

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Task Environment Centered Design of Organization

Permalink

<https://escholarship.org/uc/item/46b0p0bx>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 19(0)

Author

Decker, Keith

Publication Date

1997

Peer reviewed

Task Environment Centered Design of Organizations

Keith Decker (DECKER@CIS.UDEL.EDU)
Dept. of Computer and Information Sciences
University of Delaware, Newark, DE 19716 USA

Abstract

The central idea, which is not new for those who study human organizations but which is sometimes forgotten by computer agent researchers, is that the design of coordination mechanisms cannot rely on the principled construction of the agents alone, but must rely on the structure and other characteristics of the task environment. Such dependencies include the structure of the environment (the particular kinds and patterns of interrelationships that occur between tasks) and the uncertainty in the environment (both in the *a priori* structure of any episode within an environment and in the outcomes of an agent's actions). In this talk, I will briefly describe a modeling framework, TÆMS, for representing abstract task environments. TÆMS has been used both for environment modeling/simulation *and* as an internal representation for computer agents to plan, schedule, and coordinate their activities with other agents (human or computer). I'll describe examples of both of these uses. This written summary provides a background bibliography, and pointers to the work discussed in the talk.

Introduction

The design of organizations or other coordination mechanisms for groups of agents depends in many ways on the agent's task environment. Just two of these dependencies are on the structure of the tasks and on the uncertainty in the task structures. The task structure includes the scope of the problems facing the agents, the complexity of the choices facing the agents, and the the particular kinds and patterns of interrelationships that occur between tasks. A few examples of environmental uncertainty include uncertainty in the *a priori* structure of any particular problem-solving episode, in the actions of other agents, and in the outcomes of an agent's own actions. These dependencies hold regardless of whether the system comprises just people, just computational agents, or a mixture of the two. For example, the presence of both uncertainty and high variance in a task structure can lead a system of agents to perform better by using coordination algorithms that adapt dynamically to each problem-solving episode (Decker & Lesser, 1993). Designing organizational coordination mechanisms also depends on non-task characteristics of the environment such as communication costs, and properties of the agents themselves. Representing and reasoning about the task environment must be part of any computational theory of coordination.

TÆMS (Task Analysis, Environment Modeling, and Simulation) was developed as a framework with which to model

and simulate complex, computationally intensive task environments at multiple levels of abstraction and from multiple viewpoints. It is a tool for building and testing computational theories of coordination. TÆMS is compatible with both formal computational agent-centered approaches and experimental approaches. The framework allows us to both mathematically analyze (when possible) and quantitatively simulate the behavior of multi-agent systems with respect to interesting characteristics of the computational task environments of which they are part. We believe that it provides the correct level of abstraction for meaningfully evaluating centralized, parallel, and distributed control algorithms, negotiation strategies, and organizational designs.

The use of TÆMS to model external environments has led to its use by computer agents as an *internal*, subjective model of the external environment. As part of an agent's internal representation of its environment, a TÆMS model allows an agent to reason about how multiple, interacting decision criteria change in response to actual and possible local and non-local agent actions. The most significant internal-agent work using TÆMS—which I will not discuss—has been the development of an agent local activity scheduler (Garvey & Lesser, 1993; Garvey, Humphrey, & Lesser, 1993; Wagner, Garvey, & Lesser, 1997).

Background

Artificial Intelligence, growing as it has from the goal of modeling *individual* intelligence, or at least replicating or augmenting it, has focused primarily on representations of individual choice and action. A large effort has gone into describing the principled construction of agents that act rationally and predictably based on their beliefs, desires, intentions, and goals (Wooldridge & Jennings, 1995). Fairly recently, researchers concerned with real-world performance have also realized that Simon's criticisms and suggestions about economics (March & Simon, 1958) also hold for many realistically situated individual agents—perfect rationality is not possible with bounded computation (e.g., (Boddy & Dean, 1989)). Distributed AI has too often kept the individualistic character of its roots, and focused on the principled construction of individual agents. It hasn't even, so far, really concerned itself with the questions of bounded rationality in real-time problem solving when it comes to the principled

construction of individual agents¹. Worst of all, it has failed yet to bring the environment to center stage in building and analyzing distributed problem solving systems.

