

Kathleen Carley

Les Gasser

This chapter provides an overview of the Computational Organization Theory Field.

7.1 Introduction

From the hospital, to the schoolroom, to the boardroom people find that the actions they take affect and are affected by various organizations, and the norms, procedures, culture, and members of those organizations. In order to navigate through an organizational world, agents (human and artificial) need social and organizational intelligence. This organizational intelligence comprises many dimensions, including communication capabilities, knowledge about who knows what, knowledge about norms, procedures, and culture of the organization, and more.

The ability of an organization to act is certainly dependent on the intelligence of the agents within it. However, organizations, and multi-agent systems in general, often show an intelligence and a set of capabilities that are distinct from the intelligence and capabilities of the agents within the system. It is not difficult to find multi-agent systems that display non-random and repeated patterns and processes of action, communication, knowledge, and memory (beyond the lifetime of a single agent) regardless of whether or not the agents are human; i.e., multi-agent systems exhibit an internal organizational design. Such designs can structure activities and attention within the system, control the actions of the system as a corporate entity, and may emerge spontaneously or be imposed. From country to country, culture to culture, task to task, and agent type to agent type, we find both differences and commonalities in the designs of organizations, in the patterns and processes connecting individual agents. In order to navigate through this environment and achieve results not achievable by individual agents, or to exhibit capabilities not held by individual agents, organizations and indeed all multi-agent systems, need to act as intelligent information processors, capable of acting as a single corporate entity, and to coordinate individual agents using some type of organizing principle or design. Research in the computational organization area employs computational techniques to theorize about and analyze organizations and the processes of organizing.

The goal of this chapter is to describe what can be done and what others have done in this area: the underlying principles, assumptions, and concerns, and

the major streams of work. After reading this chapter we hope you will have gained insight into the aims, findings and new possibilities of this field. We also hope you will have developed a preliminary understanding of the nature of computational organizational models and your own ideas about how to construct virtual experiments using such models.

7.1.1 What is an Organization?

A classic response to the question “What is an organization?” is “I know it when I see it.” Indeed, every text book in organizational theory provides a definition of organizations. Unfortunately, there is no wide consensus on the definition of “organization,” and indeed as theorists reason about organizations trying to answer fundamentally different questions, they construct different definitions of the basic phenomenon. While there is no single definition of organizations that is uniformly agreed to, there are general tenets that are more or less shared. In general, organizations are characterized as:

- large-scale problem solving technologies
- comprised of multiple agents (human, artificial, or both)
- engaged in one or more tasks; organizations are systems of activity
- goal directed (however, goals can change, may not be articulable, and may not be shared by all organizational members)
- able to affect and be affected by their environment
- having knowledge, culture, memories, history, and capabilities distinct from any single agent
- having legal standing distinct from that of individual agents

One rationale for the existence of organizations qua organizations is that they exist to overcome the limitations of individual agency¹. From this viewpoint, there are four basic limitations: cognitive, physical, temporal, and institutional.

1. Cognitive Limitations—Agents as boundedly rational actors have cognitive limitations and therefore must join together to achieve higher-levels of performance.
2. Physical Limitations—Agents are limited physically, both because of their physiology and because of the resources available to them, and therefore must coordinate their actions, e.g., to achieve higher-levels of productivity. All action takes place situated in specific space-time locations, and agents are limited (e.g. by relativity limits) in their access to other spacetime locations; this fundamental locality means that and distributed action is fundamentally a multiagent—and hence potentially

1. Other rationales include human needs for social affiliation, and the simple non-teleological emergence of patterns of activity in complex environments. However, here we focus on the standard functional approach.

organized—phenomenon.

3. Temporal Limitations—Agents are temporally limited and therefore must join together to achieve goals which transcend the lifetime of any one agent.

4. Institutional Limitations—Agents are legally or politically limited and therefore must attain organizational status to act as a corporate actor rather than as an individual actor.

There is a plethora of ways in which organizations are constituted to overcome limitations of individual agency. Researchers in various areas refer to the way in which an organization is organized as the form, structure, architecture or design of that organization. Decades of research in this area have repeatedly shown that there is no single correct or proper organizational design.² That is, there is no single design that yields the optimal performance under all conditions. Which design is optimal depends on a variety of factors including the specific task or tasks being performed, the intelligence, cognitive capabilities, or training of the agents, the volatility of the environment, legal or political constraints on organizational design, and the type of outcome desired (e.g., efficiency, effectiveness, accuracy, or minimal costs). Researchers' recognition of how organization performance differentially depends upon multiple factors has led to the notion of "contingency theories" of organization. From an organizational engineering perspective, locating an optimal organizational design for a specific, multi-dimensional situation is the aim. From a theoretical perspective locating the general principles and tradeoffs underlying organizational design in a multidimensional space is key.

Consequently, research in this area has often focused on the search for general principles of organizing and the conditions under which these principles do or do not apply. For example, two such linked principles are task or occupational specialization (namely that individuals can become more effective when they are expert in particular activities requiring particular and limited types of knowledge), and the division of labor (namely that appropriate division of tasks, knowledge, and skill among agents in an organization can improve performance e.g., by limiting task and knowledge dependencies). In general, organizations which employ specific and productive instances these principles are able to overcome the limitations of individual agency, coordinate individual actions, and leverage training costs, skill development, and resources in such a way that the organization as a whole achieves higher levels of performance than are otherwise achievable. However, over-specialization and excessive division can reduce performance and flexibility by de-skilling individuals, decreasing attention due to boredom, and increasing decision making time, and by actually increasing coordination costs in situations of uncertainty or failure.

2. This includes field and survey research on actual human organizations, laboratory experiments on human groups, virtual experiments using computational models, and analyses using mathematical models.

7.1.2 What is Computational Organization Theory?

Researchers in the field of Computational Organization Theory (COT) use mathematical and computational methods to study both human and automated organizations as computational entities. Human organizations can be viewed as inherently computational because many of their activities transform information from one form to another, and because organizational activity is frequently information-driven.

Computational Organization Theory (COT) attempts to understand and model two distinct but complementary types of organization. The first is the natural or human organization which continually acquires, manipulates, and produces information (and possibly other material goods) through the joint, interlocked activities of people and automated information technologies. Second, COT studies artificial computational organizations comprising multiple distributed agents which exhibit collective organizational properties (such as the need to act collectively, an assignment of tasks, the distribution of knowledge and ability across agents, and constraints on the connections and communication among agents). Researchers use computational analysis to develop a better understanding of the fundamental principles of organizing multiple information processing agents and the nature of organizations as computational entities. The general aims of research in this area is to build new concepts, theories, and knowledge about organizing and organization in the abstract, to develop tools and procedures for the validation and analysis of computational organizational models, and to reflect these computational abstractions back to actual organizational practice through both tools and knowledge.

Research in this area has resulted in a large number of models, each with its own special characteristics. For example, one useful model of information-seeking, decisionmaking, and problem-solving activity in organizations is distributed search. Since formal computational models of search are well understood, modeling organizational activity as search can provide a clear and tractable explanatory framework. New approaches to control or task allocation in distributed search frameworks can, by analogy, provide suggestive new approaches to these problems in human organizations, e.g. in the development of new organizational forms or for reasoning about the effects of alternative strategic decisions. In the end, distributed search models provide just one type of abstraction that is useful for reasoning about problems of both human organizations and computational ones, and so help to unify thinking about both types.

Some other research projects with particular illustrative models are listed in Table 7.1.

