Toward an Automatically Generated Theory of Coordination — Empirical Explorations

Eun Gyung Kim[1,2] and Nicholas V. Findler[1]

Abstract. We define the role and mechanisms of coordination in Intelligent Agent Societies (IASs). We then outline our approach and the underlying design principles aimed at the automatic generation of a theory of coordination. Such theory would assist in designing new IASs and provide trouble-shooting tools for suboptimally functioning IASs. We also describe in this paper the decisions that have been made in this endeavor. We have been able to show via a simplified model that the approach is feasible and can produce results. The project is one of the few examples of work in Experimental AI.

1 INTRODUCTION

Social structures may enhance the coordination of agent activity, such as message management, and the allocation of resources and tasks. Such structures are alliances, coalitions, teams and markets of which only the first grouping is considered for the time being. The structures are external to and independent of individual agents, and would allow the scaling-down of complex systems consisting of large number of agents. By reducing the danger of combinatorial explosion in dealing with the problems of agent cognition, cooperation and control, we expect to be able to manage the emergent behavior of individual agents and of alliances of agents.

An *alliance* is a temporary group formed voluntarily by agents whose goal structures (AND-OR trees) are similar enough. The agents give up, while in the alliance, some of their own goals and fully cooperate with the other members of the alliance. They stay in the alliance as long as it is in their interest, after which they may join another alliance or stay on their own. (It is possible to impose some payment as an entry fee into and some penalty for departing from the alliance.) Two or more alliances may also merge into one, based on a membership vote (counting on majority or plurality).

Further, cognitive activities — in particular, decision making, plan generation and execution — are usually performed jointly by groups of agents relying on distributed knowledge, skills and resources. There is, symbolically speaking, a single group to deal with and to respond to, instead of an indefinitely large number of individual agents. This mode of operation leads to higher efficiency as well as to the possibility of graceful degradation; i.e., whenever a small number of operating units become dysfunctional, other units can take over their responsibilities while the whole system does not crash but produces useful results, perhaps at a slower pace and of lower quality.

It will be helpful to provide a precise definition of *emergent phenomena* since its interpretation varies in different disciplines. The term refers to the appearance of patterns (of properties, actions, results, information, knowledge) that are not apparent at lower levels. Individual agents and groups of agents may well be aware of the possibility of emergence and could strive to enhance or to diminish it and its effects. We can talk about a *reasoning horizon* within which agents can predict emergent phenomena and their effects. A usually more limited domain is the *control horizon* within which the agents can successfully influence higher level events by lower level activity. (In the project discussed, only the control horizon is relevant.) Notice that the formation of groups of agents, the actions jointly decided upon and the norms are emergent phenomena because they are not under the control of individual agents.

Finally, we note that one can envisage an arbitrary number of levels of abstraction at which emergent phenomena may occur whereas in most traditional disciplines, such as thermodynamics, there is a distinction between only two, a macro and a micro level.

2 THE OBJECTIVES OF COORDINATION

Coordination is understood to consist of a set of mechanisms necessary for the effective operation of IASs. It is defined as the process of managing dependencies between activities [1,2]. Its fundamental components include the allocation of scarce resources, communication between the agents about intermediate results, coordination goals, capabilities and plans, status of the different aspects of the environment, ordering/suggesting/accepting solution methods to sub-problems, asking for and offering help, providing some meta-level information (e.g., optimum message routing strategies under different conditions), etc.

Coordination is needed and is usually available also in those cases in which there is no full cooperation among the agents or groups of agents. In a human society, for example, competition is constrained by consumer protection, various government agencies and anti-trust laws. People and organizations antagonistic to one another may interact via prescribed legal channels. Coordination theory can be defined as a set of axioms, mathematical and logical constructs, and analytical techniques used to create a model of dependency management in IASs. We have been engaged in a project to create an experimentally obtained theory of coordination.

3 ON THE AUTOMATIC GENERATION OF A THEORY OF COORDINATION

Our investigation is based on an easily modifiable and parametrizable generic IAS, the *P-System* (*P* stands for production), a metaphorical and abstract version of our earlier system, the Distributed Control of Nationwide Manufacturing Operations [3,4,5]. The P-System, sharing characteristic properties with most, if not all, IASs, is used for a series of statistically designed experiments. In the course of running the P-System under different conditions, we observe and measure data from which certain high-level, emergent variables can be created. We infer, from the statistical analysis of the data, characteristic and important descriptors of the organization and functioning of IASs. It is expected that the resulting relationships between the evaluated observations and the respective properties of the agent societies may produce a satisfactory theory of coordination, which in turn would generate design tools and guidelines for the construction of new IASs, and trouble-shooting tools for analyzing existing IASs.