In contrast, the organizational science community has since the 60's (e.g. (Lawrence & Lorsch, 1967)) regarded the task environment as a crucial, central variable in explaining complex systems, and a whole branch of research has grown up around it (contingency theory). Representations in this community are rarely mathematically formal in nature but rather try to present very rich descriptions using terms such as uncertainty, decomposability, stability, etc.

TÆMS, as a framework to represent coordination problems in a formal, domain-independent way, is unlike any existing computational representation that is focussed on coordination issues. The form of the framework is more detailed in structure than many organizational-theoretic models of organizational environments, such as Thompson's notions of pooled, sequential, and reciprocal processes (1967), Burton and Obel's linear programs (1984), or Malone's queueing models (1987), but is influenced by them, and by the importance of environmental uncertainty and dependency that appear in contingency-theoretic and open systems views of organizations (Lawrence & Lorsch, 1967; Galbraith, 1977; Stinchcombe, 1990; Scott, 1987). As a problem representation for computational tasks, it is richer and more expressive than game theory representations (Rosenschein & Zlotkin, 1994). For example, a typical game or team theory problem statement is concerned with a single decision; a typical TÆMS objective problem solving episode represents the possible outcomes of many sequences of choices that are interrelated with one another (e.g., "schedules"). TÆMS can represent a game theoretic problem, and we could boil down a single decision made by an agent faced with a TÆMS task structure into a game theoretic problem.² Because TÆMS is more expressive, we can use it to operationalize some of the rich but informal concepts of organizational science, especially those that focus on various dependencies and uncertainties that are the basis of (for example) both contingency theoretic and transaction cost economic (Williamson, 1975; Moe, 1984) views of organizations. An example of this is our recreation of Burton and Obel's experiments on *decomposability* (Burton & Obel, 1984; Decker, 1997).

As a tool for building and testing computational theories of coordination, the TÆMS framework can, for example, support the construction of ACTS theory instances (Carley & Pritula, 1994). In ACTS theory organizations are viewed as collections of intelligent agents who are cognitively restricted, task oriented, and socially situated. TÆMS provides ways to think about and represent environmental constraints (task characteristics and social characteristics involving communication links and what information and what possible actions

¹On the other hand, work on (mostly standalone) robotic agents has wrestled with these questions, e.g., (Simmons et al., 1997)

²TÆMS does not say how agents make their decisions. It is perfectly reasonable for an (computer) agent to use game-theoretic reasoning processes.

are available to what agents). While simple models can some times be solved analytically (Decker & Lesser, 1993), many complex models require simulation techniques. Compared to other organizational simulations such as (Lin & Carley, 1993) or (Levitt et al., 1994), TÆMS provides a much more detailed model of task structures, and does not provide any fixed *agent* model.

The contingency theory observation that no single organization or coordination mechanism is 'the best' across environments, problem-solving instances, or even particular situations is also common in the study of multi-agent cooperative distributed problem solving (Fox, 1981; Durfee, Lesser, & Corkill, 1987; Durfee & Montgomery, 1991; Decker & Lesser, 1993). Key features of task environments demonstrated in both these threads of work that lead to different coordination mechanisms include those related to the structure of the environment (what we will call *task interrelationships*) and environmental *uncertainty*.

Short Overview of TÆMS

The principle purpose of a TÆMS model is to analyze, explain, or predict the performance of a system or some component. While TÆMS does not establish a particular performance criteria, it focuses on providing multi-criteria performance information such as the temporal intervals of task executions, and the *quality* of the execution or its result. *Quality* is an intentionally vaguely-defined term that must be instantiated for a particular environment and performance criteria—there could be a whole vector of result/state attributes, over which an agent would presumably express its preferences. Examples of *quality* vector attributes include the precision, belief, or completeness of a task result. TÆMS models describe how several quantities—the quality vector produced by executing a task, the time taken to perform that task, the time when a task can be started, its deadline, and whether the task is necessary at all—are affected by the execution of other tasks.

