

What is Coordination Theory and How Can It Help Design Cooperative Work Systems?

Thomas W. Malone and Kevin Crowston

Center for Coordination Science (E53-333)
Massachusetts Institute of Technology
Cambridge, MA 02139

It is possible to design cooperative work tools based only on “common sense” and good intuitions. But the history of technology is replete with examples of good theories greatly aiding the development of useful technology. Where, then, might we look for theories to help us design computer-supported cooperative work tools? In this paper, we will describe one possible perspective—the interdisciplinary study of coordination—that focuses, in part, on how people work together now and how they might do so differently with new information technologies.

In one sense, there is little that is new about the study of coordination. Many different disciplines—including computer science, sociology, political science, management science, systems theory, economics, linguistics, and psychology—have all dealt, in one way or another, with fundamental questions about coordination. Furthermore, several previous writers have suggested that theories about coordination are likely to be important for designing cooperative work tools (e.g., [Holt88], [Wino86]).

We hope to suggest here, however, that the potential for fruitful interdisciplinary connections concerning coordination is much greater than has as yet been widely appreciated. For instance, we believe that fundamentally similar coordination phenomena arise—unrecognized as such—in many of the fields listed above. Though a single coherent body of theory about coordination does not yet exist, many different disciplines could both contribute to and benefit from more general theories of coordination. Of particular interest to researchers in the field of computer-supported cooperative work is the prospect of drawing on a much richer body of existing and future work in these fields than has previously been suggested.

In this paper, we will first describe what we mean by “coordination theory” and give examples of how previous research on computer-supported cooperative work can be interpreted from this perspective. We will then suggest one way of developing this perspective further by proposing tentative definitions of coordination and analyzing its components in more detail.

What is coordination?

We all have an intuitive sense of what the word “coordination” means. When we attend a well-run conference, when we watch a winning basketball team, or when we see a smoothly functioning assembly line we may notice how well coordinated the actions of a group of people seem to be. Often, however, good coordination is nearly invisible, and we sometimes notice coordination most clearly when it is lacking. When we spend hours waiting on an airport runway because the airline can't find a gate for our plane, when the hotel room we thought had been reserved for us is sold out, or when a company fails

repeatedly to capitalize on innovative ideas its researchers develop we may become very aware of the effects of poor coordination.

In order to proceed it is helpful to have a more precise idea of what we mean by "coordination." Appendix A lists a number of definitions that have been suggested for this term. The diversity of these definitions illustrates the difficulty of defining coordination, and also the variety of possible starting points for studying the concept. For our purposes here, however, we believe it is most useful to start with the following "common sense" definition of coordination taken from a dictionary [Amer81]:

the act of working together harmoniously.

We will refer to this as the "broad" definition of coordination, and will suggest a more restrictive "narrow" definition below. It is important to note here, however, that we intend for working together "harmoniously" to include conflict as well as cooperation. Even when a group of actors has strong conflicts of interest or belief, they may still produce results that observers would judge to be "good" or "harmonious." For example, different groups in a company often compete for budget resources and people, and this competition sometimes contributes to the company's ability to produce useful products.

What is coordination theory?

We define *coordination theory* as a body of principles about how activities can be coordinated, that is, about how actors can work together harmoniously. It is important to realize that there is not yet a coherent body of theory in this domain. However, there are theories, concepts, and results from many different fields that could both contribute to and benefit from the development of such general theories.

For instance, it is clear that questions about how people coordinate their activities are central to parts of organization theory, sociology, social psychology, anthropology, linguistics, law, and political science. Important parts of economics and management science also analyze how people can coordinate their work with a special focus on rational ways of allocating resources. Computer science does not deal primarily with people, but different computational processes must certainly "work together harmoniously," and as numerous observers have pointed out, certain kinds of interactions among computational processes resemble interactions among people (e.g., [Fox81], [Hewi86], [Hube88], [Mill88], [Smit81]).

These potential overlaps suggest that coordination theory will be like other interdisciplinary fields that arise from the recognition of commonalities in problems that have previously been considered separately in different fields. For instance, the field of cognitive science grew out of the recognition by researchers in several different fields (e.g., psychology, computer science, and linguistics) that they were dealing separately with similar problems: how can information processing systems (people or computers) do things like use language, learn, plan, remember, and solve problems (e.g., see [Gard85], [Norm80])? Most observers would agree that progress in the new field has benefited significantly from emergent cross disciplinary connections, and the paradigms used have in turn been quite influential in the older fields [Gard85].