The approach is analogous to theory formation in Physics where experimental results may suggest novel conceptual frameworks that have relevance to phenomena beyond those appearing in the original experiments. In other words, we hope that the theory being developed will help in understanding coordination *in general*, as well as form the basis of models of coordination *for specific applications*. We note

[1]Department of Computer Science & Engineering, Arizona State University, Tempe, AZ, USA.
[2] On sabbatical leave from the Korea University of Technology and Education, School of Information Technology, Chonan P.O.BOX 55, Chonan, Chungnam 330-600, Republic of Korea.

that related empirical work has been done by several researchers, such as the pioneering projects on the evolution of cooperation by Axelrod and associates [6,7].

The prototypical *P-System* creates the following environment:

- There is a critical dependency, possible conflicts and contests between subsystem controller agents.
- Communication between agents is asynchronous. Messages can be broadcasted at large, or sent to selected groups of agents or to an individual one on the basis of need-to-know and qualified-to-know.
- The sequence of manufacturing operations of a given product defines a hierarchical network of tasks, the P-tree, that corresponds (is homomorphic) to the problem solving network needed by the planning process (see Figure 1). Leaf nodes reference raw materials or sub-components provided by other producers. Higher-level process nodes correspond to manufacturing/assembly operations. Each node may also be associated with an OR-subtree (alternative tasks can accomplish the given job), but are not shown in Figure 1.

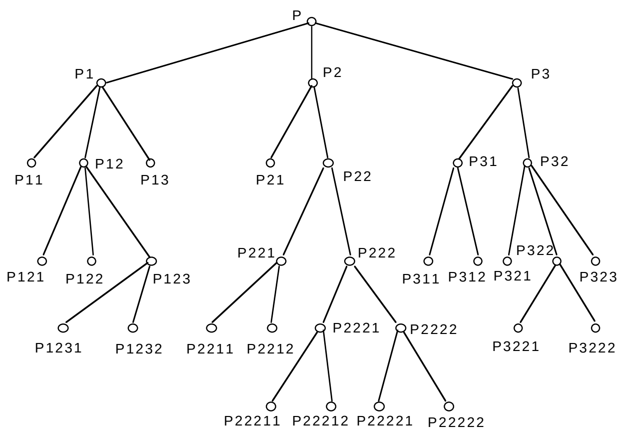

Figure 1. The metaphorical production plan is an AND-tree

- Planning is equivalent to assigning resources to the (metaphorical) manufacturing and assembly operations over space and time. An agent with a higher-priority task (see below) can obtain a needed resource from another agent with a lower-priority task. The latter task can be performed with a less satisfactory resource (more expensive or slower) or shuttling a resource between two nodes at opportune time points (see Resource Availability cases below).
- There may be priority-oriented or deadline-oriented tasks. The former implies that each task associated with the completion of a component must be well-coordinated with the completion of its sibling, ancestor and descendent components. The deadline-oriented categories of tasks have a deadline by which they have to be completed to satisfy the constraints of the final product.
- *Availability of resources* (tools for the assembly/ manufacturing operations) may change intermittently or regularly. Idle resources already allocated and the (temporary) storage of components also cost money. The total range of resource availability has four subranges: (1) *infeasible*: the production process cannot function for lack of indispensable resources; (2) *deficient*: the production process can function only if some resources are transferred between process nodes at opportune moments (balancing the costs of transfer and component storage); (3) *scarce*: some tasks must be allocated suboptimal resource types; (4) *abundant*: every task can be allocated the optimum resource type.
- There are three tiers in resource taxonomy: (1) resource category (every item in it can be used for the same task or tasks); (2) a resource category contains one or more resource types (a given task can be performed by each but at different cost or time levels, depending on the type chosen); (3) one or more resource instances exist within each resource type these are the ones actually allocated to tasks.
- Agents are associated with each process node; each resource category and, within it, each resource type; each task category and, within it, each task type. Communication among them, over limited bandwidth, includes requests for and provision of information, resource or action; allocation of task or resource to agents; a piece of information; an acknowledgment; etc.
- *Manager Agents* have different responsibilities. A top-level *Monitor Manager* collects and processes information from the other Managers, and stores it in its knowledge base. The *Message Manager* intercepts each message, records the agent IDs that originate, transmit and receive it, as well as categorizes and stores the messages according to their contents type in a *Message Database* (crucial for the development of the trouble-shooting tool.) The *Coordinative Process Manager* is concerned with solution synthesis, reinforcement (e.g., the support in the coordinator-coworker relationship), and scheduling processes (tasks and resources). The *Manager of Negotiating Processes* assesses how agreements and decisions are made and kept. The *Manager of Neutral Processes* observes the cost and the effect of learning processes. The *Constraints Manager* identifies the cost/benefit ratio of inherent and imposed constraints (capabilities, classes, timing, costs, capacities, resource availability). The collected information is processed by the *Statistical Analyzer Manager*.
- Two different objective functions can be used. The P-System is to produce a given number of final products either (1) at a minimum cost within a "reasonable" period of time, or (2) at a "reasonable" cost within a minimum period of time. Either of these requires an optimum allocation schedule of the manufacturing/assembly operations and resources to individual agents over space and time, while satisfying a set of constraints.