A TÆMS model of environmental and task characteristics has three levels: *generative*, *objective*, and *subjective*. The *generative* level describes the statistical characteristics of objective problem instances (called *episodes*) in a domain. A generative level model consists of a description of the generative processes or distributions from which the range of alternative problem instances can be derived, and is used to study performance over a range of problems in an environment. The *objective* level describes the essential, 'real' task structure of a particular problem-solving situation or instance over time. Typically no agent ever has access to this complete and total information in the model or simulation. Finally, the *subjective* level describes the agents' view of the situation. A subjective level model is essential for evaluating coordination algorithms, because while individual behavior and system performance can be measured objectively, agents must make decisions with only subjective information.³ Obviously, when

³In organizational theoretic terms, subjective *perception* can be

TÆMS is used as an internal agent representation language in the real, non-simulated world, only the *subjective* information actually exists!

A problem instance (called an *episode*) is defined as a set of task groups, each with a deadline. The task groups may arrive at different times. A task group consists of a set of tasks related to one another by a subtask relationship that forms an acyclic graph (usually a tree). The circles higher up in the tree represent various subtasks involved in the task group, and indicate precisely how quality will accrue depending on what leaf tasks are executed and when. Tasks at the leaves of the tree (without subtasks) represent *basic actions* or *executable methods*, which are the actual computations or actions the agent will actually execute (in the figure, these are shown as boxes). The arrows between tasks and/or methods indicate other task interrelationships where the execution of some method will have a positive or negative effect on the quality or duration components of another method. This notation and associated semantics are formally defined in (Decker, 1995).

Hospital Patient Scheduling Example

Let's look at a brief example of a TÆMS task structure model in terms of its ability to reason about organizational decision making. The following description is from an actual case study (Ow, Prietula, & Hsu, 1989):

Patients in General Hospital reside in units that are organized by branches of medicine, such as orthopedics or neurosurgery. Each day, physicians request certain tests and/or therapy to be performed as a part of the diagnosis and treatment of a patient. [...] Tests are performed by separate, independent, and distally located ancillary departments in the hospital. The radiology department, for example, provides X-ray services and may receive requests from a number of different units in the hospital.

Furthermore, each test may interact with other tests in relationships such as enables, requires—delay (must be performed after), and inhibits (test A's performance invalidates test B's result if A is performed during specified time period relative to B). Note that the unit secretaries (as scheduling agents) try to minimize the patients' stays in the hospital, while the ancillary secretaries (as scheduling agents) try to maximize equipment use (throughput) and minimize setup times.

Figure 1 shows an subjective TÆMS task structure corresponding to an episode in this domain, and the subjective views of the unit and ancillary scheduling agents after four tests have been ordered. Note that quite a bit of detail can be captured in just the 'computational' aspects of the environment—in this case, the tasks use peoples' time, not a computer's. However, TÆMS can model in more detail the

used to predict agent actions or *outputs*, while unperceived, objective environmental characteristics affect performance (or *outcomes*) (Scott, 1987).

physical resources and job shop characteristics of the ancillaries if necessary (Decker, 1995). Such detail is not necessary for us to analyze the protocols developed by (Ow et al., 1989), who propose a primary unit-ancillary protocol and a secondary ancillary-ancillary protocol.

We use *min* (AND) to represent quality accrual because in general neither the nursing units nor ancillaries can change the doctor's orders—all tests must be done as prescribed. We have added two new non-local effects: *requires—delay* and *inhibits*. The first effect says that a certain amount δ of time must pass after executing one method before the second is enabled. The second relationship, *A inhibits B*, means that *B* will not produce any quality if executed in a certain window of time relative to the execution of *A*, and can be defined in a similar manner.

