, they can usually coordinate faster. Inertia is also shown to be crucial for the scalability of coordination. Further, our experiments verify the undesirable effects on coordination of uncertainty and delay in updating knowledge.

We study coordination, empirically, in the context of a decentralized resource allocation problem. Coordination is required as agents share the resources in the system. The concepts in our work, however, apply to other domains as well. For example, we can use these concepts to help coordinate agents that are playing a soccer game. Coordination is required as agents in a team share the goal of putting the ball in opposing teams' net and of preventing the opposing team from doing the same. Often, an agent carrying the ball gets open, and one of the opposing agents must reach out to block it. However, if the agents in the opposing team have low inertia, more than one agent may reach out simultaneously, which leaves another agent open and sets up a pass.

Chapter 2 describes the simplified setup used for empirical study of coordination in decentralized systems, including the probabilistic decision protocol used by the agents and the key attributes studied using this setup. The relationship of these attributes to the coordination process is studied in Chapter 3. This chapter also studies the scalability of the various experimental results.

Chapter 4 maps the rich terrain of coordination in decentralized systems using simple qualitative rules that yield heuristics for designing multiagent systems. Chapter 5 discusses some relevant conceptual issues, mentions the pertitent literature and concludes with a description of the future research direction.

Chapter 2

Empirical Study of Coordination

The experimental framework used in this study generalizes over the one used by Sen *et al* [1994, 1996, 1998]. The setup consists of an array, each of whose elements is thought of as a resource. Figure 2.1 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 assumed that the quality of a resource received by an agent varies inversely with the number of agents using that resource. Figure 2.2 shows the utility accruing to an agent based on such a model. Thus each agent would like to be using a resource that is used by as few agents as possible. Further, it is assumed that all resources are equivalent and mutually interchangeable.

Agents in the system gradually disperse from the more crowded resources towards the less crowded ones. Equilibrium is achieved when the agents are uniformly

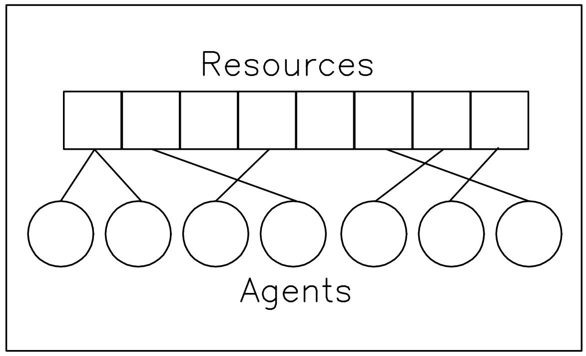

Figure 2.1: Resources and agents

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 complementary rather than same decisions on part of the agents. In general, complementary decisions are more interesting, because they cannot be hardwired in some trivial mechanism.

Whereas Sen *et al.* considered only knowledge (which is coupled in their setup with choice as we show later), we considered in our setup the other key concepts that were briefly mentioned in Chapter 1. These are further elaborated upon later in this chapter. The knowledge available to an agent is measured in terms of the number of resources whose occupancy is known to the agent. An agents' choice, on the other hand, corresponds to the number of resources it *can elect* to move to. Thus, knowledge and choice are orthogonal properties. Figure 2.3 illustrates the knowledge and choice windows for an agent currently located at resource i. For the specific case shown, the agent knows about fewer resources than it can choose to move to. Thus, some of its potential actions are marked by uncertainty.

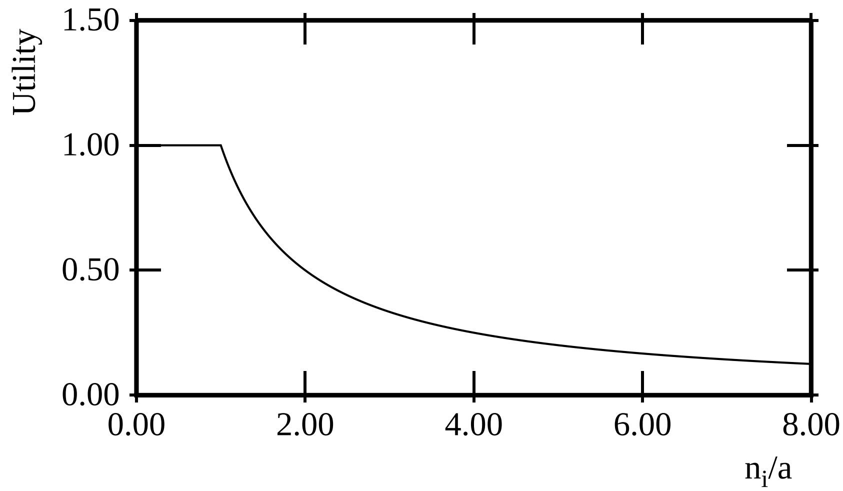

Figure 2.2: Utility accruing to an agent at resource i (n_i is the number of agents at resource i and a is the average number of agents per resource in the system)

In initial experiments, the knowledge and choice windows of an agent were symmetrically distributed around its current location, as shown in Figure 2.3. In later versions, we allowed for the knowledge and choice windows to be skewed with respect to each other and the agents' current location. This, however, proved to have no significant bearing on the trends observed.

2.1 Decision Protocol

We postulate 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 to stochastically decide whether to move to a new location and, if so, which one. All agents use the same probabilistic decision

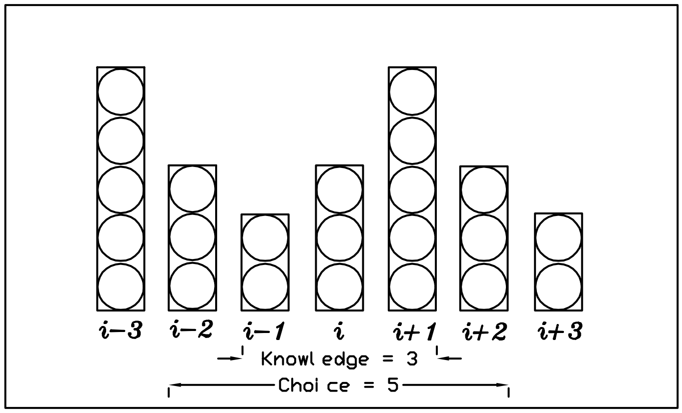

Figure 2.3: Knowledge and choice windows

function and only move to less occupied resources than their present resource. 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} values are treated as weights.