Computational organization theories are most often grounded in existing cognitive, knowledge-based, information-processing theories of individual behavior. However, COT extends this to an organizational level [e.g., 55] and gives precision to the notion of bounded rationality by specifying the nature of the boundaries [7]. The basic tenets of the organizational information-processing perspective developed since the later 1940s are:

Table 7.1 Illustrative Models

Model	Author
Garbage Can	Cohen, March and Olsen (1972)
AAIS	Masuch and LaPotin (1989)
ELM	Carley (1992)
HITOP	Gasser and Majchrzak (1992)
Plural-Soar	Carley et al. (1992)
VDT	Cohen (1992), Levitt et al. (1994)
Organizational Consultant	Baligh, Burton and Obel (1990, 1994)
Action	Gasser and Majchrzak (1994)
Orgahead	Carley and Svoboda (1996)
TAEMS	Decker
Sugarscape	Epstein and Axtell (1996)
Cultural Transmission	Harrison and Carrol (1991)

- **Bounded rationality:** Organizational agents are boundedly rational. There are two types of bounds — limits to capabilities and limits to knowledge. Capabilities depend on the agents' cognitive, computational, and/or physical architecture. Knowledge depends on the agents' ability to learn and the agents' intellectual history. The agents' position in an organization influences to which information an agent has access. Thus, an agents' knowledge of how to do specific tasks, of how its specific organization operates, and indeed of how organizations operate in general, is a function of what positions the agent has held.
- **Information ubiquity:** Within organizations large quantities of information in many different forms are widely distributed across multiple agents. The information may not necessarily be correct.
- **Task orientation:** Organizations and the agents within them are continually engaged in performing one or more tasks. The tasks that in which an organization is engaged require these agents to communicate, build on, analyze, adapt or otherwise process organizational information using various technologies, and to search out new information and new solutions.
- **Distributional constraints:** Organizational performance is a function of what information is shared by whom, when, and of the process of searching for that information. An organizations' culture is the distribution of the knowledge and process across the agents within it. This distribution affects the extent and character of socially shared cognition, team mental models, group information processing, and concurrent information analysis.
- **Uncertainty:** Uncertainty about task outcomes, environmental conditions, and about many other aspects of organizational life influences organizational activity. Distributed computational models such as distributed search or distributed constraint satisfaction pose distribution itself as a source of uncertainty: distribution can render critical uncertainty-reducing information less available because of the

cost of seeking, transmitting, or assimilating it, and because of the overhead of coordinating information needs across agents.

- **Organizational intelligence:** Organizational intelligence resides in the distribution of knowledge, processes, procedures across agents and the linkages among agents. Organizations redesign themselves and their vision of their environments on the basis of the information available to them, with the aim of enabling them to better search for or process information. Such redesign is part of organizational learning processes. It can alter an organization's intelligence, and may or may not improve organizational performance.
- **Irrevocable change:** As agents and organizations learn, their intelligence is irrevocably restructured. This one-directional evolution means that the kind and order in which things are learned—particular histories—can have dramatic consequences.
- **Necessity of Communication:** In order to function as a corporate unit, agents within an organization need to communicate. This communication may take place explicitly by sending and receiving messages or implicitly by perceiving the actions of others.

In addition to this neo-information-processing view of organizations researchers in this area share a series of implicit background assumptions. These are:

- **Modelability:** Organizational phenomena are modelable.
- **Performance differential:** It is possible to distinguish differences in organizational performance.
- **Manipulability:** Organization are entities that can be manipulated and transformed.
- **Designability:** Organizations are entities that can be designed. This is not to say that organizations do not evolve, nor that they cannot be found in nature, for assuredly both events occur. However, they can also be consciously designed and redesigned: organizational transformations can be purposeful and principled.
- **Practicality:** Organizational transformations (based on the design or manipulation of models) can be transferred into and implemented in actual practice.
- **Pragmatism:** The costs of modeling and researching organizations using computational methods are relatively lower than the costs of manipulating or researching similar aspects of actual organizations in vivo, and the benefits gained outweigh the costs.

These assumptions that underlie the research in the computational organization theory area are the result of a fundamentally interdisciplinary intellectual history. Research in this area draws on work in distributed artificial intelligence (DAI), multi-agent systems, adaptive agents, organizational theory, communication theory, social networks, and information diffusion. One of foundational works in this area is *The Behavioral Theory of the Firm* [13] in which a simple information processing model of an organization is used to address issues of design and performance.

While the strongest roots are in the information processing [55,43,59,18,13] and social information processing [54] tradition, current models also have roots in the areas of resource dependency [49], institutionalism [51] population ecology [28], and symbolic interaction [20]. Formalisms and specific measures of organizational design are drawn from the work in the areas of coordination [40] social networks [60] and distributed control [12,15,36].

7.1.3 Why take a Computational Approach?

Organizations are heterogeneous, complex, dynamic nonlinear adaptive and evolving systems. Organizational action results from interactions among adaptive systems (both human and artificial), emergent structuration in response to non-linear processes, and detailed interactions among hundreds of factors. As such, they are poor candidates for analytical models. Because of the natural complexity of the object of study, existing models and theories of organization are often vague, intuitive, and under-specified. Scientific progress will be more readily achievable if the theories are more explicit and well defined. Computational theorizing helps to achieve this.

Computational analysis is an invaluable tool for theory building and examining issues of organizational dynamics as it enables the researcher to generate a set of precise, consistent and complete set of theoretical propositions from basic principles even when there are complex interactions among the relevant factors. Computational models allow researchers to show proofs of concept and to demonstrate whether or not completely modelable factors can generate certain phenomena. In this way, computational models can be used to show the potential legitimacy of various theoretical claims in organization science.

Theoretical computational models can be used to demonstrate lower bounds or tractability of organizational information processing phenomena (e.g., minimal information necessary to reach distributed agreement or awareness [26], or the tractability of an organizational decision or negotiation processes [52]. Experimental and empirically-based models can also provide computationally-plausible accounts of organizational activity [29,14].

7.2 Organizational Concepts Useful in Modeling Organizations

In order to model an organization the following factors must be modeled at some level of detail: the agents, the organization's design, the task or environment, the technology, and the stressors. Organizations can use different configurations of agents, designs, tasks, and technology to accomplish the same goal. In fact, one of the major issues in the computational organization area is determining what configuration makes sense when and the relative costs and benefits of various configurations.

Models in the COT area vary dramatically in the level of detail in which the

agent, design, task, and technology are modeled. The better or more detailed these underlying model the more precise the predictions possible from the model. Models run the gamut from simple abstract models of generic decision making behavior (such as the Garbage Can Model and CORP) to detailed models of specific organizational decisions or decision making processes (such as VDT and Hi-TOP). The simpler more abstract models are typically referred to as intellectual models. For these models a central research goal is theory building: to discover general principles underlying organizational behavior. The more detailed models may allow the researcher to use the model to emulate specific organizations by entering specific authority structures and/or procedures. For these models a key research goal is organizational engineering: to examine whether or not the performance of a specific organization will be affected by making some specific change such as re-engineering the task in a particular way or adding a new technology.

7.2.1 Agent and Agency

Organizations are composed of agents. These agents may be human, artificial, or both. Agents take action. Agents can make decisions. The actions that the agents are capable of, and the decisions that they make, depend on their capabilities and knowledge, the situation in which they are embedded, and the task(s) they are performing. Agents are boundedly rational. Each agent occupies a position in the organization. This position defines what task(s) the agent does, whom the agent must communicate with, whom the agent reports to, who reports to the agent, and so forth. Agents have specific knowledge, skills, and capabilities. Classes of agents can be defined on the basis of differences in position, knowledge, skills, or capabilities. Importantly, the agent's knowledge is potentially comprised not just of task-based or technical knowledge but also of social or organizational knowledge.