A Summary of TÆMS-related work

Analysis in TÆMS: The methodology we have been building uses the TÆMS framework and other DAI formalisms to build and chain together statistical models of coordination behavior that focus on the sources of uncertainty or variance in the environment and agents, and their effect on the (potentially multi-criteria) performance of the agents. For example, we have used this methodology to develop expressions for the expected value of, and confidence intervals on, the time of termination of a set of agents in any arbitrary simple distributed sensor network environment that has a static organizational structure and coordination algorithm (Decker & Lesser, 1993). We have also used this model to analyze a dynamic, one-shot reorganization algorithm (and have shown when the extra overhead is worthwhile versus the static algorithm) (Decker, 1995). In each case we can predict the effects of adding more agents, changing the relative cost of communication and computation, and changing how the agents are organized (in this case, by changing the range and overlap of their capabilities). These results were achieved by direct mathematical analysis of the model and verified through simulation in TÆMS.

Agent Internal Architectures: We have used TÆMS as a core element in the design of computational agent architectures (DECAF (Decker et al., 1995) and RETSINA (Sycara et al., 1996; Decker et al., 1997a)). These computer agents may work only with other computer agents, or with people (see the examples below). A complete agent comprises seven parts (some of which can be omitted in certain task environments). The central component is the belief knowledge base, which stores a representation of the current agent objectives, the structure of the proposed tasks to achieve those objectives, and annotations on these tasks such as local and non-local commitments, using TÆMS or a TÆMS-like language. The other components are an agent communication component (using KQML, a standard agent communication language), a decision-making component (that uses decision theory to choose what ob-

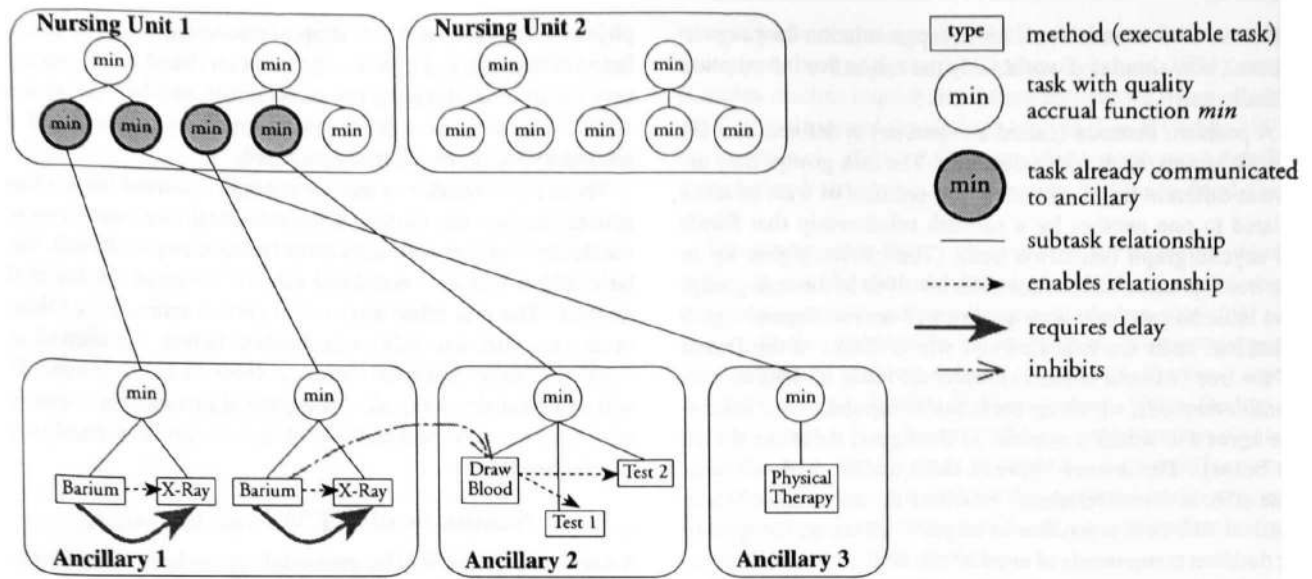


Figure 1: High-level, subjective task structure for a typical hospital patient scheduling episode. The top task in each ancillary is really the same objective entity as the unit task it is linked to in the diagram.

jectives to achieve), a planning component (that builds or retrieves task structures to achieve objectives), a local scheduling component (that orders and locates in time basic executable tasks), an execution monitoring component (that can check task executions for progress and/or missed deadlines), and a coordination component named GPGP (that helps agents to coordinate their actions by communicating task structures and commitments to certain tasks).