In order to attain a high-level of generality, we have originally defined 25 tweakable entities (TEs) that characterize tasks, agents, resources, skills, production processes, relative cost functions, events and constraints. For each experiment (a run of the P-System), particular TE values are automatically selected by the Experimental Design Generator according to a multi-tier, balanced, incomplete, factorial design (see Section 4.2).

4 EMPIRICAL EXPLORATIONS

Significant effort has been spent on identifying a reliable but combinatorially not explosive technique to obtain results that can show the method of computation and prove the feasibility of the approach. We list some of the decisions made along this line.

4.1 The Quality Measure of Coordination

After experimenting with some preliminary choices, we have defined the Quality Measure of Coordination (QMC), based on the concept of synchronization and supply balancing, as

$$QMC = \frac{\sum_{(l)} l . \sum_{(j)} t_{lj}^{*} / t_{lj}}{\sum_{(l)}\sum_{(j)} l_j}$$

Here l is the level number, lj references the j-th process node from the left at level l, t_{lj}^{*} is the best possible time associated with assembly/manufacturing at node lj, and t_{lj} is the actual time used after local and global optimization (these terms refer to systematic resource exchanges when needed). The weighting factor in the numerator, l, expresses the fact that deficient synchronization and supply balance has a detrimental, cascading effect on coordination at the levels above the process node in question. Thus the lower the node level, the more serious the effects are. The denominator normalizes QMC to the range [0,1].

The optimization of QMC under different conditions and its relationship to judiciously chosen functional combinations of TEs, the emergent variables, play a central role in this project.

4.2 The Statistical Design of the Experiments

As noted before, the TEs determine the functional and operational characteristics of the P-System and produce emergent behavior. Each emergent variable, a particular TE or the combination of a few, represents an *aspect* to be discussed later. Most TEs are quantitative. However, there are a few categorical variables of qualitative nature; e.g., whether resource and task allocation is preplanned or reactively performed during the run of the P-System. To develop a demonstration system, we have drastically reduced the number of TEs to 12. These can be resource-, task- or inventory-related. The choice was partly based on multiple regression and correlation analysis. We have also left out all event- and agent-related TEs, and postponed the study of reactive planning and message traffic until later.

The statistical design of experiments needed some novel ideas. The design process consists of two phases. The first phase deals with a subset of the Total Number of Qualitative Designs of experiments, TNQD, which we will call NQD. We identify five equally-sized *qualitative* subranges of the total quantitative range of each TE: very low, low, medium, high, very high. Thus, while TNQD = 5^{12} (there are 12 quantitative TEs currently), an arbitrary subset of this can be obtained by generating NQD random numbers between 1 and TNQD, each pointing to a qualitative design selected. (This heuristic must be used since there are no symmetrical, balanced, fractional experimental designs available for such high number of control variables.) The second phase of the design process leads to actual quantitative computer-based experiments. Every qualitative design obtained in the first phase is used to carry out a sufficiently large number of quantitative experiments (NQE); i.e., runs of the P-System. We assign randomly chosen quantitative values to each TE within their respective current qualitative subrange. The QMC values are computed while the objective function can be minimum total time or minimum total cost. This approach assures a high level of generality in the findings.

An interesting heuristics leads to the meaningful ratio between NQD and NQE. The problem boils down to the question: in order to obtain optimum precision in the results, given the total length of computing time, how much time should be spent in considering different qualitative designs and how many quantitative experiments should be performed for each qualitative design. The basis of the heuristics is that the total variance of the resulting QMC values is made up of three components: the first one is due to the different qualitative designs selected, the second one is due to the different actual quantitative experiments that belong to the same qualitative design, and a third one is a random scatter not related to any manageable factor. The calculation of the first two components is now explained by the following example. Assume that we have NQD= 4 (four qualitative deigns are selected) and NQE=3 (each qualitative design has three actual quantitative experiments associated with it). In the following Table 1, QMC is denoted by Q.