In coordination theory, the common problems have to do with coordination: How can overall goals be subdivided into actions? How can actions be assigned to groups or to individual actors? How can resources be allocated among different actors? How can information be shared among different actors to help achieve the overall goals?

In its attempts to find generalizations that apply across disciplines and across levels of analysis, coordination theory resembles earlier work on systems theory and cybernetics (e.g., [Beer67], [Boul56], [Emer69], [Forr80], [vonB50], [Wien61]). We are significantly better equipped for the task of identifying and analyzing coordination processes now,

however, than systems theorists were several decades ago. For instance, new qualitative languages from computer and cognitive sciences (such as object inheritance networks and Petri nets) seem especially promising as tools for formalizing “mid-level theories” like Winograd and Flores’ [Wino86] “conversations for action.” These qualitative mid-level theories are more specific than the quantitative abstractions of systems theory, but more general than specific case studies.

What isn’t coordination theory?

If coordination theory can draw upon so many different fields, is it any more than just the union of these fields? How can we look at a theory and decide whether it is or is not an example of coordination theory? While it is certainly not helpful to include everything in coordination theory, neither do we think it is essential to draw sharp boundaries between what is and is not coordination theory. Instead, as in cognitive science and many other fields, we think certain characteristic questions and approaches will come to typify central examples of coordination theory. For example, theories that apply to only one kind of actor will probably be less important to coordination theory than theories that can be applied to several kinds of actors.

Previous examples of coordination theory and CSCW

With the definition of coordination theory we have just presented, it is clear that some of the work already done in the field of computer-supported cooperative work can be viewed as examples of the use of coordination theory. Even though these authors did not use the term “coordination theory,” each of the following examples involves using ideas about coordination from other disciplines to help develop cooperative work tools:

- (1) Holt [Holt88] describes a theoretical language used for designing coordination tools that is based, in part, on ideas about Petri nets, a formalism widely used in computer science to represent process flows in distributed or parallel systems. This language is part of a larger theoretical framework called “coordination mechanics.”
- (2) Winograd and Flores ([Flor88], [Wino87], [Wino86]) have developed a theoretical perspective for analyzing group action based heavily on ideas from linguistics (e.g., [Sear75]) about different kinds of “speech acts,” such as “requests” and “commitments.” This perspective was a primary basis for designing the Coordinator, a computer tool that helps people make and keep track of requests and commitments to each other.
- (3) Malone [(Malo90)] describes how ideas from organization theory about flexible organizational structures called “adhocracies” [Mint79] and ideas from artificial intelligence about “blackboard architectures” for sharing information among program modules ([Erma80], [Nii86]) contributed to the design of the Information Lens, a system for helping people share information in organizations [Malo87].
- (4) Conklin & Begeman [Conk88] and Lee [Lee90a] describe systems to help groups of people record the structure of arguments (e.g., positions, arguments, and counterarguments) that are based in part on ideas from philosophy and rhetoric about the logical structure of decision-making.
- (5) Turoff [Turo83] used ideas about prices and markets to suggest a computer-based system to help people to exchange services within organizations.

Clearly, drawing a line around these examples and calling them “coordination theory” does not, in itself, provide any benefit. Nor does using ideas about coordination from

other disciplines provide any guarantee of developing useful cooperative work tools. Nevertheless, we feel that considering these examples within the common framework of coordination theory provides two benefits: (1) it suggests that no one of these perspectives is the complete story, and (2) it suggests something about how we might look for more insights of the sort that many people feel have resulted from these previous examples.

In particular, the perspective of coordination theory suggests (1) that we should look to previous work in various disciplines for more insights about coordination, (2) that we should attempt to develop frameworks or concepts that will facilitate such interdisciplinary transfers, and (3) that we should attempt to develop new concepts and theories focused specifically on the questions of coordination that seem central to building cooperative work tools. In the next section, we will take a step in this direction.