Generalized Partial Global Planning: The important thing to remember is that no one agent has a global picture of what every other agent could do/is doing. The complete task structure is broken up across potentially many agents. Some might have the big picture, but no details; others might have all the details but only for some small portion of the problem. The key observation is that whenever a coordination relationship (subtask, enables, facilitates, etc.) extends between the part of a task structure known by one agent and that known by another agent, that there exists an opportunity (or perhaps a requirement) for coordination of the activities of the two agents. The GPGP family of coordination mechanisms includes agent communication behaviors to tease out these spanning coordination relationships, and provides individual mechanisms to react and respond to each possible coordination relationship. Different environments will require different subsets of coordination mechanisms for good performance. Several of our papers have dealt with specifying and analyzing these coordination mechanisms, and learning the situations and environments in which they are useful.

The GPGP algorithm family specifies three basic areas of the agent's coordination behavior: how and when to communicate and construct non-local, partial task struc-

ture views of the current problem solving situation; how and when to exchange the partial results of domain problem solving; how and when to make and break *commitments* to other agents about what results will be available and when. We have experimented with 7 different coordination mechanisms, including a set meant to emulate the original PGP algorithm (Decker & Lesser, 1995), discovery of the need for "successor-side" coordination mechanisms in the distributed data processing domain (Prasad, Decker, Garvey, & Lesser, 1996), and ongoing analysis and evaluation of mechanisms for the hospital patient scheduling problem presented earlier.

WARREN: A multi-agent financial portfolio management organization (Sycara et al., 1996; Decker et al., 1997a). Many computational agents work together in a dynamically organized team in order to efficiently and robustly retrieve changing stock prices, news, and fundamental data from various locations on the Internet, and to analyze that data in various contexts in order to provide an up-to-date financial picture. The WARREN system is *open*, meaning agents can come and go at any time. A WARREN organization consists of a portfolio interface agent for each user, two task agents for fundamental stock analysis and price-news graphing, a news information agent for Dow-Jones and Clarinet news, several different stock ticker information agents, and two EDGAR information agents assigned to the SEC's electronic archives for quarterly and annual reports. Organizational information agents include a "matchmaker" or yellow-pages agent that helps an information requestor find the appropriate information server in the dynamic, open system. Other organizational agents include "brokers" or middle-managers that can help bal-

ance workloads in some subsystems. One important requirement was robustness, so that when any WARREN agent leaves the system (or crashes) the remaining agents reorganize so as to carry on as effectively as possible. We have analyzed and experimentally verified some simple models of the performance of these simple alternate organizational forms with respect to characteristics such as efficiency, adaptability, robustness, and privacy/security (Decker, Sycara, & Williamson, 1997b).

MADEsmart: A project for coordinating mixed human- and computational-agent systems in concurrent engineering design (Obrst et al., 1997). The initial domain problem is the design of helicopter body panels using composite materials. MADEsmart seeks to partially automate the integrated product teams used to organize design engineers through the use of multi-agent approaches. For example, associated with each human engineer in an integrated product team is a user assistant agent that can interact with that engineer. Other agents wrap around existing computationally intensive resources such as composite fiber placement simulations and the COSTADE design cost analysis tool, which uses an existing FORTRAN-based model.

We plan to eventually apply our scheduling technologies to intelligent user interfaces. The user assistant agent will help a user to schedule his or her activities at the workstation and display that schedule in a meaningful and expressive form that can be queried and explained. Of course, users will have significant freedom in the ordering of their activities—the purpose of the Local Schedule Display is to make sure that tasks are not forgotten, that time critical or critical enabling tasks are identified to the user, and that facilitating or other soft-related tasks are also identified.