The variance due to different qualitative designs is Var(QD) = Var(Av(Q_{ij})), which is the variance of the items in the third column. Further, the variance due to different quantitative experiments is Var(QE) = Av(Var(Q_{ij})), which is the average of the items in the fourth column. The ratio Var(QD)/Var(QE) should be the ratio of the times allocated to QD and QE, respectively — the idea being that the larger the variance of an item is, the more time should be spent on its measurement.

4.3 Other Miscellaneous Decisions

We had to decide on the number of final products (NFP) to be produced by the P-System in the experiments while QMC having a fairly steady value beyond NFP. There are several factors that may cause difficulties in this regard. Such are the "running-in" and "running-out" times for the P-System. During the former period, the assembly/manufacturing operations start at the bottom of the P-tree and gradually reach its root from where the first final product leaves. From this point on, the P-System works continuously and the completion rate of the final product should be constant. (One must bear in mind, the Java has some overhead time used for various house-keeping chores and garbage collection that add a random component to processing time at irregular intervals.) Similar issues arise during the running-out phase when node activity gradually disappears from the bottom of the tree upwards. We have found that continuous and steady production is experienced after the first three final products leave the root of the P-tree; thus NFP=4 was chosen.

It was important to characterize each member of the set of experiments by the Resource Availability category it belongs. We can then make relevant conclusions concerning the effects of less than abundant resources — the infeasible, deficient and scarce cases — on production time and cost aspects.

5 THE COMPUTATIONAL APPROACH

After several shorter trial runs, we had a realistic run over 4 full days on an IBM PC with 500 MHz clock speed. Having identified the ratio r=Var(QD)/Var(QE)~60, we specified NQD= 11,000 (out of which 2,400 qualitative designs proved to be in the *Infeasible Resources* category). Instead of the "expected" NQE=871,200, we obtained 860,282. In addition to the 12 TEs considered in the experiments, the successfully completed experiments were also assigned a value of the categorical variable RA (Resource Availability) as deficient=1, scarce=2 or abundant=3. The average execution time of a qualitative experiment was 334.51 ms and the size of the output file was 61.8 MB.

Table 2 describes the currently used list of TEs, the RA and QMC, their notation (the *x'*s), and their role in the P-System. We emphasize that the reduced set of independent variables is not sufficient to produce acceptable results because of their numerous latent relationships among themselves and with the dependent variable QMC. As stated before, our major objective in the work described is to prove the feasibility of the method.

The Principal Component/Factor Analysis program and other parts of the SPSS statistical package have produced, among others, the following relevant results: (1) A 14x14 Correlation Matrix of the 13 independent variables and QMC. (2) Six factor*s*, *F*'s, linear combinations of the *x*'s, have zero correlations among themselves and high correlations with QMC. (3) A series of regression functions, *models*, that connect QMC, on one hand, and — via the *F* factors — the *x* variables, on the other. The final, 24-th, model, chosen by us, has the highest number of terms, 24, each with statistically significant correlation with QMC and no cross-correlation among the terms.

The six mutually orthogonal factors "explaining" the QMC are:

$F_1=b_1 + b_2.X_{13} + b_3.X_3$	$b_1 = 0.736$,	$b_2 = 0.123$,	$b_3 = -0.017$,
$F_2=b_4 + b_5.X_2 + b_6.X_1$	$b_4 = 0.891$,	$b_5 = -0.471$,	$b_6 = -0.010$,
$F_3=b_7 + b_8.X_5$	$b_7 = 0.775$,	$b_8 = -0.004$,	
$F_4=b_9 + b_{10}.X_{11}+ b_{11}.X_8$	$b_9 = 0.773$,	$b_{10} = 0.009$,	$b_{11}= 0.003$,
$F_5=b_{12} + b_{13}.X_9$	$b_{12} = 0.778$,	$b_{13} = -0.016$,	
$F_6 =b_{14}+ b_{15}.X_7$	$b_{14} = 0.775$,	$b_{15} = -0.004$.	