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

where α , β and γ 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 $\gamma = 1$.

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 given by

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

2.2 Key Attributes

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

2.2.1 Knowledge

Knowledge refers to a reduction in uncertainty perceived by the agent. Traditionally, knowledge is believed to help coordination. The amount of knowledge available to an agent in our setup, 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. The variables r_i and r_j, which refer to the number of agents at resources i and j, respectively, are based on what an agent located at resource i knows about the environment. If location j is within the agents' knowledge window, then r_j is the actual value of resource occupancy.

2.2.2 Choice

Choice has to do with the number of actions among which an agent is allowed to choose. In other words, by choice, we mean raw physical choice. Intuitively, too many choices would more likely disrupt rather than help coordination. 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. If resource j is not in the choice window, then r_j is not used, and $p_{ij} = 0$, even if the resource j is within its knowledge window.

If, however, resource j is in the choice window of an agent, but not in its knowledge window, r_j is 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.

2.2.3 Sharing of Knowledge

The amount of shared knowledge among the agents corresponds to overlapping knowledge windows, as shown in Figure 2.4. 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.

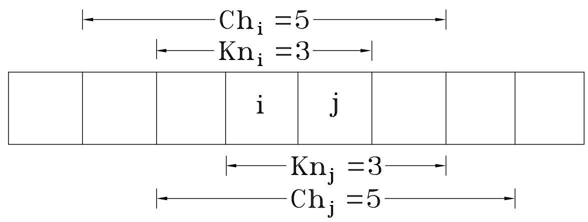

Figure 2.4: Knowledge (Kn), choice (Ch), and knowledge sharing of agents at resources i and j

2.2.4 Inertia

Inertia refers to the reluctance that 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 have low inertia may never achieve coordination. In current setup, inertia is captured using probability p_{ii}, which reflects the tendency of an agent to stay in its resource even if better alternatives are known to it.

It turns out that the above protocol maximizes an agents' inertia for problems of small dimensions, i.e., when few resources are involved. 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 even further up, resulting in an inertia of 1 at equilibrium. An inertia of 1 for all agents denotes convergence, because then none of them move.

We control inertia in our experiments using the parameter α. Inertia, as our experiments show, can facilitate coordination. A system whose agents have low inertia may exhibit chaotic behavior, and may never achieve coordination. On the other extreme, a very high inertia would lead to an inactive system.

2.2.5 Precision

Instead of defining convergence as precise convergence, we found it convenient to allow some imprecision. Imprecision is the distance from a perfectly coordinated state, i.e., the minimum number of agent relocations required to coordinate. Thus, a state would be deemed acceptable (and the simulation would halt) if the number of agents occupying each resource were within an acceptable level of imprecision.

Introducing imprecision 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 states in which the system may get stuck, e.g., if almost all of the resources were being used optimally, but say one of the resources was under-used and another, faraway resource was over-used. 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.

2.2.6 Delay in Updating Knowledge

Delay is introduced when each agent has access to the relevant state of the system, but only from earlier times. Such delays can have an undesirable effect on coordination. We introduce a delay τ in our system, quite like Huberman and Hogg's approach [1988], by allowing agents the knowledge about occupancy of resources in their knowledge window that is τ time steps old or outdated.

2.2.7 Scalability

Scalability of our setup can be viewed in terms of the total number of resources in the system and the average number of agents per resource. We study the role of the various attributes described above in enhancing the scalability of coordination in decentralized multiagent systems.

Chapter 3

Experimental Results

This chapter presents and analyses the results from our experiments. As mentioned in Chapter 2, our experimental framework generalizes over the one used by Sen *et al*. Whereas they consider only knowledge (which we find is coupled in their setup with choice), we consider several other key attributes. When these enhancements are eliminated, we do indeed achieve results similar to those of Sen *et al.*, but in light of our more extensive exploration, are forced to different conclusions.

3.1 Choice and Knowledge

Some of our experimental results are displayed in Tables 3.1 and 3.2. The tuple in each caption indicates, respectively, the number of resources, the number of agents, the initial deviation (the distance of agent distribution from a coordinated state), and the imprecision tolerated. Each result entry shown in the tables represents an average of 20 simulation runs.

Know-ledge	Choice								
	2	3	4	5	6	7	8	9	10
2	76	73	189	250	778	1060	2341	3793	12749
3		55	106	232	839	1020	2075	4684	9006
4			105	246	884	2321	4444	10658	12818
5				297	1006	1363	6375	8763	18616
6					1234	2403	5051	11036	16381
7						2061	5409	9056	38603
8							7491	14387	32242
9								18041	35064
10									45262

Table 3.1: Number of steps to coordination $\langle 10, 30, \pm 15, \pm 0 \rangle$

We always average the results over several runs, but it takes more runs for the results to be reliably duplicated if the imprecision is set low, especially for a problem of larger size. 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. Further, the results based on larger (tolerated) imprecision are generally more robust. This is because, as our experiments showed, in going from almost coordinated state to perfectly coordinated state, the system exhibits sustained oscillations for long periods of time. This point is elaborated upon below.

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.)

Know-ledge	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 3.2: Number of steps to coordination $\langle 16, 48, \pm 24, \pm 4\rangle$

As seen from the tables, the trends shown by the diagonal elements duplicate those achieved by Sen *el al.* However, for each of these elements, imparting more knowledge also implies imparting greater choice. Sen *el al.* attribute this trend (increasing the choice and knowledge simultaneously increases the time to coordinate) to knowledge alone, as knowledge and choice are tied together as a single variable in their setting. However, knowledge and choice are orthogonal concepts. Thus, we do not support their conclusion that increasing knowledge *alone* causes a loss of the effectiveness of coordination.

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 [Simon, 1996, p. 14]. We conjectured that the inherent symmetry in our problem might be causing this. When we tried to break this symmetry by offsetting the agents' choices and knowledge, however, it had no substantial effect on the above behavior. So we discarded that conjecture. We probe the reasons for this surprising behavior in the next section.