In most organizational models agents are viewed as cooperating together to achieve some collective goal such as producing more widgets, finding the best path through a maze, filling orders, or classifying objects. However, agents need not be cooperative. Competition among agents may emerge for a variety of reasons, one of the most common being an organizational accounting system that reimburses individuals on the basis of individual contribution or performance and not just the overall organization performance. Similarly, in most organizational models agents are viewed as essentially honest and as not knowingly communicating incorrect information or decisions. However, this need not be the case. Ironically, for many organizational tasks, the task itself may be so ambiguous, that the agent may not be able to detect whether or not other agents are lying. Or, the task may be so constraining that lying is immediately obvious.

Models of organizations have represented organizations as single decision makers and as collections of decision makers. Within the multi-agent organizational models agents have been modeled in a variety of ways. For example, in the Garbage Can Model agents are characterized by a small set of parameters and vectors such as their energy, what problems they consider salient, what problems they can act on,

and so forth. In this case the agent's abilities are represented as numerical arrays or values. For example, what problems are salient to an agent is represented by which cells in an agent by problem matrix are ones, and agent energy is simply a numeric value. In VDT the agent is modeled as an in-box, an out-box, a set of preferences for how to handle information, a set of skills, and so forth. Whereas, in Plural-Soar each agent is modeled as a separate soar agent. In this case the agent's knowledge is a set of rules in a series of problem spaces.

7.2.2 Organizational Design

The organization's design is the procedures and rules that embody organizational knowledge and the set of connections linking agents and tasks. The procedures and rules range from task-based rules, to accounting procedures, to procedures for hiring, firing, promoting, and moving agents about. The set of linkages among agents and tasks are often described as structures or networks. There are multiple such structures in the organization. The structure most familiar to anyone who has worked in an organization is the authority structure. This is often called the formal organization chart. However, the authority structure is simply the formal structure. In addition, there are a myriad of other inter-linked structures that constrain and provide opportunities for action, and which are typically not perceived by those who work in the organization. One such additional structure is the informal structure — the network of interactions among agents such as the friendship network, and the advice network. In addition to these structures is the task structure (the precedence ordering among subtasks), the task-resource structure (defining which resources are needed for which task), the task-skill structure (defining what skills are needed for which task), the resource access structure (defining which agent has access to which resources), the skill structure (defining which agent has which skills), the task assignment structure (defining which agent is assigned or is allowed to work on which subtasks), and so on.

The two most typical ways of conceptualizing the organization's design is as a set of attributes (such as centralized or not centralized, or density of communication ties) or as a set of matrices. The attribute approach is used in Hi-TOP, the Organizational Consultant, and AAIS. The matrix approach is used in the Garbage Can Model, CORP, and VDT. Illustrative structures in both their matrix and network form are shown in Figure 7.1.

7.2.3 Task

The organization and its members are engaged in one or more tasks. These tasks may be composed of subtasks which may themselves be further subdivided. These tasks may have dependencies among them. Thompson identified three such dependencies—pooled (the results from two or more tasks are jointly needed to perform a different task), sequential (two or more subtasks must be performed in a specified sequence), and reciprocal (two tasks depend jointly on each other).

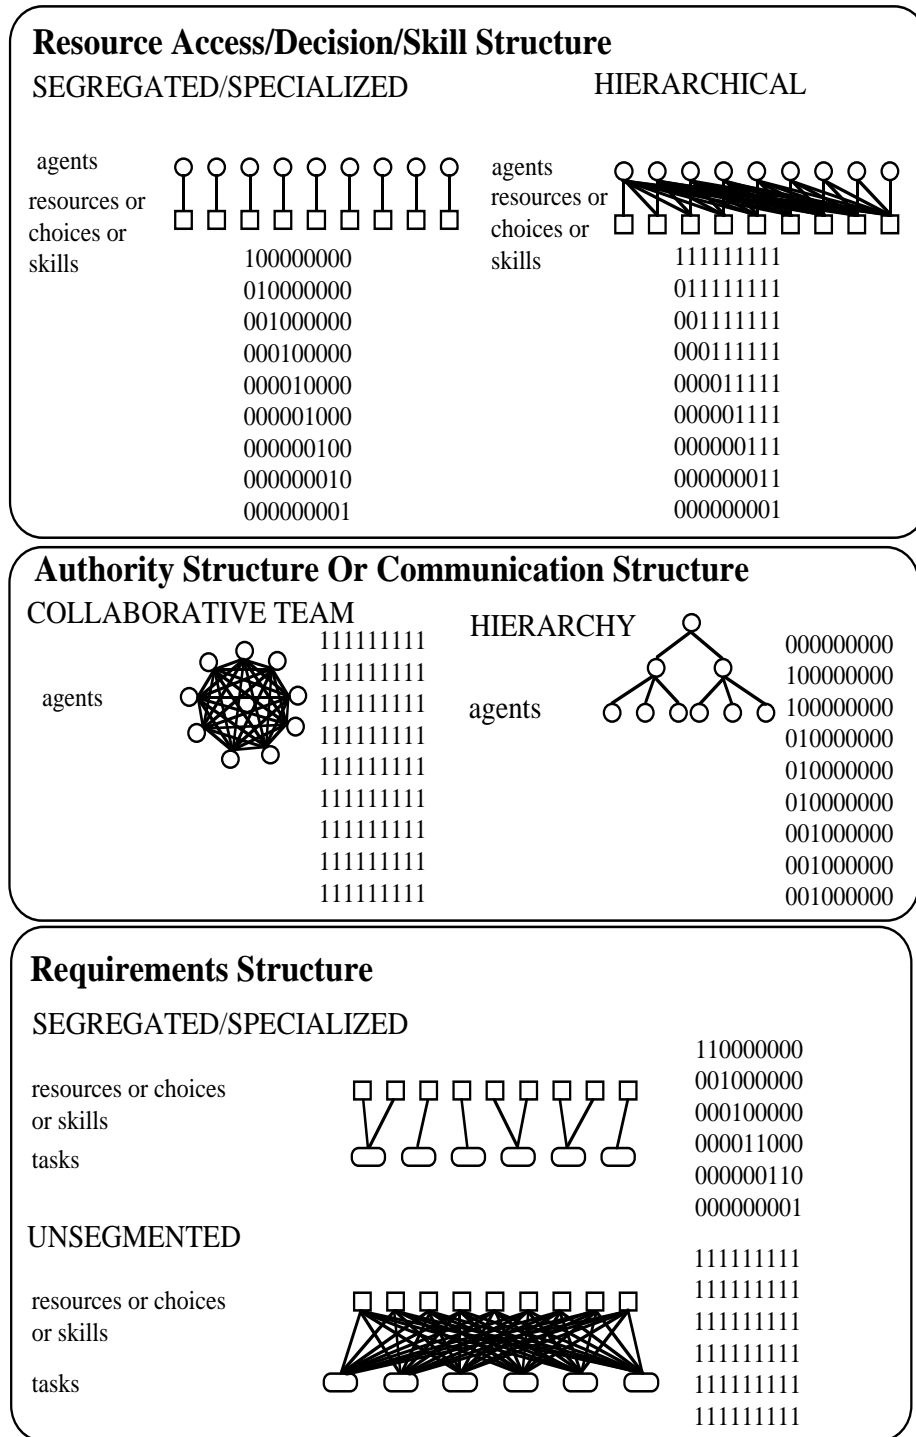


Figure 7.1 Illustrative structures that comprise an organization's design.

Organizations are expected to improve their performance if they match the rest of their structure to the underlying task structure.