Table 1. Illustration for the decision on how to divide up total processing time

Qualitative Design	QMC Values in Quantitative Experiments	Average of QMCs	Variance of QMCs
1	Q_{11}, Q_{12}, Q_{13}	Av(Q_{1j})	Var(Q_{1j})
2	Q_{21}, Q_{22}, Q_{23}	Av(Q_{2j})	Var(Q_{2j})
3	Q_{31}, Q_{32}, Q_{33}	Av(Q_{3j})	Var(Q_{3j})
4	Q_{41}, Q_{42}, Q_{43}	Av(Q_{4j})	Var(Q_{4j})

Table 2. The dependent variable, the independent variables and their system function

Variables	Notation	System Function
TE1	x_1	Number of Resource Categories
TE2	x_2	Number of Resource Types per Resource Category
TE3	x_3	Number of Resource Instances per Resource Type
TE4	x_4	Cost of Resource Type working per task difficulty and per unit time
TE5	x_5	Time necessary for Resource Type to accomplish work per task difficulty
TE6	x_6	Time necessary to transfer a Resource Instance between two process nodes
TE7	x_7	Cost necessary to transfer a Resource Instance between two process nodes
TE8	x_8	Number of Task Categories
TE9	x_9	Number of Task Types per Task Category
TE10	x_{10}	Number of skills OR-ed per Task Type
TE11	x_{11}	Storage cost per subcomponent piece per unit time
TE12	x_{12}	Maximum storage size
RA	x_{13}	Resource Availability (categorical variable)
QMC	Q	Quality Measure of Coordination

The Regression Function of Model 24 is accepted as:

$$Q = a_0 + a_1.F_1 + a_2.F_2 + a_3.F_3 + a_4.F_4 + a_5.F_5 + a_6.F_6 + a_7.F_1{}^\wedge 2 + a_8.F_2{}^\wedge 2 + a_9.F_3{}^\wedge 2. + a_{10}.F_4{}^\wedge 2 + a_{11}.F_5{}^\wedge 2 + a_{12}.F_6{}^\wedge 2 + a_{13}.F_1.F_2 + a_{14}.F_1.F_3 + a_{15}.F_1.F_4 + a_{16}.F_1.F_5 + a_{17}.F_1.F_6 + a_{18}.F_2.F_3 + a_{19}.F_2.F_4 + a_{20}.F_2.F_5 + a_{21}.F_2.F_6 + a_{22}.F_3.F_4 + a_{23}.F_3.F_5 + a_{24}.F_3.F_6 + a_{25}.F_4.F_5 + a_{26}.F_4.F_6 + a_{27}.F_5.F_6$$

with

a_0= 0.784, a_1=0.161, a_2=-0.343, a_3=-0.135, a_4=-0.137, a_5=0.150, a_6=-0.047, a_7=0.007, a_8=-0.059, a_9=-0.033, a_{10}=-0.022, a_{11}=-0.02, a_{12}=0.000, a_{13}=-0.04, a_{14}=0.215, a_{15}=-0.175, a_{16}=-0.218, a_{17}=0.053, a_{18}=0.006, a_{19}=0.00, a_{20}=-0.003, a_{21}=-0.003, a_{22}=0.040, a_{23}=0.027, a_{24}=0.00, a_{25}=-0.029, a_{26}=0.011, a_{27}=0.120.

Based on

$$\frac{\partial}{\partial x_k} Q(F_1,\ldots F_6) = \sum_{i=1}^{6} \frac{\partial Q}{\partial F_i} \frac{\partial F_i}{\partial x_k}$$

we have set each equation (k=1, 13) to zero:

$$\begin{array}{ll}
b_6\,(a_2 + a_{13}.F_1 + 2a_8.F_2 + a_{18}.F_3 + a_{20}.F_5 + a_{21}.F_6) & = 0 \\
b_3\,(a_1 + 2a_7.F_1 + a_{13}.F_2 + a_{14}.F_3 + a_{15}.F_4 + a_{16}.F_5 + a_{17}.F_6) & = 0 \\
b_8\,(a_3 + a_{14}.F_1 + a_{18}.F_2 + 2a_9.F_3 + a_{22}.F_4 + a_{23}.F_5) & = 0 \\
b_{15}\,(a_6 + a_{17}.F_1 + a_{21}.F_2 + a_{26}.F_4 + a_{27}.F_5) & = 0 \\
b_{11}\,(a_4 + a_{15}.F_1 + a_{22}.F_3 + 2a_{10}.F_4 + a_{25}.F_5 + a_{26}.F_6) & = 0 \\
b_{13}\,(a_5 + a_{16}.F_1 + a_{20}.F_2 + a_{23}.F_3 + a_{25}.F_4 + 2a_{11}.F_5 + a_{27}.F_6) & = 0
\end{array}$$

The above equations can be solved for the *F*'s and, subsequently, for the variables x_i by using, for example, the mathematical programming package Maple. Checking whether *Q* is maximum is a little more complicated but several techniques are available in the literature on numerical optimization.