Observant readers would notice that the trends shown by the diagonal or even first row in Tables 3.1 and 3.2 are not monotonously increasing. Though such a behavior can occur because of randomizations, we found that this behavior was common to many more results than we show here. On further investigation, we found that one of the key attributes in our framework, inertia, caused this behavior. Inertia is controlled by parameter α as described in Chapter 2. We discuss inertia in detail later but suffice it to say here that a lower value of inertia produces the monotonously increasing trends that we expected, though, the values of individual elements in the tables increase considerably.

3.2 Sharing of Knowledge

When the local knowledge of agents is increased, an interesting hidden effect occurs. This is the amount of *sharing* that an 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 in the system. Once the sharing is factored in, we can explain the decreased effectiveness of coordination when we hold constant the extent of choices available to each agent.

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 total 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 $\Theta(k^3)$. This result leads us to the following hypothesis.

H1. Increasing the knowledge while holding the choice constant increases the convergence time proportional to the cube of the size of the knowledge window.

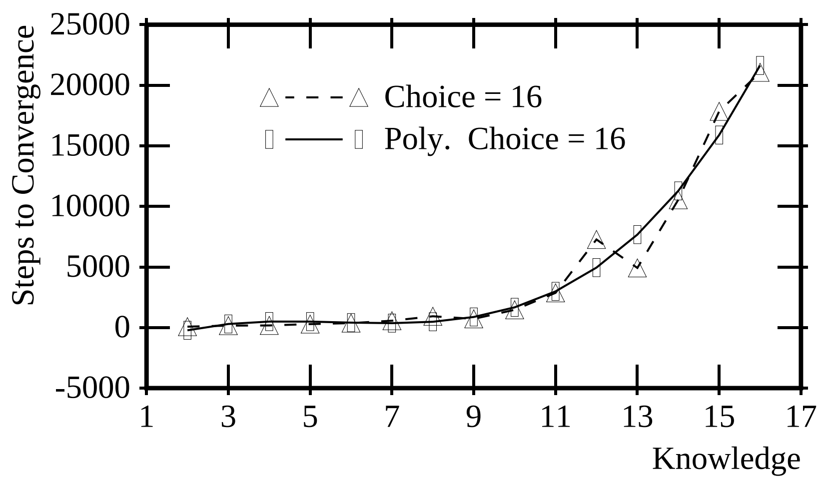

Figure 3.1: Effect of sharing of knowledge $\langle 16, 48, \pm 24, \pm 4\rangle$

Figure 3.1 which is based on the last column of Table 3.2, shows that the time to convergence indeed 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 column (constant, maximal choice) of Table 3.2. This figure indicates that sharing of knowledge 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.

Sharing and choice, however, are not the only reasons that slow down the coordination. Other attributes, especially inertia and imprecision act in concert, as demonstrated in the following sections. This is also evident from the last column of Table 3.1. The time to convergence, in this case, is an order or two lower than the sharing metric. This, however, is easily explained when we compare the problem sizes represented in Tables 3.1 and 3.2. Given the same values of parameters (α, β, and γ) for both sets of results, the inertia neutralizes the effects of sharing to a much lesser degree in the larger problem than it does in the smaller problem.

The inertia also has a bearing on the results of Table 3.2. Although, the time to convergence for the outer three columns has the same order as the sharing metric, it has a lower order than the sharing metric for the inner columns. This, however, is not surprising. Agents, in the case of inner columns, have lower choice and therefore higher inertia that neutralizes the effect of sharing to a greater extent than in the case of outer columns. The role of inertia in achieving coordination is studied in greater detail in Section 3.4.

3.3 Precision

In Chapter 2 we defined imprecision as the distance from a perfectly coordinated state, i.e., the minimum number of agent relocations necessary to lead the system from a given state to a perfectly coordinated state. Reducing the required precision in this manner, enhances the scalability of coordination. In other words, as the required quality of the coordination increases, the cost in time to coordinate becomes extremely high. We studied this observation further by delineating the effect of the deviation from coordination at the start of each simulation run. For a system with 30 agents and 10 resources, the deviation ranges from 1 (almost coordinated) to 15 (maximally uncoordinated).

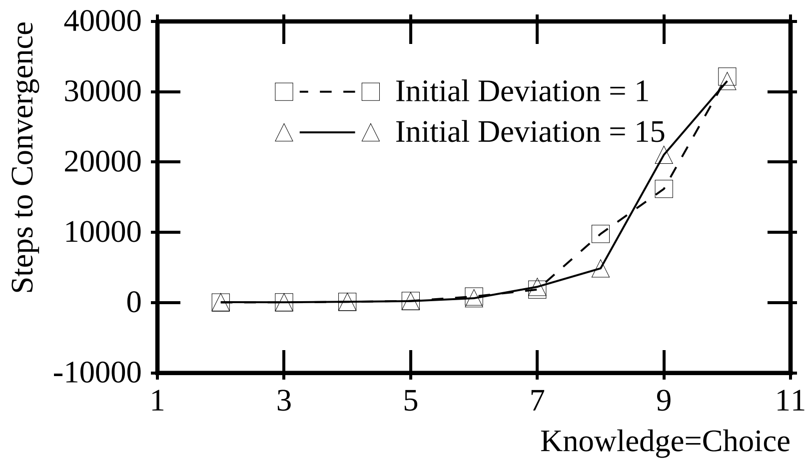

Figure 3.2: Effect of initial deviation $\langle 10, 30, \pm 0 \rangle$

Figure 3.2 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. In other words, 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 3.3 supports this claim. The exponential variation occurs, because during the last little bit of precision, the system begins to exhibit divergent properties (i.e., becomes increasingly unstable). We

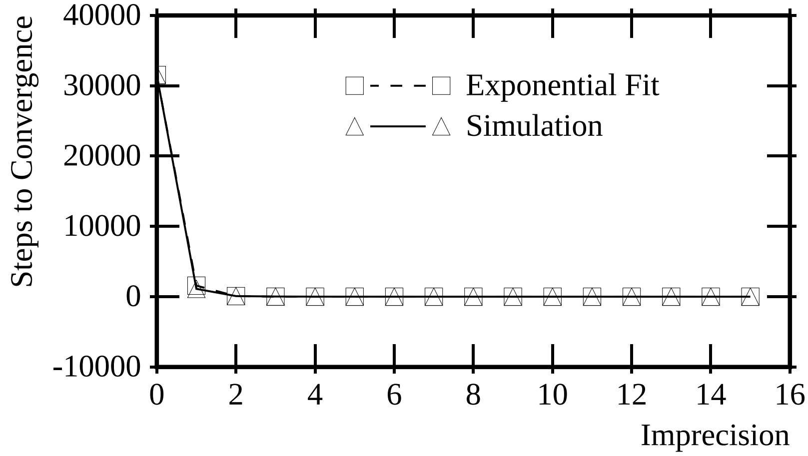

Figure 3.3: Effect of imprecision $\langle 10, 30, \pm 15 \rangle$

explore the reasons behind such a behavior in Section 3.6. 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 3.4 and 3.5 demonstrate, for small choice 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.