TOWARD A FRAMEWORK FOR COORDINATION THEORY

So far, we have claimed that many different disciplines can contribute to our understanding of coordination and that better understanding of coordination will help build useful cooperative work tools. But is it really sensible to use the term “coordination” in describing all the different kinds of phenomena to which we have alluded? For that matter, is there anything in common among these different phenomena, other than some occasional similarities in terminology? As a first step toward answering these questions, we will present in this section our preliminary efforts toward developing a framework for analyzing coordination. This framework is not a “theory of everything;” it is only one approach which we have found helpful in seeing the relationships between different views of coordination.

Components of coordination

According to our broad definition of coordination above, coordination means “the act of working together harmoniously.” What does this broad definition of coordination imply? First of all, what does the word “work” imply? The same dictionary defines “work” as “physical or mental effort or activity directed toward the production or accomplishment of something” [Amer81]. Thus there must be one or more *actors*, performing some *activities* which are directed toward some ends. In what follows, we will sometimes refer to the ends toward which the activities are directed as *goals*. By using the word “harmoniously,” the definition implies that the activities are not independent. Instead, they must be performed in a way that helps create “pleasing” and avoids “displeasing” outcomes, that is, that achieves the goals. We will refer to these goal-relevant relationships between the activities as *interdependencies*. These components and the coordination processes associated with them are summarized in Table 1. (See [Bali86], [Bali81], [Barn64;], [Malo87b], [Malo88], [McGr84], [Mint79] for related decompositions of coordination.)

| <i>Components of coordination</i> | <i>Associated coordination processes</i> |
|-----------------------------------|---|
| Goals | Identifying goals |
| Activities | Mapping goals to activities (e.g., goal decomposition) |
| Actors | Selecting actors Assigning activities to actors |
| Interdependencies | “Managing” interdependencies |

Table 1. Components of coordination.

For example, an automobile manufacturing company might be thought of as having a set of goals (e.g., producing several different lines of automobiles) and a set of actors (e.g., people) who perform activities that achieve these goals. These activities may have various kinds of interdependencies such as using the same resources (e.g., an assembly line) or needing to be done in a certain order (e.g., a car must usually be designed before it is built).

One use of this set of components of coordination is to help facilitate conceptual transfers between disciplines. For instance, elsewhere [Malo88], we have shown how research in selected areas of economics and artificial intelligence can be compared in terms of these dimensions. This comparison suggested a novel insight for economic theorists about the importance of product descriptions, as well as prices, in coordinating resource allocation in markets.

Coordination is attributed to a situation by observers

It is important to realize that the actors involved in a situation may or may not all agree on the identification of all these components. Instead, one or more of these components may be attributed by an observer in order to analyze the situation in terms of coordination. For instance, we may sometimes analyze everything that happens in a manufacturing division as one "activity", while at other times, we may want to analyze each station on an assembly line as a separate "activity."

One very important case of this occurs when the actors have *conflicting goals*, but we choose to analyze the results of their behavior in terms of how well it achieves some goals in which we are interested. For instance, even though two designers on a project team may have strongly opposing views about how a product should be designed, we can evaluate their collective behavior in terms of the quality of the final design. Another important example of conflicting goals occurs in market transactions: All the participants in a market might have the goal of maximizing their own benefits, but we, as observers, can evaluate the market as a coordination mechanism in terms of how well it achieves some global goal such as allocating economic resources to maximize consumer utilities (e.g., [Debr59]).

In practice, situations in which actors have at least partly conflicting goals are nearly universal, and mixtures of cooperation and conflict are quite common (e.g., [Cibo87], [Will85], [Sche69]). When we analyze the coordination in these situations, we must (at least implicitly) evaluate the actors' collective behavior in terms of how well it achieves some overall goals (which may or may not be held by the actors themselves).

A narrow definition of coordination

The broad definition of coordination we have been using includes almost everything that happens when actors work together: setting goals, selecting actors, and performing all the other activities that need to be done. For some purposes, it is useful to be able to focus explicitly on the elements that are unique to coordination, that is, on the aspects of "working together harmoniously" that are not simply part of "working." In our analysis of the broad definition above, the element of coordination that was implied by the word "harmoniously" was interdependencies. Therefore, when we want to focus specifically on the aspects of a situation that are unique to coordination, we will use the following narrow definition of coordination:

the act of managing interdependencies between activities performed to achieve a goal.