A few more words should be said about certain plausible *aspect*s of agent behavior in Intelligent Agent Societies. Their numerical evaluation will be possible when we introduce additional TEs that appear to be potentially relevant to the behavior of the P-System. The most plausible aspects in an intuitive order are:

Autonomy and power: This reflects the degree to which agents are controlled by other agents. Exogenous factors may give certain agents the ability to enforce their decisions on other agents (i.e., forced cooperation). This phenomenon is most clearly apparent in studies of political and other social systems [8-11]. It introduces the element of *force majeure* into coordination problems and makes coordination processes somewhat simpler. Relative power affects subtask assignment, agent strategies and interaction methods.

Accuracy of knowledge about other agents: Agents in an IAS generally build and uses models of other agents to forecast others' actions [12-14]. The more accurate the model, the greater the ability of the agent to predict global requirements and tailor its own activity accordingly. Accurate knowledge allows the adoption of effective interaction strategies and improved conflict resolution.

Connectivity: The connectivity of an agent can be defined as the number of other agents with which it can interact [15,16]. The higher the connectivity of agents, the likelier they are to obtain the help/information they need, and the better the chances of efficient coordination and optimal solutions to problems. Connectivity thus directly affects coordination.

Agreement between agents' goals: The relations between the goals of different agents are important because they underlie the interactions between agents (interaction with the purpose of achieving goals). Extremes may range from a cooperative problem-solving system where all agents will cooperate on a common goal to a purely competitive system where each agent seeks to maximize its expected utility, at the expense of any and all other agents [9,17,18].

Coverage: The knowledge particular agents have, the processes they can execute, and the data they have access to, all influence the assignment of tasks to agents. The coverage of an IAS relative to a task reflects the amount of the task that can be performed by the society [19]. If many agents can perform a task, the coverage can be said to be higher than if only one or a few agents can do it. The higher the coverage, the more flexible the system and the fewer the constraints on task assignment and similar coordination problems.

*Cost versus Time***:** We will be able to identify the Cost function when Total Time is minimized and the Time function when Total Cost is minimized — an important practical issue when realistic compromises are sought between time and cost.

6 CONCLUSIONS

In general, we can state that coordination is a combination of a variety of mechanisms aimed at substituting for the unattainable perfect world of complete and uptodate knowledge of goals, plans, actions and interactions as well as of agents' unlimited processing and communication power. This is done by means of an appropriate and adaptive organizational structure (well-balanced division of labor and flexible interaction among agents), exchanging meta-level information (e.g., control information, planning methods, credible commitments, joint model building of the environment), and reducing logical coupling and resource dependencies of agents (effective techniques for task allocation, resolving resource conflicts and logical contradictions, and the like).

In this exploratory work, we have outlined and proven the feasibility of an approach that can lead to quasi-optimum coordination in a characteristic subset of Intelligent Agent Societies. Because of the limited number of control variables (TEs) incorporated in the system at this stage, we do not present the quantitative aspects of their role in the management of coordination.

7 RELATED WORK

Much of the existing research in DAI adopts a solution-oriented approach, as opposed to a theory-oriented approach, and is directed at demonstrating the validity of constructs and approaches for modeling specific phenomena or solving specific classes of problems. There are, however, several exceptions as follows.

Huberman and associates [20-22,13] have worked on statistical physics-based models of IASs in relation to resource contention and predictive behavior. Rosenschein and associates[23-26] have related game theory to interaction protocols and coalitions. Gasser [27] gives a detailed account of the range of DAI approaches to coordination. Lesser and Decker [12,19,28-30] conducted research directed at the design of coordination mechanisms using their effectiveness on the characteristics of the tasks and the environment. Nagendra Prasad and Lesser [31] use the "facilitates" and "enables" relationships in agents learning coordination strategies. Malone and Crowston have introduced basic concepts of coordination science [1,2,9,17,18,32]. Their work is based on characterizing dependencies and analyzing specific processes for managing them. Durfee and Montgomery reduce coordination to search [33]. Malone [32] discusses hierarchies and markets based on the latter three activities. Carley and Prietula [8] emphasize the increasingly complex nature of IASs because intelligent agents tend to act in parallel and to adapt to the behavior of other agents. Jennings [10,34] has proposed techniques dealing with commitments (pledges to undertake a specific course of action) and conventions (means of monitoring commitments in changing circumstances). Findler and associates [35-39] examined variations in system behavior in terms of variations of the precision of agents' models. This work simulated the development of patterns of behavior and the onset of chaotic regimes. With certain specific assumptions, use was made of interactions between agents and about individual agents' models of the decision procedures of other agents. Coordination was taken the process by which an agent reasons about its local actions and the anticipated actions of others in trying to ensure community acts in a coherent manner. This was shown to be the key to achieving the overall objective of coordination. The *raison d'être* of multi-agent systems is that no individual agent has sufficient competence, resources, or information to solve the entire problem by itself.