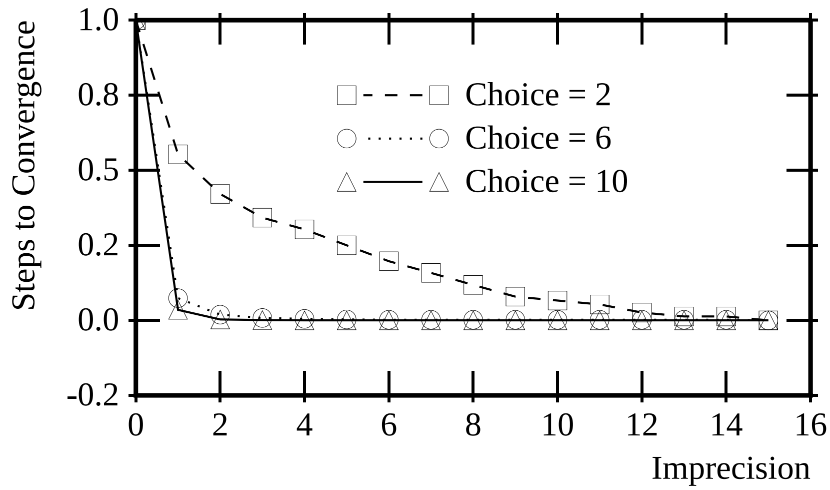

Figure 3.4: Effect of choice relative to imprecision $\langle 10, 30, \pm 15\rangle$ (each curve is normalized to 1)

3.4 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 2.1 in Chapter 2, it should be clear that, in general, as the number of choices increases, $\sum_j f_{ij}$ increases, and consequently the inertia (i.e., p_{ii}) decreases. This reason, especially when coupled with an imprecision of 0, can prevent coordination even for moderately large dimensions.

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 3.6 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

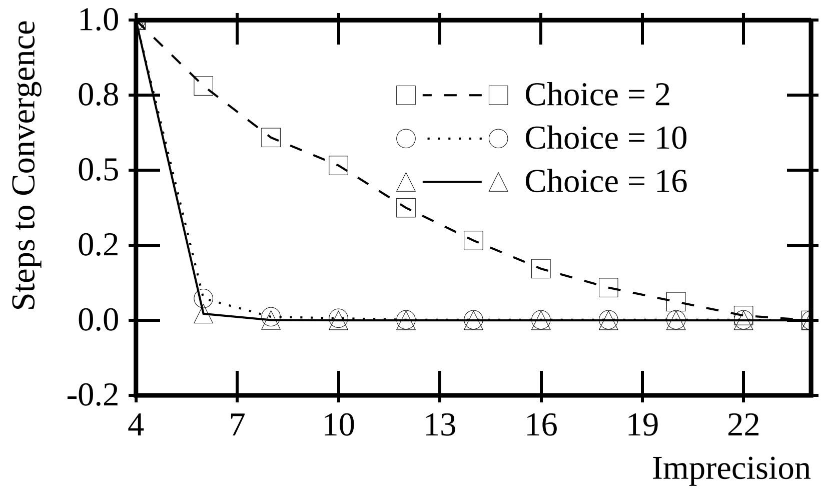

Figure 3.5: Effect of choice relative to imprecision $\langle 16, 48, \pm 24 \rangle$ (each curve is normalized to 1)

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

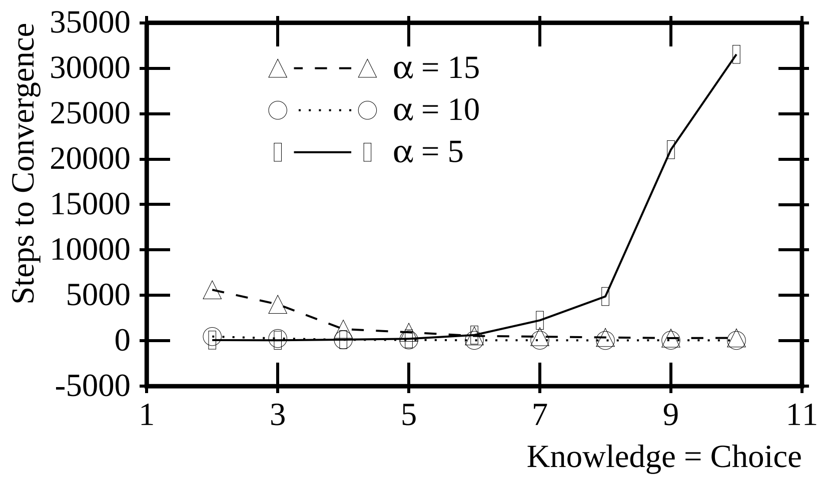

Figure 3.6: Effect of inertia (α) $\langle 10, 30, \pm 15, \pm 0 \rangle$

resources they occupy.

3.5 Delay in Updating Knowledge

Delay (τ) is introduced by allowing agents the knowledge about occupancy of resources in their knowledge window that is τ time steps old or outdated. The results obtained from our simulations with the introduction of delay in updating knowledge are summarized in Tables 3.3-3.5. These results show that for a given inertia (here $\alpha = 9$), delay in information access causes the coordination to slow down due to sustained oscillations in the system. This verifies a similar theoretical result obtained by Huberman and Hogg [1988].