Future Directions

TÆMS is a framework for describing complex task environments. When combined with traditional DAI tools for describing coordination algorithms and agents, it provides a basis for analysis and/or simulation in any standard Common Lisp environment. When analyzing an existing or proposed organizational design or coordination algorithm, we advocate an approach that focuses first on behavior due to coordination relationships in certain situations and then expanding this model to incorporate the sources of uncertainty that are present. Such a process may iteratively refine the organization or algorithm—with a parameterized algorithm one might approach this as a pure parameter optimization problem. On the other hand, when trying to design an organization or coordination algorithm for a given environment, one can also start with both the coordination relationships and the uncertainties present, and add features to deal with each explicitly (our implementation of the GPGP algorithm, which features several independent ‘plug-in’ modules to deal with different classes of interrelationships, is a case in point).

TÆMS is also a language that can be used internally by agents to represent and reason about their subjective

views of the current problem-solving situation. Significant work has been done on scheduling agent activities efficiently when agent meta-knowledge and preferences are dynamically changing during the problem-solving process. However, much work remains to be done in linking traditional (or non-traditional) planning work to such architectures and dynamic environments. Finally, we need to build, model, analyze, and thus better understand larger, more complex multi-agent organizations—going beyond simple dynamically made teams and simple brokered hierarchies. This work will progress both in the study of totally artificial agent organizations (e.g. WARREN) and mixed human- and computational-agent organizations (e.g. MADEsmart).

Many of these papers are available from the author's web site, <http://www.cis.udel.edu/~decker/>, which also contains pointers to the UMass and CMU agents web sites.

References

- Boddy, M. & Dean, T. (1989). Solving time-dependent planning problems. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, pp. 979–984 Detroit.
- Burton, R. M. & Obel, B. (1984). *Designing Efficient Organizations: Modelling and Experimentation*. Amsterdam: North Holland.
- Carley, K. & Prietula, M. (1994). ACTS theory: Extending the model of bounded rationality. In K. Carley & M. Prietula (Eds.), *Computational Organization Theory*, chap. 4, pp. 55–89. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Decker, K. S. (1995). *Environment Centered Analysis and Design of Coordination Mechanisms*. Ph.D. thesis, University of Massachusetts. <http://dis.cs.umass.edu/~decker/thesis.html>.
- Decker, K. S. (1997). Task environment centered simulation. In M. Prietula, K. Carley, & L. Gasser (Eds.), *Simulating Organizations: Computational Models of Institutions and Groups*. AAAI Press/MIT Press. Forthcoming.
- Decker, K. S. & Lesser, V. R. (1993). An approach to analyzing the need for meta-level communication. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, pp. 360–366 Chambéry, France.
- Decker, K. S. & Lesser, V. R. (1995). Designing a family of coordination algorithms. In *Proceedings of the First International Conference on Multi-Agent Systems*, pp. 73–80 San Francisco. AAAI Press. Longer version available as UMass CS-TR 94–14.
- Decker, K. S., Lesser, V., Prasad, M. N., & Wagner, T. (1995). MACRON: an architecture for multi-agent cooperative information gathering. In *Proceedings of the CIKM-*