Last but not least, we single out the excellent book by Cohen on empirical methods in Artificial Intelligence [40].

ACKNOWLEDGEMENTS

We thank Peter Heck and Vinutha Ram for earlier programming assistance, Prof. Dennis Young for valuable advice on statistical processing, and Dr. Raphael Malyankar for ideas on coordination.

REFERENCES

[1] T. W. Malone, Analogies between human organizations and artificial intelligence systems: Two examples and some reflections. In M. Masuch & M. Warglien, eds., *Artificial Intelligence in Organization and Management Theory*, Ch. 1, 21-40. North-Holland, Amsterdam, 1992.

[2] T. W. Malone & K. Crowston, The interdisciplinary study of coordination. *ACM Computing Surveys*, 26(1):87-119, March 1994.

[3] N. V. Findler, Distributed knowledge-based systems in manufacturing. In G. Rzevski, ed., *Artificial Intelligence in Manufacturing*, 17-33. Springer Verlag, New York, 1989.

[4] N. V. Findler & J. Gao, Dynamic hierarchical control for distributed problem solving (*Data and Knowledge Engineering*, 2:285-301, 1987).

[5] N. V. Findler & Q. Ge, Perceiving and planning before acting - an approach to enhance global network coherence. *Internat. J. of Intelligent Systems*, 4:459-470, 1989.

[6] R. Axelrod, *The Evolution of Cooperation*. Basic Books, NY, 1984.

[7] R. Axelrod & W. Hamilton, The evolution of cooperation. *Science*, 211:1390-1396, 1981.

[8] K. Carley & M. J. Prietula, ACTS Theory: Extending the Model of Bounded Rationality. In K. Carley & M. J. Prietula, eds., *Computational Organizational Theory,* 55-87. Lawrence Erlbaum, Hillsdale, NJ, 1994.

[9] K. Crowston, A taxonomy of organizational dependencies and coordination mechanisms. CCS Working Paper 174, Center for Coordination Science, MIT, 1994.

[10] N. R. Jennings, Commitments and conventions: The foundation of coordination in multi-agent systems. *Knowledge Engineering Rev.*, 81323-250, 1993.

[11] T. W. Sandholm & V. R. Lesser, Coalition formation among computationally bounded agents. *Artificial Intelligence,* 94, 99-137,1997.

[12] K. S. Decker, E. H. Durfee, & V. R. Lesser, Evaluating research in cooperative distributed problem solving. In L. Gasser & M. N. Huhns, eds., *Distributed Artificial Intelligence*, vol. 2, Ch. 19, 485-519. Morgan Kaufmann , San Mateo, CA, 1989.

[13] B. A. Huberman & N. Glance, Diversity and collective action. In H. Haken & A. Mikhailov, eds., *Interdisciplinary Approaches to Nonlinear Systems*. Springer, Berlin/Heidelberg, 1993.

[14] D. Wilkinson, Cohesion and Conflict: Lessons from the Study of Three-Party Interaction. St. Martin's Press, New York, NY, 1976.

[15] A. Rapoport. N-Person Game Theory: *Concepts and Applications*. Uni. of Michigan Press, Ann Arbor, MI, 1970.

[16] E. Rasmusen, *Games and Information: An Introduction to Game Theory*. Blackwell, Oxford, 2nd edition, 1994.

[17] K. Crowston, Modeling coordination in organizations. In M. Masuch & M. Warglien, eds., *Artificial Intelligence in Organization and Management Theory*, Ch. 9, 215-234. North-H2lland, Amsterdam, 1992.

[18] K. Crowston, Evolving novel organizational forms. In K. Carley & M. J. Prietula, eds., *Computational Organizational Theory,* 19-38.. Lawrence Erlbaum, Hillsdale, NJ, 1994.