Unlike Huberman and Hogg who show that oscillations (and hence a slowdown in

knowledge	choice								
	2	3	4	5	6	7	8	9	10
2	303	209	132	81	89	114	88	84	127
3		196	119	101	76	62	104	215	103
4			87	74	79	67	80	141	114
5				68	41	47	107	71	106
6					55	48	50	66	84
7						46	59	54	86
8							36	46	60
9								41	46
10									35

Table 3.3: Effect of delay, $\tau = 0$ $\langle 10, 30, \pm 15, \pm 0 \rangle$

knowledge	choice								
	2	3	4	5	6	7	8	9	10
2	309	227	158	213	286	301	523	526	1631
3		198	169	158	195	346	577	478	1213
4			95	173	217	289	500	763	1186
5				89	228	209	303	424	825
6					90	160	319	435	878
7						206	312	378	763
8							288	373	562
9								559	751
10									446

Table 3.4: Effect of delay, $\tau = 2$ $\langle 10, 30, \pm 15, \pm 0 \rangle$

coordination) occurs due to delays alone, we have shown earlier that such slowdown can also occur despite complete knowledge and no delays provided inertia is set to a low value. This behavior was not observed by Huberman and Hogg as they selected a high enough inertia in their formulation by setting the value of time interval to a sufficiently small value such that at most one agent changed its strategy in any timestep.

Conversely, Kephart *et al.* [1989] selected a low enough inertia through a parameter β, which in fact combines the effects of inertia with delay. However, as inertia is

knowledge	choice								
	2	3	4	5	6	7	8	9	10
2	320	234	365	354	706	1603	3637	5388	10105
3		303	243	351	875	1683	3200	5220	10129
4			162	430	663	1395	3023	10310	13690
5				449	809	1893	3135	8332	8288
6					907	1757	4451	8344	16041
7						5575	8312	14589	38881
8							15917	33174	47620
9								81961	72640
10									84863

Table 3.5: Effect of delay, $\tau = 4$ $\langle 10, 30, \pm 15, \pm 0 \rangle$

coupled with delay in their formulation, they do not explicitly study the above case.

Tables 3.3-3.5 show that apart from slowing down the coordination, delay also reverses the trends of their columns. Thus delay has an effect that in many ways is quite opposite to that of inertia. This intuition is captured by Kephart *et al.* through their parameter β in which delay and inertia are invertly related.

3.6 Scalability

Recall that reducing the required precision enhances the scalability of coordination, especially when a low value of inertia is used. Almost all the effort is consumed in progressing from a deviation of 1 (almost coordinated) to 0 (perfectly coordinated) as explained earlier based on Figure 3.2. Using a moderately high value of inertia also enhances the scalability of coordination. In this section, we explore the reasons for these behaviors, both analytically as well as empirically.

Given that most of the coordination effort is consumed in the last bit of coordination, we study analytically the setup shown in Figure 3.7 for the case where all the agents in the system have global knowledge. For a system with n resources and $a * n$

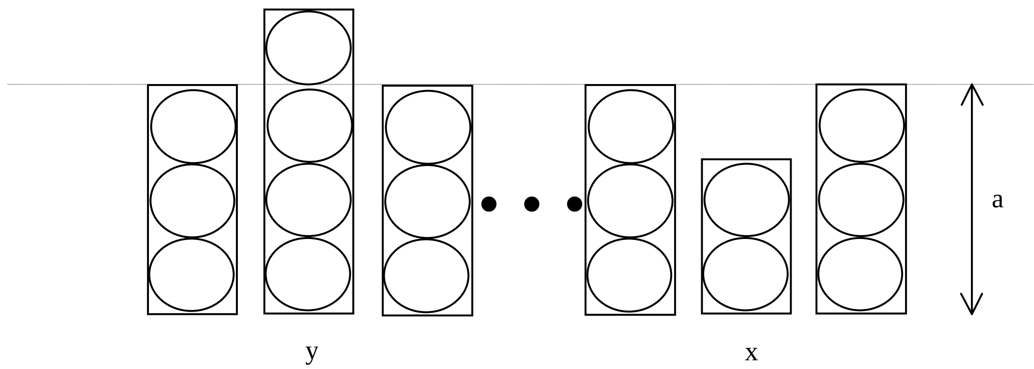

Figure 3.7: System setup used to explore scalability

agents, a being the average number of agents per resource, the sizes of knowledge and choice windows are fixed at n. The convergence characteristics of this system can be studied in terms of the expected number of agents in the various resources at the next timestep.

Based on our decision protocol, the $a(n-2)$ agents occupying $(n-2)$ resources in Figure 3.7 either stay at their current location or move to the resource x. The weights for these actions are given as

$$f_{ii} = 1$$

$$f_{ix} = 1 - \frac{1}{1 + \exp(\frac{1-a}{2})}, \forall i : i \neq x, i \neq y$$

and the probability of moving from resource i to resource x is given by

$$p_{ix} = \frac{f_{ix}}{1 + f_{ix}}$$

Resource x is occupied by $(a-1)$ agents who, due to our protocol, do not move in the next time step. The $(a+1)$ agents occupying resource y can either stay at their current location or move to any other resource. The weights for these agents are

$$f_{yy} = 1$$

$$f_{yx} = 1 - \frac{1}{1 + \exp(1 - \frac{\alpha}{2})}$$

$$f_{yi} = 1 - \frac{1}{1 + \exp(\frac{1-\alpha}{2})}, \forall i : i \neq x, i \neq y$$

and the resulting probabilities are given by

$$p_{yy} = \frac{1}{1 + f_{yx} + (n-2)f_{yi}}$$

$$p_{yx} = \frac{f_{yx}}{1 + f_{yx} + (n-2)f_{yi}}$$

The expected number of agents at resource x at the next time step is computed as follows. Let num_x be a random variable corresponding to the number of agents at resource x in the next time step. As num_x follows a binomial distribution, the expected value of num_x is computed as given below.

$$E[num_x] = (a-1) + a(n-2)p_{ix} + (a+1)p_{yx}$$

Similarly, the expected number of agents at resource y at the next time step is given by

$$E[num_y] = (a+1)p_{yy}$$

The expected number of agents at resource x, computed using the above equation with α equal to 4 and 10 are summarized in Tables 3.6 and 3.7, respectively. Table 3.6 shows that for low values of inertia (characterized by $\alpha = 4$), the expected occupancy of the least occupied resource x (with an occupancy of $a-1$ at the current time step) increases well beyond the average occupancy a. Further, the higher the values of a and n, the higher is the expected occupancy. These results indicate that a low inertia has negative implications for scalability. The argument in favor of high inertia for improved scalability is well supported by Table 3.7. Later in this section, we verify these conclusions, using empirical results obtained from our simulations.