One way of characterizing the organization's environment is as the set of tasks the organization faces. For example, in the car industry can be viewed as a repeated design task. Task environments vary in the degree of repetition, volatility, bias, and complexity. Repetition: whether the set of tasks are repetitive (you do the same thing over and over), quasi-repetitive (you do the same type of thing repeatedly but the specific instances or details are different), or non-repetitive (each task is unique). Volatility: the rapidity with which the task environment changes. Bias: the extent to which all possible tasks, regardless of task features, have the same outcome or solution. Complexity: the amount of information that needs to be processed to do the task. Typical task environment changes are oscillating (e.g., seasonal firms such as swimming pool dealers often have an alternate task, such as selling Christmas goods, that they do in the alternate season), step change (e.g., when a new manufacturing technology is introduced firms change from one task to another), and gradual (e.g., as minor variations are made in non-production technologies).

The performance of the organization can be measured with respect to the task or tasks it is performing. There are three types of performance measures: effectiveness (is the task being performed well), efficiency (is the task being performed in such a way that output is maximized relative to some input), and perceived effectiveness (is the organization perceived as performing well by one or more stakeholders such as the general public, the government, the board of directors, or the media). For many tasks in which the product is generated by the group as a whole, while it might be possible to measure the overall organization's performance, in real human groups it is often impossible to objectively measure the actual contribution of any one member. With respect to effectiveness, three aspects of effectiveness are: relative performance (how well is the organization performing compared to other organizations), accuracy (how many decisions are being made correctly), and timeliness (how rapidly are decisions being made). For particular tasks or industries there are often entire literatures on how specifically to measure performance in that situation.

Within the COT area there are two strands of research on tasks. Some models, such as VDT, use very detailed models of specific organizational tasks focusing on the dependencies among subtasks but leaving the content of what is done in the subtask otherwise unspecified. In contrast, other models such as CORP use highly stylized tasks that, although retaining key features of actual tasks, differ in detail and complexity from those done in actual organizations. These highly stylized tasks are often referred to as canonical tasks. A set of such tasks are emerging. This set includes: the sugar-production task, the maze, the binary classification task, the radar task, the warehouse task, and the meeting scheduling task. In Figure 7.2 the binary choice task and the warehouse task are illustrated.

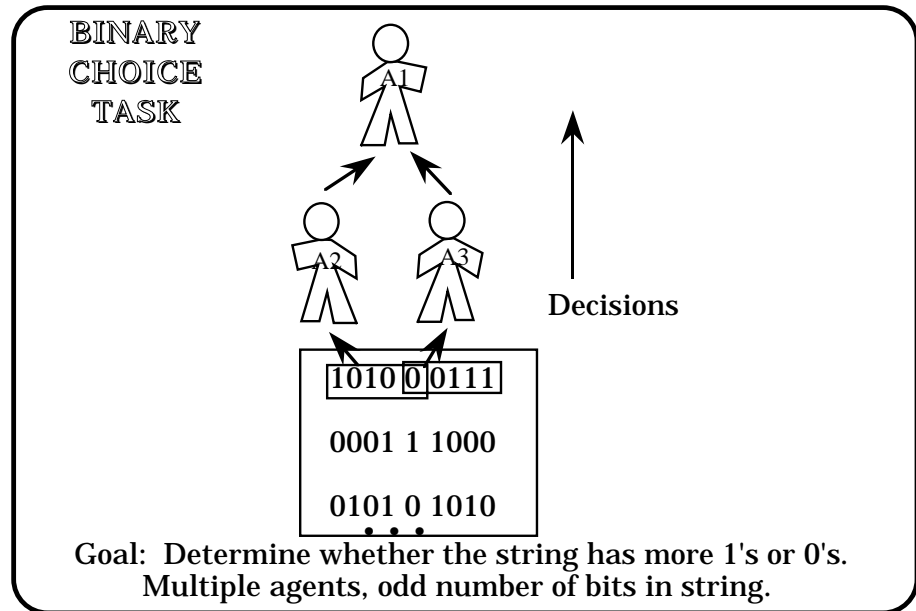
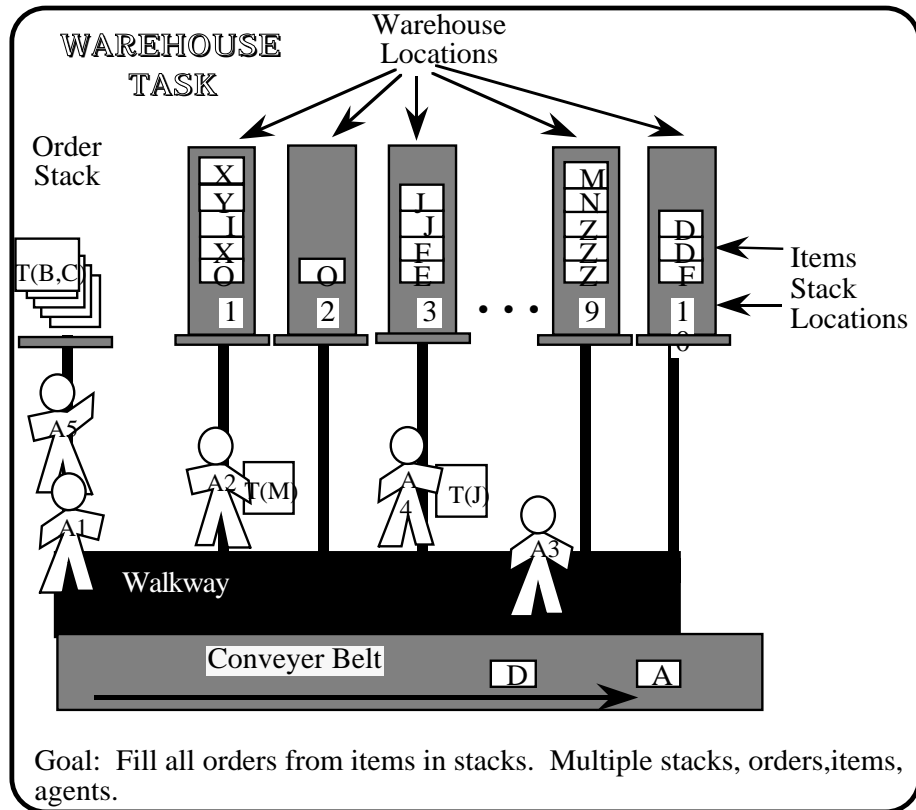


Figure 7.2 Illustrative tasks.

7.2.4 Technology

Research on organizational behavior needs to account for the role of information and telecommunication technologies. Most COT models ignore issues of technology within the organization. In a few cases, researchers have begun to look at how different technologies for processing, storing, retrieving, or communicating information might affect the actions of the individuals within the organization and/or the organization's overall performance. Two different approaches to examining the impact of communication technologies have appeared: technology as tool and technology as agent.

One approach to modeling technology is to treat it as a tool and to differentiate tools in terms of the attributes such as access, speed, synchronicity, and recordability. This approach is taken in the Virtual Design Team (VDT) [32,37]. Within VDT, the organizational agents have a suite of communication technologies available to them like telephone, face-to-face meetings, and email. The agent also has a preference to send certain information via certain technologies. Using VDT the researcher can begin to examine how altering the way in which memos are sent, e.g., by paper or email, may affect the speed with which the organization solves a task and the amount of rework that needs to be done.

A second approach to modeling technology is to treat it as an artificial agent (and as altering the information processing capabilities of existing human agents who have access to the technology). This approach is taken in the constructural model [33]. Within the constructural model the agents have an information processing or communication technology available to them which affects how many other agents they can simultaneously communicate with, whether the communicated message must be the same to all receivers, how much information the agent can retain, and so on. The technology itself may act as an agent with particular information processing and communicative capabilities. For example, a book as agent can interact with many other agents at the same time, send a different message to each receiver, can survive longer in the population than human agents, and, unlike human agents, cannot learn.