- 95 Workshop on Intelligent Information Agents Baltimore, MD.
- Decker, K. S., Pannu, A., Sycara, K., & Williamson, M. (1997a). Designing behaviors for information agents. In *Proceedings of the 1st Intl. Conf. on Autonomous Agents*, pp. 404–413 Marina del Rey.
- Decker, K. S., Sycara, K., & Williamson, M. (1997b). Middle-agents for the internet. In *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence* Nagoya, Japan.
- Durfee, E. H. & Montgomery, T. A. (1991). Coordination as distributed search in a hierarchical behavior space. *IEEE Transactions on Systems, Man, and Cybernetics*, 21(6), 1363–1378.
- Durfee, E. H., Lesser, V. R., & Corkill, D. D. (1987). Coherent cooperation among communicating problem solvers. *IEEE Transactions on Computers*, 36(11), 1275–1291.
- Fox, M. S. (1981). An organizational view of distributed systems. *IEEE Transactions on Systems, Man, and Cybernetics*, 11(1), 70–80.
- Galbraith, J. (1977). *Organizational Design*. Reading, MA: Addison-Wesley.
- Garvey, A., Humphrey, M., & Lesser, V. (1993). Task interdependencies in design-to-time real-time scheduling. In *Proceedings of the Eleventh National Conference on Artificial Intelligence*, pp. 580–585 Washington.
- Garvey, A. & Lesser, V. (1993). Design-to-time real-time scheduling. *IEEE Transactions on Systems, Man, and Cybernetics*, 23(6), 1491–1502.
- Lawrence, P. & Lorsch, J. (1967). *Organization and Environment*. Cambridge, MA: Harvard University Press.
- Levitt, R., Cohen, P., Kunz, J., Nass, C., Christiansen, T., & Jin, Y. (1994). The virtual design team: Simulating how organizational structure and communication tools affect team performance. In K. Carley & M. Prietula (Eds.), *Computational Organization Theory*. Lawrence Erlbaum Associates.
- Lin, Z. & Carley, K. (1993). Proactive or reactive: An analysis of the effect of agent style on organizational decision-making performance. *International Journal of Intelligent Systems in Accounting, Finance, and Management*, 2(4), 271–289. Special issue on “Mathematical and Computational Models of Organizations: Models and Characteristics of Agent Behavior”.
- Malone, T. W. (1987). Modeling coordination in organizations and markets. *Management Science*, 33, 1317–1332.
- March, J. G. & Simon, H. A. (1958). *Organizations*. New York: Wiley.
- Moe, T. M. (1984). The new economics of organization. *American Journal of Political Science*, 28(4), 739–777.
- Obrst, L., Woytowicz, M., Rock, D., Lander, S., Gallagher, K., & Decker, K. S. (1997). Agent-based integrated project teams. Submitted, 1997 ASME Design Engineering and Computers in Engineering Conference; Engineering Information Management Symposium.
- Ow, P. S., Prietula, M. J., & Hsu, W. (1989). Configuring knowledge-based systems to organizational structures: Issues and examples in multiple agent support. In L. F. Pau, J. Motiwalla, Y. H. Pao, & H. H. Teh (Eds.), *Expert Systems in Economics, Banking, and Management*, pp. 309–318. Amsterdam: North-Holland.
- Prasad, M. N., Decker, K. S., Garvey, A., & Lesser, V. (1996). Exploring organizational designs with TÆMS: A case study of distributed data processing. In *Proceedings of the Second International Conference on Multi-agent Systems* Kyoto, Japan.
- Rosenschein, J. S. & Zlotkin, G. (1994). *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*. Cambridge, Mass.: MIT Press.
- Scott, W. R. (1987). *Organizations: Rational, Natural, and Open Systems*. Englewood Cliffs, NJ: Prentice-Hall, Inc.
- Simmons, R., Goodwin, R., Haigh, K., Koenig, S., & O’Sullivan, J. (1997). A layered architecture for office delivery robots. In *Proceedings of the 1st Intl. Conf. on Autonomous Agents*, pp. 245–252 Marina del Rey.
- Stinchcombe, A. L. (1990). *Information and Organizations*. Berkeley, CA: University of California Press.
- Sycara, K., Decker, K. S., Pannu, A., Williamson, M., & Zeng, D. (1996). Distributed intelligent agents. *IEEE Expert*, 11(6), 36–46.
- Thompson, J. D. (1967). *Organizations in Action*. New York: McGraw-Hill.
- Wagner, T., Garvey, A., & Lesser, V. (1997). Complex goal criteria and its application in design-to-criteria scheduling. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence* Providence.
- Williamson, O. E. (1975). *Markets and Hierarchies: Analysis and Antitrust Implications*. New York: The Free Press.
- Wooldridge, M. & Jennings, N. (1995). Intelligent agents: Theory and practice. *The Knowledge Engineering Review*, 10(2), 115–152.