	n		
a	10	16	20
2	3.7642	5.5309	6.7313
3	6.0970	8.7612	10.5675
4	8.4299	11.9915	14.4036

Table 3.6: Expected number of agents at resource x ($\alpha = 4$)

	n		
a	10	16	20
2	1.2227	1.3503	1.4356
3	2.3259	2.5178	2.6460
4	3.4291	3.6853	3.8564

Table 3.7: Expected number of agents at resource x ($\alpha = 10$)

Similarly, Tables 3.8 and 3.9 show the expected number of agents at resource y for α values of 4 and 10, respectively. However, these tables show that the expected occupancy of the most occupied resource y (with an occupancy of $a+1$ at the current time step) doesn't have negative implications for scalability with low or high inertia.

	n		
a	10	16	20
2	1.0996	2.1411	2.5192
3	1.4661	2.8547	3.3589
4	1.8326	3.5684	4.1986

Table 3.8: Expected number of agents at resource y ($\alpha = 4$)

We verify the above conclusions (drawn analytically), using the empirical results obtained from our simulations and shown in Tables 3.10-3.13. The results in these tables are based on a tolerated imprecision of 0. Given a value of α, the number of steps to coordination increase as the problem size increases (Tables 3.11 and 3.12). Further, as Tables 3.10-3.13 show, the number of steps to coordination decrease as α increases for each problem size.

a	n		
	10	16	20
2	2.7128	2.8312	2.8916
3	3.6170	3.7750	3.8554
4	4.5213	4.7187	4.8193

Table 3.9: Expected number of agents at resource y ($\alpha = 10$)

knowledge/ choice	a		
	2	3	4
2	85	84	79
3	81	68	58
4	62	38	91
5	57	70	127
6	49	152	540
7	76	192	1366
8	73	445	3821
9	99	834	13885
10	132	1632	27961

Table 3.10: Effect of scalability ($n = 10, \alpha = 6$)

3.7 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 polynomial 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.*
- We observed that keeping the knowledge and choice windows of an agent symmetrically distributed around its current resource yield the same trends

knowledge/ choice	a		
	2	3	4
2	210	184	166
3	169	130	116
4	60	63	61
5	59	53	50
6	39	68	40
7	34	54	75
8	42	57	57
9	41	55	135
10	33	44	113

Table 3.11: Effect of scalability ($n = 10, \alpha = 8$)

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 γ) in our protocol. The results highlighted an interesting interplay among the various bases of coordination. Varying α alone, however, provides representative results.
- We also calculated the average system utility based on Figure 2.2 of Chapter 2, for each of our runs. It was observed that the longer agents took to converge the lower the corresponding system utility was. Table 3.14,

knowledge/ choice	a		
	2	3	4
2	410	384	367
3	430	404	321
4	150	169	148
5	124	145	142
6	85	89	199
7	103	181	308
8	131	232	407
9	129	324	802
10	180	481	1520
11	189	574	3082
12	188	753	3808
13	205	1626	14571
14	231	1755	23855
15	374	2851	50523
16	618	4369	96834

Table 3.12: Effect of scalability $(n = 16, \alpha = 8)$

which corresponds to the results of Table 3.2 of Section 3.1, illustrates this observation in ample measure. This, however, was no surprise. Therefore, we did not always present those results and relied only upon the time to convergence to make our claims related to system performance.

knowledge/ choice	a		
	2	3	4
2	1090	987	880
3	1250	899	778
4	415	245	281
5	285	191	170
6	191	198	138
7	137	174	99
8	152	102	138
9	97	140	110
10	87	114	117
11	87	87	145
12	112	118	135
13	69	125	253
14	106	102	249
15	66	147	258
16	111	113	319

Table 3.13: Effect of scalability ($n = 16, \alpha = 10$)

Know-ledge	Choice							
	2	4	6	8	10	12	14	16
2	0.768	0.771	0.809	0.822	0.828	0.821	0.824	0.814
4		0.795	0.815	0.828	0.822	0.819	0.813	0.810
6			0.816	0.826	0.820	0.814	0.804	0.801
8				0.813	0.806	0.802	0.793	0.787
10					0.787	0.782	0.779	0.773
12						0.761	0.759	0.756
14							0.745	0.739
16								0.732

Table 3.14: Average system utility based Figure 2.2 $\langle 16, 48, \pm 24, \pm 4\rangle$

Chapter 4

Rules for Effective Coordination

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 previously been empirically studied in such systems. Others—knowledge and choice—have been studied but, as our analysis showed, the relationships involved are richer than believed. Trends due to inertia and precision can dominate and sometimes reverse the simpler trends due to knowledge and choice.

4.1 Mapping the Terrain

The following simple rules summarize our qualitative results.

R1. Low inertia & low imprecision $\Longrightarrow$ knowledge sharing governs behavior $\Longrightarrow$ local knowledge & limited choice perform better

R2a. Moderately high inertia $\Longrightarrow$ extent of knowledge or choice is less important

R2b. High imprecision $\Longrightarrow$ extent of knowledge or choice is less important

R3. Very high inertia $\Longrightarrow$ system inactivity

The above rules demarcate the most important regions of our terrain. Figure 4.1 illustrates the corresponding regions. Rule $R1$, supported by Figure 3.1, is mapped to Region I in Figure 4.1. To achieve effective coordination in this region, agents must limit their knowledge as well as choice. The results of Sen *et al.* lie within this region. Figures 3.4-3.5 and 3.6 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 4.1, imply that knowledge and choice are less relevant for coordination. Rule $R3$ is intuitively obvious and is represented by region III in Figure 4.1.

The value used for inertia in the different regions of Figure 4.1 must be adjusted to account for delay in updating knowledge. As pointed out in Chapter 3, the higher the delay the higher must be the corresponding value of inertia at each of the boundaries separating regions I, IIa, and III. The values of inertia and imprecision at each of the region boundaries in Figure 4.1 must also account for the problem size measured in terms of number of agents and resources in the system. The larger the problem size the larger should be the inertia and imprecision.