7.3 Dynamics

Many of the key issues in the COT area center around organizational dynamics. How do organizations learn? How do organizations evolve? What new organizational designs will emerge in the next decade or century? To examine such issues many researchers use complex adaptive approaches or optimization approaches, such as genetic algorithms, genetic programming, neural networks, and simulated annealing. This work demonstrates that interactions between agent architecture, the way agents are coordinated, and the way the agents and the coordination structure adapt and change over time affect organizational behavior. Another common result is that for collections of even minimally intelligent agents organization sponta-

neously emerges. Indeed, many studies show that hierarchy spontaneously emerges.

Organizations are seen to be potentially dynamic in many ways. Organizations are capable of being redesigned (change in who reports to whom) or re-engineered (change in the who does what). Agents within the organization are capable of changing; e.g., by learning or, in the case of artificial agents, by reconfiguring themselves. The processes, communications, or types of interactions can change. And so forth. There are a variety of processes that affect the organization's ability to adapt. For example, in order to achieve new levels of performance organizations often need to engage in an exploration process where they examine new technologies, new ways of doing business, new designs, and so on. However, organizations can also improve performance by exploiting current knowledge and technologies and getting better at what they do best. Exploration and exploitation are somewhat at odds with each other and organizations need to balance these two forces for change [41]. Organizational adaptation is not guaranteed to improve performance. Organizations are typically more likely to redesign themselves when they are new and such adaptations may in turn cause the organization to fail. Such early failure is referred to as the liability of newness.

A variety of different approaches to organizational dynamics have been taken. To illustrate some of the issues involved two different approaches to will be briefly described. The first of these examples is based on the idea of flexible agents - agents which can restructure themselves in response to changes in the environment. The second of these examples is based on the idea of dual-level learning — organizations in which agent level and structural level learning occur in response to changes in the environment.

Ishida, Gasser and Yokoo [30,31; see also 23] demonstrated the potential for flexible organizational structures to emerge from flexible agents. The basic idea underlying their approach is that the agents are flexible entities and the interactions are the stable foundation. Agents were modeled as a set of problem-solving production rules; i.e. mappings of rules to rule-collections. These mappings were flexible. Interactions were modeled as production rules; i.e. under certain metalevel control (rule-firing) conditions, this knowledge (LHS clause or input) interacts with this other knowledge (another LHS clause or input) to produce this knowledge (a RHS clause, which maps into an LHS clause for another rule or an output). These production rules never changed. Agents flexibly emerged out of a fabric of interactions, and re-configured their local knowledge and the structure or pattern of interactions among themselves in response to changes in the environment. The actual number and character of agents changed over time as did the organizational structure. The simulations showed that really effective organizations tended to learn adaptations over time; i.e., on repeated exposure to similar (oscillating) environmental changes these organizations went through fewer adaptations. In this sense the organizations of agents learned how to learn.

Carley and Svoboda [9] used the ORGAHEAD model of organizational change to demonstrate the importance of learning histories and that organizations in which both the agents and the structure were flexible and could learn over time were

not guaranteed to improve their performance. The ORGAHEAD model is based on the social conception of organizations as inherently complex, computational and adaptive in which knowledge and learning occurs at multiple levels. Within ORGAHEAD organizational action results from both the behavior of multiple agents and the structure (or network) connecting these agents and the way in which knowledge is distributed among these agents. Agents learn through gaining new knowledge through experience. This is implemented using a standard experiential learning model in which agents add new knowledge to old knowledge and continually update the probability with which they take certain actions. Learning occurs at the structural level—by altering procedures and linkages among the agents (such as who reports to whom and who does what)—as the organization redesigns and restructures itself. This strategic learning is implemented as a simulated annealing algorithm. In this case there is a CEO (the annealer) that proposes a change, evaluates the potential impact of this change on the organization by trying to anticipate what will happen in the future, and then decides whether or not to accept the change according to the Metropolis criteria. According to the Metropolis criteria the change is always accepted if it is anticipated to improve performance and is accepted but with decreasing likelihood over time if the change is risky and is anticipated to decrease performance. The results from these studies show that not all adaptation is advantageous. Individual and structural learning clash; e.g., organizations re-engineer themselves for better performance only to lose the lessons of experience learned by various agents as those new agents are moved to different tasks or leave the organization. Because of such learning clashes adaptation is often maladaptation, the history of how and when the organization adapted is as much a determinant of performance as the organization's design, and truly adaptive organizations whose performance actually improves over time are those which balance the two types of learning.

7.4 Methodological Issues

There are numerous methodological issues involved in the development and testing of computational models of organizations and organizing. There are three, however, that require special mention. The first, is the use of virtual experiments to examine the implications of the computational model. The second, has to do with validation, verification and the relation of the computational models to data on organizations. The third, has to do with the role of development tools and frameworks to mitigate the costs of developing these models.

7.4.1 Virtual Experiments and Data Collection

Unlike early models in this area, such as the Garbage Can Model, today's computational models of organizations are often sufficiently complex that they cannot be completely analyzed. For example, the parameter space of set of options is fre-

quently so extensive that the researcher cannot explore all possible input combinations to determine the performance of the system. Nevertheless, a large number of combinations need to be explored as the performance of the system may change drastically for different combinations of inputs. One of the reasons that modern models are so complex is that organizations themselves are complex. Another, is that models are often designed by representing process. As such, the same model can be used to address a number of questions about organizations and organizing.

To address this problem, researchers in this area run virtual experiments. A virtual experiment is an experiment in which the results are gathered via simulation. In running a virtual experiment the researcher sets up a series of simulations to address a specific question. For the virtual experiment the researcher chooses a small set of parameters—perhaps three—and then varies these systematically over some range. All other parameters are typically held constant or allowed to vary randomly in classic Monte Carlo fashion. Statistical procedures for designing and analyzing the resulting data can be used for virtual experiments just as they can for experiments using humans in the laboratory.

For example, imagine that the computational organizational model allows the researcher to control the number of agents, the way agents make decisions (following standard operating procedures or based on experience, how the agents can send messages (such as face-to-face one-on-one or group meetings, email to one other agent or to all other agents), the complexity of the task, the complexity of the organization's authority and communication structure, and a variety of other parameters or options. Such a model could be used to address a number of research questions including: (1) How large does the organization need to be to reap the benefits of email? And (2) for the same task are there different combinations of technology, authority structure, and communication structure that lead to the same level of performance? To address the first question the researcher might vary the size of the organization from say 2 to 200 in increments of 20 (11 cells) and may consider all four communication technologies. This would be a 11x4 experimental design. To address the second question the researcher might consider all four communication technologies, two different authority structures (e.g., team and hierarchy), and two different communication structures (e.g., a completely connected structure like everyone-to-everyone and one that follows the authority structure (only communication is to or from manager). This would be a 4x2x2 design. In each case some number of simulations would be needed to be run for each cell, with the number chosen based on the required power of the test.