[19] K. S. Decker, Environment-Centered Analysis and Design of Coordination Mechanisms. PhD dissertation, Uni. of Massachusetts at Amherst, May 1995. Dept. of C.S

[20] T. Hogg & B. A. Huberman, Controlling chaos in distributed systems. *IEEE Trans. on Systems, Man and C ybernetics*, 21(6):1325-1332, 1991.

[21] B. A. Huberman & N. Glance, Evolutionary games and computer simulation. *Proc. of the Nat. Ac. of Sciences (USA)*, 90(16):7715-7718, 1993.

[22] B. A. Huberman & T. Hogg, The behavior of computational ecologies. In B. A. Huberman, ed., *The Ecology of Computation*, 77-115. Elsevier Science B. V., Amsterdam, 1988.

[23] J. S. Rosenschein & M. Genesereth, Deals among rational agents. In B. A. Huberman, ed., *The Ecology of Computation*, 117-132. Elsevier, Amsterdam, 1988.

[24] G. Zlotkin & J. S. Rosenschein, Coalition, cryptography, and stability: Mechanisms for coalition formation in task oriented domains. In *Proc. of the Nat. Conf. on Artificial Intelligence*, 432-437, Washington, DC, 1994.

[25] J. S. Rosenschein & G. Zlotkin, Rules of Encounter: Designing Conventions for Automated Negotiation among Computers. MIT Press, Cambridge, MA, 1994.

[26] G. Zlotkin & J. S. Rosenschein, Negotiation and conflict resolution in non-cooperative domains. In *Proc. of the Eighth Nat. Conf. on Artificial Intelligence*, 100-105, Boston, MA, 1990.

[27] L. Gasser, DAI approaches to coordination. In N. M. Avouris & L. Gasser, eds., *Distributed Artificial Intelligence: Theory and Praxis*, 31-51. Kluwer, Boston, MA, 1992.

[28] K. S. Decker & V. R. Lesser, Quantitative modeling of complex computational task environments. In *Proc. of the 11-th Nat. Conf. on Artificial Intelligence*, 217-224, Washington, DC, 1993.

[29] K. S. Decker, Task environment centered simulation. In M. Prietula, K. Carley & L. Gasser, eds., *Simulating Organizations: Computational Models of Institutions and Groups*, 1996.

[30] K. S. Decker, TfiMS: A framework for analysis and design of coordination mechanisms. In G. O'Hare & N. Jennings, eds., *Foundations of Distributed Artificial Intelligence*. Wiley, 1995.

[31] M. V. Nagandra Prasad & V. Lesser, The use of meta-level information in learning situation-specific coordination. *Proc. of the Internat. Joint Conf. on Artificial Intelligence*, Nagoya, Japan, pp. 640-546, 1997.

[32] T. W. Malone, Modeling coordination in organizations and markets. *Management Science*, 33(10):1317-1332, 1987.

[33] E. H. Durfee & T. A. Montgomery, Coordination as distributed search in a hierarchical behavior space. *IEEE Trans. on Systems, Man and Cybernetics*. SMC-21: 1363-1378, 199].

[34] N. R. Jennings, Controlling cooperative problem solving in industrial multi-agent systems using joint intentions. Artificial Intelligence, 75:195-240, 1993.

[35] N. V. Findler & R. M. Malyankar, An empirical approach to a Theory of Coordination Part I: Design principles and first results (*Computational and Mathematical Organization Theory*, 6**:**119-144, 2000).

[36] N. V. Findler & R. M. Malyankar, Emergent behavior in societies of heterogeneous, interacting agents: Alliances and norms. In R. Conte & N. Gilbert, eds., *Artificial Societies: The Computer Simulation of Social Life*, Ch. 11, 212-237. Uni. College (London) Press, London, 1995.

[37] N. V. Findler & R. M. Malyankar, Social structures and the problem of coordination in intelligent agent societies (Invited talk at the *Special Session on Agent-Based Simulation, Planning and Control in IMACS World Congress 2000*, Lausanne; 2000; Proc. on CD: Paper 122-12)

[38] R. M. Malyankar, N. V. Findler & P. S. Heck, The effects of satisficing models on coordination. In *Satisficing Models: Papers from the 1998 AAAI Spring Symposium*, TR SS-98-05, pp. 39-45, AAAI Press: Palo Alto, CA, 1998.

[39] R. M. Malyankar & N. V. Findler, A methodology for modeling coordination in intelligent agent societies*. Computational and Mathematical Organization Theory,* 4, 317-345, 1998.

[40] P. R. Cohen, *Empirical Methods for Artificial Intelligence*. The MIT Press, Cambridge, MA, 1995.