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 tolerance for 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

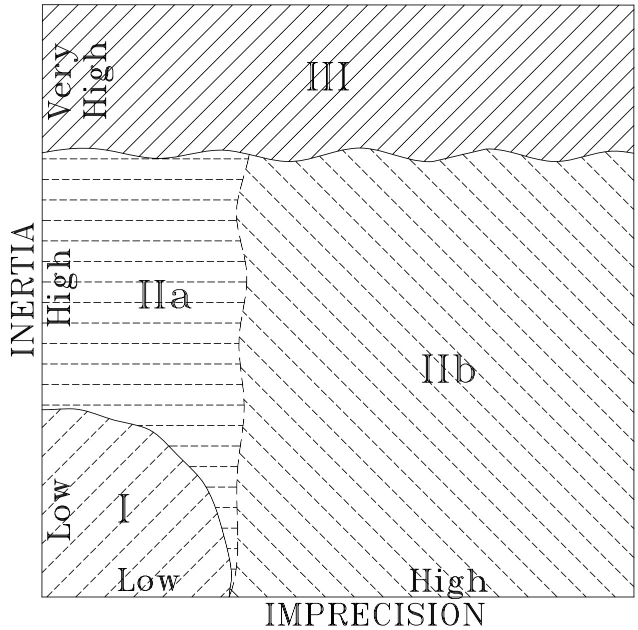

Figure 4.1: Mapping the terrain of decentralized systems

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.

Chapter 5

Conclusions

We explore the key features of decentralized systems and their relationship to coordination. For this purpose, we use an experimental framework that generalizes over the one used by Sen *et al.* The experimental results indicate that in particular five major attributes of decentralized systems 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. The subtle interplay of above attributes, not all of which were considered in most previous studies, makes coordination in decentralized multiagent systems a richer phenomena than previously believed.

The experiments show that for a low value of inertia, coordination slows down (due to sustained oscillations in the system) when the choices available to agents increase. When shared knowledge increases, then too coordination slows down. For a higher inertia, however, the above trends are seen to reverse. Inertia or higher patience by the agents tends to damp out the oscillations that cause the system to slow down.

It is also shown that perfect coordination is often inordinately more time-consuming than slightly imperfect coordination. Inertia together with imprecision, are shown to be crucial for the scalability of coordination. We verify this conclusion this both analytically and empirically.

Our experiments also verify the well-known but undesirable effects of uncertainty and delay in updating knowledge, on the coordination process. We show that for a given inertia, delay in updating knowledge causes the coordination to slow down. We, however, show that delay alone does not cause the oscillations and resulting slow down, as often believed. The chief culprit is a low value of inertia used by agents in their decision process.

Our experiments show that precision and inertia, primarily, 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.

5.1 Discussion

In this study, we developed empirical results about coordination in a simple setting involving multiple, potentially conflicting, autonomous agents. Despite its simplicity, our setup 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 conceptual conclusions were not supported.

Unlike Sen *et al.* and Huberman and Hogg [1988], we allow agents to choose among actions that are not necessarily rational (dictated by their knowledge). Further, Huberman and Hogg model knowledge differently from us as well as Sen *et al.* in that knowledge in their setup measures the accuracy with which an agent perceives (knows) its payoffs. This implies an assumption that agents are able to perceive their payoffs and are able to evaluate the payoffs for different strategies. We, like Sen *et al.*, model knowledge using a more direct, simpler approach.

We also capture the interplay of delay and inertia in our setup. Unlike Huberman and Hogg [1988] who show that oscillations and hence slowdown of coordination occurs due to delays alone, we have shown that such slowdown also occurs despite complete knowledge and no delays provided inertia is set to a low value. Kephart *et al.* [1989] select a low enough inertia in one of their parameters. However, as inertia is coupled with delay in their formulation, they do not explicitly study the above case.

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, the agents are non-adaptive in their decision-making protocol. Given the well-known limitations of reinforcement learning, 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. Further, these 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.

Given the simplicity of setup used, our results should not be taken as etched in stone. However, 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 [1995] 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.

5.2 Literature

In addition to the works mentioned before, 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 its total spending. 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. To achieve effective coordination in their system, one of the system parameters (α) is set to a very high value initially and subsequently reduced to a much lower value. In our setup this amounts to selecting a low inertia early on, which is then increased to a higher value later in the simulation, to facilitate convergence.

Chavez *et al.* [1997] introduce Challenger, a multiagent system that performs distributed resource allocation with each agent individually managing a local resource and relying only on locally available information. Challenger is similar to market-based system of Kuwabara *et al.*, in that the agents act as buyers and sellers in a marketplace, always trying to maximize their own utility. A key difference is in the addition of a learning behavior to Challenger's agents for allowing better performance under a wide range of conditions.

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 locations.

Knowledge has long been recognized as a key factor in AI planning and is often characterized as either domain knowledge or control knowledge. Domain knowledge describes the world and the actions that are available to the planner. Control knowledge indicates how the planner controls its search for a plan. Minton [1988] describes an explaination-based approach to learn search control knowledge in the planning system, PRODIGY. He observes that PRODIGY can sometimes learn search control rules that degrade rather than improve the system's performance. PRODIGY discards such useless even harmful rules, though the problem of discarding harmful knowledge is not trivial. This is yet another instance of how additional knowledge can sometimes be detrimental.

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.

When agents are loosely coupled, communication is often expensive relative to computation [Bond and Gasser, 1988]. This has led to active exploration of coordination techniques requiring little or no communication. Fenster *et al.* [1995] draw upon the intuition of communication-free human interactions to develop a focal point algorithm and simulate its applicability in randomly generated, albeit, simplified problem domains. The focal point algorithm applies to any world consisting of several objects each having various properties (measured using predicates) and in which the agents want to choose a common (same) object without communicating with each other. For each agent to select the same focal points, the authors make use of intuitive properties such as uniqueness (rarity), symmetry, and extremeness. This technique assumes complete and perfect knowledge of the world on part of the various agents though it allows local (independent) representation of this knowledge in each agent.

Rachlin *et al.* also show how agents can achieve coordination without explicit communication using their A-Team architecture [1998]. 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.

Shehory *et al.* present an approach to load balancing based on agent cloning [1998]. They treat load balancing problems by considering that agents are overloaded with tasks while the resources that the agents use may be idle. They implement agent cloning mechanism in their RETSINA infrastructure to remedy local agent overloads. Overloaded agents create new agents or clones to perform excess tasks using the unused resources on the system. To decide when to clone, a stochastic model of decision making based on dynamic programming is used.