7.4.2 Validation and Verification

Computational organization theory is grounded theory. That is the models that embody the theory are informed by and tested against empirical data. This grounding is done using various validation and verification procedures. In the COT area three types of validation are particularly important: theoretical, external, and cross-model. Theoretical verification has to do with determining whether the model is

an adequate conceptualization of the real world for assessing the key issue being addressed. The adequacy of the conceptualization is often determined on the basis of whether or not a set of situation experts consider the model to have captured the main factors that they observe in organizations. External validation has to do with determining whether or not the results from the virtual experiments match the results from the real world. Finally, cross-model validation has to do with determining whether or not the results from one computational model map on to, and/or extend, the results of another model.³

For both theoretical and external validation the real world may be a human organization, a laboratory experiment, or an organization of artificial agents, and so on. Organizations leave “traces” of their activities such as accounting records, stockholder reports, technical connections among parts, operating procedures, web pages, etc. These can be analyzed using computational methods. Such data can also be captured, mapped, analyzed, and linked to other computational models either as input or as data against which to validate the computational models. Such data helps to form and test computational theories of organization and organizing.

7.4.3 Computational Frameworks

One of the pressing issues in the COT area is the development of a general testbed or framework that has the appropriate building blocks to minimize the time required to develop organizational and social models. A variety of tools are beginning to appear; as yet, however, no one tool dominates. Among the existing tools are: MACE, SDML, Multi-agent Soar, and SWARM.

7.4.3.1 *M*

ACE

MACE [19,20] was one of the first general (domain-independent) testbeds for modeling multi-agent systems. It is a truly concurrent distributed object system. MACE introduced the idea of using agents for all phases of system construction, user interaction, and experiment management, as well as for the basis of the modeled system itself. Thus, the testbed and the experiment became an integrated multi-agent organization. For example, “user interface agents” are used as actual asynchronous user-to-system and system-to-user representatives, interpreters, translators, displays, managers, and so forth.

MACE also included explicit social modeling concepts drawn from social theory. The first such idea is the recursive composition of agents so that a group can itself be treated as an agent with distributed internal structure. Secondly, the “social worlds” concept of symbolic interactionists like Anselm Strauss is operationalized as knowledge-based agent boundaries. Each agent defines a set of acquaintances.

3. Cross-model validation is also called docking [2].

This acquaintanceship knowledge, rather than explicit constraints or laws or testbed programming structures, defines the boundaries of communication and interaction (hence the social structure). This concept provides a clean semantic model for flexible agent organizations. Finally, MACE used “modeling other agents” as its foundation of social knowledge and social structure: drawing on the ideas of G.H. Mead and the symbolic interactionists (Mead: ‘taking the role of the other’ as a foundational principle that can unify mind, self, and society over time), and also based on Hewitt’s ideas of Actor acquaintances (a much simpler notion, basically just message addresses). MACE included specific facilities to model a number of features of other agents (like goals, roles, skills, etc.) in special a acquaintance database, and it used these to structure individual interactions and thus to establish social structure defined as patterns of interaction over time. This idea of modeling others and acquaintances has now become commonplace within CS/AI; however, few researchers recognize the link they are making to social theory.

7.4.3.2 SDML

SDML (Strictly Declarative Modeling Language) [47,48,16] is a multi-agent object-oriented language for modeling organizations. SDML is particularly suited for modeling multi-agent systems in which the agents interact in a team (flat) or hierarchical organizational structure. SDML is effectively theory-neutral with respect to the cognitive capabilities of the agent. It is flexible enough to represent both simple agents and more sophisticated agents as well as the linkages among them. SDML currently includes various libraries for alternate architectures such as genetic programming and Soar. These libraries facilitate exploring the interaction between agent cognition and organizational design.

Key social ideas are captured in the SDML architecture. For example, social agents are capable of distinguishing between explanation and action. The declarative representation within SDML makes this possible. Within SDML agents in the same class can be represented by sharing rules between them. Another key idea in organizations is that within the organization there are predefined linkages among agents and predefined roles in which knowledge is embedded and that constrain behavior. From this perspective, structure is patterns of positions or roles over time. This notion of structure is integral to SDML as within SDML the structure of the multi-agent system is represented as a container hierarchy. For example, agents may be contained within divisions which are contained within organizations. Containers and their associated agents are also linked by an inheritance hierarchy. Change in agents and in the linkages among them is made possible by controlling the time levels associated with agent and container data bases.

7.4.3.3 Multi-Agent Soar

Soar is a computational architecture for general intelligence [34]. Agents are goal directed and can be characterized in terms of their problem spaces, operators,

and associated preferences. Preferences can be used to represent shared norms or cultural choices. The agent's goals need not be articulable and can be automatically generated or consciously selected by the agent as deliberation ensues and. Multi-agent Soar is an approach to modeling agents as collections of Soar agents [57,58,6]. The current system facilitates inter-agent communication and does not require each agent to be a separate processor.

Soar was designed as a specification of key psychological ideas such as bounded rationality. As such, Soar can be thought of as a unified theory of cognition. Multi-agent soar is built around two core social ideas: modeling other agents and communication. In multi-agent Soar models each agent has a mental model of each other agent and the social world. This mental model may include expectations about the other agents' goals, preferences, and so forth and allows the agent to anticipate what it thinks others will do. Communication in these models is handled by passing *communique's* with content. Agents have problem spaces for determining when to communicate what to whom, how to compose and parse messages. Agent's can be interrupted by sending them new communications and can learn from the communication.

7.4.3.4 *SWARM*

SWARM is a multi-agent simulation language for modeling collections of concurrently interacting agents in a dynamic environment [56,46,1]. Both synchronous and asynchronous execution is admitted. SWARM is particularly suited to exploring complex systems composed of large numbers of relatively simple agents. Within SWARM computations can dynamically restructure themselves to accommodate changes in the input data and the objective function. One of the intended applications of SWARM is artificial life applications. A large variety of agent based models can be implemented within SWARM.

The key social idea that is captured in SWARM is the logic of collective intelligence. That is, over time systems of SWARM agents come to exhibit collective intelligence over and above the simple aggregation of agent knowledge. This notion of emergent intelligence is central to the science of complexity.

7.5 Conclusion

Computational organization theory (COT) is the study of organizations as computational entities. As noted, the computational organization is seen as taking two complementary forms: [1] the natural or human organization which is replete with information and the need to process it and [2] computational systems composed of multiple distributed agents which have organizational properties. Computational analysis is used to develop a better understanding of organizing and organizations.

Organization is seen to arise from the need to overcome the various limitations on individual agency—cognitive, physical, temporal, and institutional. Organizations,

however, are complex entities in which one or more agents are engaged in one or more tasks and where knowledge, capabilities and semantics are distributed. Thus, each organization has a design, a set of networks and procedures linking agents, tasks, resources, and skills that describes these various distributions.

Computational organizational models are grounded operational theories. In other words, unlike traditional DAI or multi-agent models COT models draw on and have integrated into them empirical knowledge from organization science about how human organizations operate and about basic principles for organizing. Much of this work follows in the information processing tradition. Many of the COT models are models composed of other embedded models. In these multi-level models, the traditional distinction between normative and descriptive often becomes blurred. For example, the models may be descriptive at the individual level—describing individuals as boundedly rational, with various built in cognitive biases—but normative at the structural level—finding the best organizational design subject to a set of task based or procedural constraints.

Computational analysis is not simply in service to organizational theorizing; rather, computational organizational theorizing is actually pushing the envelope in terms of computational tools and techniques. COT makes contributions to mainstream AI and CS, including fostering progress on such issues as: large scale qualitative simulation, comparison and extension of optimization procedures (particularly procedures suited to extremely complex and possibly changing performance surfaces); aggregation/disaggregation of distributed objects; on-line/offline coordination algorithms; organizational and multiagent learning; and semantics heterogeneity. Research in this area requires further development of the scientific infrastructure including developing: easy-to-use cost-effective computational tool kits for designing and building computational models of organizations, teams, and social systems (e.g., a multi-agent oriented language with built in task objects and communication); multi-agent logics; intelligent tools for analyzing computational models; validation procedures, protocols, and canonical data sets; managerial decision aids based on computational organization models; and protocols and standards for inter-agent communication. Key theoretical concerns in this area center around determining: what coordination structures are best for what types of agents and tasks; whether hybrid models (such as a joint annealer and genetic programming model) are better models for exploring organizational issues and for locating new organizational designs; representations for, and management of, uncertainty in organizational systems; the interactions among, and the relative advantages and disadvantages of various types of adaptation, evolution, learning, and flexibility; measures of organizational design; the existence of, or limitations of, fundamental principles of organizing; the tradeoffs for system performance of task-based, agent-based, and structure-based coordination schemes; representations for information and communication technology in computational models; and the relation between distributed semantics and knowledge on teamwork, organizational culture and performance.

Three directions that are particularly important for future research are organiza-

tional design, organizational dynamics and organizational cognition. The key issue under organizational design is not what is the appropriate division of labor, nor is it how should agents be coordinated. Rather, there is a growing understanding that there is a complex interaction among task, agent cognition or capabilities, and the other structural and procedural elements of organizational design. As such, the issue is finding what combinations of types of agents, structures (patterns of interactions among agents), and ways of organizing the task are most likely to meet the organization's goal. The key issue under organizational dynamics is not whether or not organizations adapt. Rather, the issues center on how to encourage effective learning, how to change to improve performance, how to retain capabilities and knowledge as the organization changes to address changes in the environment, and what new designs are possible. As to organizational cognition (perception, memory) there are a variety of issues ranging from how to represent organizational knowledge, to what level of sharing (of knowledge, procedures, or semantics) is necessary and by which agents to ensure effective organizational performance.

7.6 Exercises

1. *[Level 1]* Provide a critical discussion of the following statement. You do not need to know organizational theory to create good models of organizations. Anyone who has ever worked in an organization can develop such models.
2. *[Level 1]* Provide a critical discussion of the following statement. How does the organizational design and the task bound the agent? What are typical constraints and opportunities afforded the agent by the design and task? Provide at least five examples for both design and task.
3. *[Level 1]* For an organization that you are familiar what types of agents exist in that organization, what are their limitations.
4. *[Level 2]* Develop a simple model of a small group of agents (1,2 or 3) trying to collectively solve a simple canonical task such as the binary choice task or the maze task. What additional issues are involved, and what extra features does the model need, as you move from 1 to 2 to 3 agents working together to do the task? How is performance affected by the increase in the number of agents?
5. *[Level 3]* Reimplement and extend in one or more ways the garbage can model of organizational choice [11]. There are many possible extensions, some of which have been discussed in the literature. Possible extensions include, but are not limited to the following: adding a formal organization authority structure, having agents work on multiple tasks simultaneously, altering the task so that it requires specific skills and not just energy to be completed, and allowing agent turnover. Show that your model can replicate the original results reported by Cohen, March and Olsen (i.e., dock the models [2]). Then show which results are altered, or what additional results are possible, given your extension.
6. *[Level 3]* Reimplement and extend in one or more ways the CORP model of

organizational performance [38]. There are many possible extensions, some of which have been discussed in the literature. Possible extensions include, but are not limited to the following: adding an informal communication structure, having agents work on multiple tasks simultaneously, allowing agents to be promoted, altering incoming information so that it is potentially incomplete or erroneous, altering the nature of the feedback (e.g., by delaying it or making it more ambiguous), and making the agents competitive (e.g., make agents try to maximize the performance relative to other's performance). Show that your model can replicate the original results reported by Lin and Carley (i.e., dock the models). Then show which results are altered, or what additional results are possible, given your extension.

7. [Level 3] For a small organization (5 to 30 people) develop a description of its design. What is the formal organization chart? What is the informal advice network (who goes to whom for work related advice)? What are the main tasks and subtasks being accomplished? Develop a task dependency graph matrix. What are the skills or resources needed to do those tasks? Develop a resource/skill access matrix and a resource/skill requirements matrix. What were the major difficulties you encountered in locating this information for an actual organization?

8. [Level 4] Develop a comprehensive representation scheme for task or a multi-agent language for doing task based models of organizations. Consider how task is represented in various organizational models. What are the limitations or features of the various representation schemes? What features should be built into your approach? Demonstrate the strength of your approach by reimplementing one or more existing COT models.

9. [Level 4] Develop a general purpose approach for modeling telecommunication and information processing technologies in the organization. What are the critical features or components of these technologies that must be modeled? How does your approach contrast with the technology as agent approach and the technology as feature approach? What are the limitations and advantages of your approach?

References

- . Axelrod. Advancing the art of simulation in the social sciences. Working Paper 97-05-048, Santa Fe Institute, 1997.
- . Axtell, R. Axelrod, J. M. Epstein, and M. D. Cohen. Aligning simulation models: A case study and results. *Computational and Mathematical Organization Theory*, 1(2): 123–142, 1996.
- . H. Baligh, R. M. Burton and B. Obel. Devising expert systems in organization theory: The organizational consultant. In M. Masuch, editor, *Organization, Management, and Expert Systems*, pages 35–57. Walter De Gruyter, Berlin, 1990.
- . H. Baligh, R. M. Burton and B. Obel. Validating the organizational consultant on the fly. In K. M. Carley and M. J. Prietula, editors, *Computational Organization Theory*, pages 179–194. Lawrence Erlbaum Associates, Hillsdale, NJ, 1994.
- . M. Carley. Organizational learning and personnel turnover. *Organization Science*, 3(1): 2–46, 1992.

- . M. Carley, J. Kjaer-Hansen, M. Prietula and A. Newell. Plural-Soar: A prolegomenon to Artificial Agents and Organizational Behavior. In M. Masuch and M. Warglien, editors, *Distributed Intelligence: Applications in Human Organizations*, pages 87-118. Elsevier Science Publications, Amsterdam, The Netherlands, 1992.
- . M. Carley and A. Newell. The nature of the social agent. *Journal of Mathematical Sociology*, 19(4): 221-262, 1994.
- . M. Carley and M. J. Prietula, editors. *Computational Organization Theory*. Lawrence Erlbaum Associates. Hillsdale, NJ, 1994.
- . M. Carley and D. M. Svoboda, Modeling organizational adaptation as a simulated annealing process. *Sociological Methods and Research*, 25(1): 138-168, 1996.
- . P. Cohen. *The Virtual Design Team: An Information Processing Model of Coordination in Project Design Teams*. PhD thesis, Stanford University, Department of Civil Engineering, Stanford, CA, 1992.
- . D. Cohen, J. G. March and J. P. Olsen. A garbage can model of organizational choice, *Administrative Sciences Quarterly* 17(1): 1-25, 1972.
- . D. Corkill. *A Framework for organizational self-design in distributed problem solving networks*. PhD Thesis, University of Massachusetts, Dept. of Computer and Information Science, Amherst, MA, 1982.
- . Cyert and J. G. March. *A Behavioral Theory of the Firm*. 2nd Edition. Blackwell Publishers, Cambridge, MA, 1992[1963].
- ecker, 1996.
- . H. Durfee, V. R. Lesser and D. D. Corkill. Coherent cooperations among communicating problem solvers. *IEEE Transactions on Computers*, C-36: 1275-1291, 1987.
- . Edmonds, M. Scott and S. Wallis. Logic, reasoning and a programming language for simulating economic and business processes with artificially intelligent agents. In P. Ein-Dor, editor, *Artificial Intelligence in Economics and Management*, pages 221-230. Kluwer Academic Publishers, Boston, MA, 1996.
- . M. Epstein and R. Axtell. *Growing Artificial Societies*. MIT Press, Cambridge, MA.
- . R. Galbraith, *Designing Complex Organizations*, Addison-Wesley Publishing Company, 1973.
- . Gasser, C. Braganza and N. Herman. MACE: A flexible testbed for distributed AI research. In M. N. Huhns, editor, *Distributed Artificial Intelligence*, pages 119-152. Pitman Publishers, 1987a.
- . Gasser, C. Braganza and N. Herman, Implementing distributed AI systems using MACE. In *Proceedings of the 3rd IEEE Conference on Artificial Intelligence Applications*, pages 315-320. Orlando, FL, February, 1987b. (Also in Readings in DAI).
- . Gasser and A. Majchrzak. ACTION Integrates manufacturing strategy, design, and planning. In P. Kidd and W. Karwowski, editors, *Ergonomics of Hybrid Automated Systems IV*, IOS Press, Netherlands, 1994.
- . Gasser and A. Majchrzak. HITOP-A: Coordination, infrastructure, and enterprise integration. In *Proceedings of the First International Conference on Enterprise Integration*, Hilton Head, South Carolina: MIT Press, Cambridge, MA, 1992.
- . Gasser and T. Ishida. A dynamic organizational architecture for adaptive problem solving, In *Proceedings of the National Conference on AI*, pages 52-58. July, 1991.