5.3 Future Work

In closing, we mention several directions into which this work can be further expanded. To apply the gains of this work to a real-life setting requiring decentralized coordination, there is a need to develop adaptive approaches for proper selection of some of the key attributes discussed earlier. To this end, an approach such as one in which multiagent learning is supervised in certain ways, may be used.

The present work can be expanded to consider communication among agents—explicit or implicit—and to better characterize the circumstances under which it helps or disrupts coordination. Further, there is a need to consider dynamic systems in which both agents and resources are being removed and added back. Such a dynamic system is likely to further emphasize the role of inertia and imprecision.

The concepts in this work are illustrated using a decentralized resource allocation problem. While we do not investigate multiagent planning here, as sufficient pertinent literature exists, a future direction could extend the notions developed in this study to distributed problem solving and distributed constraint satisfaction problems in cooperative multiagent systems, to further test their applicability. The mapping of the various attributes from our work, to the above domain may be obtained as given below.

- Knowledge: the number of constraints exchanged among agents; the number of agents with whom an agent exchanges information determine its neighbourhood or the locality of its knowledge.
- Choice: corresponds to selection of constraints an agent chooses to satisfy (or adjust in concert with other agents or on its own).
- Inertia: corresponds to how much resistance agents have to satisfy (or adjust) constraints.
- Imprecision: the degree to which the global constraints are met (based on the number of unsatisfied constraints or a similar metric).

Bibliography

Robert Axelrod. *The Evolution of Cooperation.* Basic Books, New York, 1985.

Cristobal Baray. Effects of individual decision schemes on group behavior. In *Proceedings of the 3rd International Conference on Multiagent Systems (ICMAS)*, pages 389–390. IEEE Computer Society Press, 1998.

Alan Bond and Les Gasser, editors. *Readings in Distributed Artificial Intelligence.* Morgan Kaufmann, San Francisco, 1988.

Anthony Chavez, Alexandros Moukas, and Pattie Maes. Challenger: A multiagent system for distributed resource allocation. In *Proceedings of the International Conference on Autonomous Agents*, pages 323–331. ACM Press, 1997.

Scott Clearwater, editor. *Market-Based Control: A Paradigm for Distributed Resource Allocation.* World Scientific, 1996.

Paul R. Cohen. *Empirical Methods for Artificial Intelligence.* MIT Press, Cambridge, MA, 1995.

Keith S. Decker and Victor R. Lesser. Designing a family of coordination algorithms. In *Proceedings of the International Conference on Multiagent Systems*, pages 73–80, 1995.

Edmund H. Durfee. *Coordination of Distributed Problem Solvers.* Kluwer, 1988.

Edmund H. Durfee. Distributed problem solving and planning. In *Weiß [1999]*, chapter 3, pages 121–164. 1999.

Maier Fenster, Sarit Kraus, and Jeffrey S. Rosenschein. Coordination without communication: Experimental validation of focal point techniques. In *Proceedings of the First International Conference on Multiagent Systems*, pages 102–108, 1995.

Les Gasser and Randall W. Hill Jr. Engineering coordinated problem solvers. *Annual Review of Computer Science*, 4:203–253, 1990.

Joseph Y. Halpern and Yoram O. Moses. Knowledge and common knowledge in a distributed environment. *Journal of the Association for Computing Machinery*, 37: 549–587, 1990.

Tad Hogg and Bernardo A. Huberman. Controlling chaos in distributed systems. *IEEE Transactions on Systems, Man, and Cybernetics*, 21(6):1325–1332, 1991.

Bernardo A. Huberman and Tad Hogg. The behavior of computational ecologies. pages 77–115. Elsevier Science Publishers/North-Holland, 1988.

Francois F. Ingrand, Michael P. Georgeff, and Anand S. Rao. An architecture for real-time reasoning and system control. *IEEE Expert*, 7(6):34–44, 1992.

Jeffrey O. Kephart, Tadd Hogg, and Bernardo A. Huberman. Dynamics of computational ecosystems: Implications for dai. In L. Gasser and M. N. Huhns, editors, *Distributed Artificial Intelligence, Volume II*, pages 79–96. Pitman/Morgan Kaufmann, London, 1989.

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.

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 *Clearwater [1996]*, pages 53–73. 1996.

David K. Lewis. *Convention: A Philosophical Study.* Harvard University Press, Cambridge, MA, 1969.

Jyi-Shane Liu and Katia Sycara. Distributed problem solving through coordination in a society of agents. In *Proceedings of the 13th International Workshop on Distributed Artificial Intelligence*, 1994.

T. W. Malone and K. Crowston. The interdisciplinary study of coordination. *ACM Computing Surveys*, 26(1):87–119, 1994.

S. Minton. *Learning Search Control Knowledge: An Explaination-based Approach.* Kluwer Academic Publishers, Boston, 1988.

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, 1998.

Sudhir K. Rustogi and Munindar P. 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 149–161. Springer-Verlag, 1999a.

Sudhir K. Rustogi and Munindar P. Singh. Be patient and tolerate imprecision: How autonomous agents can coordinate effectively. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 1999b.

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.

Sandip Sen, Neeraj Arora, and Shounak Roychoudhury. Using limited information to enhance group stability. *International Journal of Human-Computer Studies*, 26 (1):69–82, 1998.

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.

Sandip Sen, Mahendra Sekaran, and John Hale. Learning to coordinate without sharing information. In *Proceedings of the National Conference on Artificial Intelligence*, pages 426–431, 1994.

Onn Shehory, Katia Sycara, Prasad Chalasani, and Somesh Jha. Increasing resource utilization and task performance by agent cloning. In *Intelligent Agents V: Proceedings of the 5th International Workshop on Agent Theories, Architectures, and Languages (ATAL-98)*, pages 413–426. Springer-Verlag, 1998.

Herbert Simon. *The Sciences of the Artificial.* MIT Press, Cambridge, MA, third edition, 1996.

Munindar P. Singh, Anand S. Rao, and Michael P. Georgeff. Formal methods in DAI: Logic-based representation and reasoning. In *Weiß [1999]*, chapter 8, pages 331–376. 1999.

Gerhard Weiß, editor. *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence.* MIT Press, Cambridge, MA, 1999.