- . S. Glance and B. A. Huberman. The outbreak of cooperation. *Journal of Mathematical Sociology*, 17(4): 281–302, 1993.
- . S. Glance and B. Huberman. Social dilemmas and fluid organizations. In K. M. Carley and M. J. Prietula, editors, *Computational Organization Theory*, pages 217–240. Lawrence Erlbaum Associates, Hillsdale, NJ, 1994.
- alpern, 1996.
- . R. Harrison and G. R. Carrol. Keeping the faith: A model of cultural transmission in formal organizations. *Administrative Science Quarterly*, 36: 552–582, 1991.
- . T. Hannan. and J. Freeman. *Organizational Ecology*. Harvard University Press, Cambridge, MA, 1989.
- uberman, XXXX.
- . Ishida, L. Gasser and M. Yokoo. Organization self-design of distributed production systems. *IEEE Transactions on Data and Knowledge Engineering*, 4(2): 123–134, 1992a.
- . Ishida, L. Gasser and M. Yokoo. An organizational approach to real-time continuous problem solving. *Japanese Society of Artificial Intelligence Journal*, 7(2): 300–308, 1992b (In Japanese).
- . Jin and R. Levitt. The Virtual Design Team: A computational model of project organizations. *Computational and Mathematical Organization Theory*, 2(3): 171–196.
- . S. Kaufer and K. M. Carley. *Communication at a Distance: The Effect of Print on Socio-Cultural Organization and Change*. Lawrence Erlbaum Associates, Hillsdale, NJ, 1993.
- . E. Laird, A. Newell, and P. S. Rosenbloom. Soar: An architecture for general intelligence. *Artificial Intelligence*, 33(1): 1-64, 1987.
- . K. Lant and S. J. Mezias. Managing discontinuous change: A simulation study of organizational learning and entrepreneurial strategies, *Strategic Management Journal*, 11: 147–179, 1990.
- . R. Lesser and D. D. Corkill. Functionally accurate cooperative distributed systems. In A. H. Bond and L. Gasser, editors, *Readings in Distributed Artificial Intelligence*, Morgan Kaufmann, Inc., San Mateo, CA, 1988.
- . E. Levitt, G. P. Cohen, J. C. Kunz, C. I. Nass, T. Christiansen and Y. Jin. The 'Virtual Design TeamU: simulating how organization structure and information processing tools affect team performance. In K. M. Carley and M. J. Prietula, editors, *Computational Organization Theory*, pages 1–18. Lawrence Erlbaum Associates, Hillsdale, NJ, 1994.
- . Lin, and K.M. Carley. Organizational response: The cost performance tradeoff. *Management Science*, 43(2): 217–234, 1997.
- . Majchrzak and L. Gasser, HITOP-A: A tool to facilitate interdisciplinary manufacturing systems design," *International Journal of Human Factors in Manufacturing*, 2(3): 255–276, 1992.
- . W. Malone. Modeling coordination in organizations and markets. *Management Science*, 33: 1317–1332, 1986.
- . G. March. Exploration and exploitation in organizational learning. In M. D. Cohen and L. S. Sproull, editors, *Organizational Learning— / Sage, Thousand Oaks, CA, 1996.*
- . G. March and R. Weissinger-Baylon, editors. Ambiguity and Command:

- Organizational Perspectives on Military Decision Making. *Pitman, Boston, MA, 1996.*
- . *G. March and H. A. Simon. Organizations. Wiley, 1958.*
- . *Masuch and M. Warglien. Artificial Intelligence in Organization and Management Theory Elsevier Science Publishers, Amsterdam, The Netherlands, 1992.*
- . *Masuch and P. LaPotin (1989). Beyond Garbage Cans: An AI Model of Organizational Choice. Administrative Science Quarterly, 34: 38–67 .*
- . *Minar, R. Burkhart, C. Langton, M. Askenazi. The Swarm simulation system: A toolkit for building multi-agent simulations. Working Paper 96-06-042. Santa Fe Institute. 1996.*
- . *Moss and O. Kuznetsova. Modelling the process of market emergence . In J. W. Owsinski and Z. Nahorski, editors, Modelling and analysing economies in transition, pages 125–138. MODEST, Warsaw. 1996.*
- . *Moss and B. Edmonds. A formal preference-state model with qualitative market judgements. Omega—the international journal of management science, 25(2): 155–169, 1997.*
- . *Pfeffer and G. R. Salancik. The external control of organizations : a resource dependence perspective. Harper Row, New York, 1978.*
- . *F. Padgett. Managing garbage can hierarchies. Administrative Science Quarterly, 25(4): 583-604, 1980.*
- . *W. Powell and P. J. DiMaggio. The New Institutionalism in Organizational Analysis. The University of Chicago Press, Chicago, IL, 1991.*
- osenschein and Zlotkin, 1994.*
- . *R. Salancik and H. Leblebici. Variety and form in organizing transactions: A generative grammar of organization. Research in the Sociology of Organizations, 6:1–31, 1988.*
- . *R. Salancik and J. Pfeffer. A social information processing approach to job attitudes and task design. Administrative Science Quarterly, 23:224–253. 1978.*
- . *A. Simon, Administrative Behavior. Free Press, New York, 1947.*
- . *Stites. Complexity. Omni, 16(8):42–52. 1994.*
- . *Tambe. Teamwork in real-world, dynamic environments. In Proceedings of the International Conference on Multi-agent Systems, 1996.*
- . *Tambe. Agent Architectures for flexible, practical teamwork. In Proceedings of the AAAI American Association of Artificial Intelligence, 1997.*
- . *Thompson. Organizations in Action. McGraw-Hill, New York, NY, 1967.*
- . *Wasserman and K. Faust. Social Network Analysis: Methods and Applications. Cambridge University Press, Cambridge, NY, 1994.*

Keywords

- computational organization theory
- network
- design
- structure
- adaptation
- linkages

Glossary of Key Terms

- computational organization theory
- network
- design
- structure
- adaptation
- linkages

Annotated List of Key Systems and Tools

- Garbage Can
- AAIS
- CORP
- HITOP-A
- Plural-Soar
- VDT
- Organizational Consultant
- Action
- Orgahead
- TAEMS
- Sugarscape
- Cultural Transmission
- MACE
- SDML
- Multi-agent Soar
- SWARM