Clearly, many important coordination situations involve multiple actors, and in our previous work (e.g., [Malo88]), we defined coordination as something that occurs only when multiple actors are involved. Since then, however, we have become convinced that the essential elements of coordination listed above arise whenever multiple, interdependent activities are performed to achieve goals—even if only one actor performs all of them.

Kinds of Interdependence

Both our definitions of coordination give a prominent role to interdependence: If there is no interdependence, there is nothing to coordinate. There is also a long history in organization theory of emphasizing the importance of interdependence in determining how organizations are structured (e.g., [Thom67], [Galb73], [Lawr67], [Pfef78], [Rock89], [Har90]). This suggests that one useful way to extend the theory of coordination is to ask what kinds of interdependence between activities are possible and how different kinds of interdependence can be managed.

Our preliminary investigations of this question have led us to believe that interdependence between activities can be analyzed in terms of *common objects* that are involved in some way in both actions. For example, the activities of designing and manufacturing a part both involve the detailed design of the part: the design activity creates the design and the manufacturing activity uses it.

These common objects constrain how each activity is performed. Different patterns of use of the common objects by the activities will result in different kinds of interdependences. For example, the parts can be manufactured only after the design is complete and the actor doing the manufacturing has received a copy. We call this pattern of usage (one task creating an object that is used by others) a *prerequisite* constraint. In general, the common object may constrain any or all of the activities that use it. In this case, for example, it might make sense for a designer to consider the constraints that the manufacturing process places on the design and to create a design that will be easier to manufacture.

Table 2 presents a preliminary list of types of interdependencies and coordination processes that can be used to manage them. The table includes both generic kinds of interdependence and specific examples of interdependence that arise in particular situations. We labeled this list “preliminary” because we suspect that there is more structure in the space of kinds of interdependence and processes than is currently reflected in the table.

One use of this table (especially the generic parts of it) may be to help show the relationships between previous work in different disciplines. For instance, much of economics is focused on analyzing market mechanisms for resource allocation, and parts of computer science have focused on questions of synchronizing activities to meet simultaneity constraints.

An even more important use of the approach suggested by this table may be to help generate possible alternative ways of coordinating in a particular situation. For instance, it may be possible to characterize a situation in terms of the kinds of interdependence it involves, and then use a “catalog” of interdependencies and their associated processes to generate a set of alternative processes that could be used to manage the interdependencies. This ability to characterize a space of possible coordination processes for a given set of activities would be useful in understanding how new kinds of coordination tools could lead to new ways of organizing cooperative work.

| <i>Kinds of Interdependence</i> | <i>Common object</i> | <i>Example of interdependence in manufacturing</i> | <i>Examples of coordination processes for managing interdependence</i> |
|---------------------------------|---|--|--|
| <i>Generic:</i> | | | |
| Prerequisite | Output of one activity which is required by the next activity | Parts must be delivered in time to be used | Ordering activities, moving information from one activity to the next |
| Shared resource | Resource required by multiple activities | Two parts installed with a common tool | Allocating resources |
| Simultaneity | Time at which more than one activity must occur | Installing two matched parts at the same time | Synchronizing activities |
| <i>Domain-specific:</i> | | | |
| Manufacturability | Part | Part designed by engineering must be made by manufacturing | Decision-making (e.g., negotiation, appeal to authority) |
| Customer relations | Customer | Both field service and sales personnel deal with same customer | Information sharing (e.g., sharing problem reports) |

Table 2. Preliminary examples of kinds of interdependence.

Example: Coordinating interdependencies between design and manufacturing

This example illustrates how knowing the interdependencies in a situation may suggest alternative ways to manage them. The example is based, in part, on extended field studies of engineering change processes in several manufacturing organizations [Crow90]. In the case of design and manufacturing, an important kind of interdependence results from the common object that is the design of the product to be manufactured.

One simple way technology can help manage these interdependencies is simply by helping to detect them in the first place. For instance, one of the applications we have investigated for the Information Lens and Object Lens systems ([Lai88], [Malo87]) is routing engineering change notices to engineers to whose work is likely be affected by a given change, even when the person making the change does not know who else it will affect.

Whether all the interdependencies are recognized or not, there seem to be at least four basic ways to manage them:

- (1) At a minimum, the designer must create a design and give it to the manufacturer to build. One simple effect of CAD systems, for example, is to make this transfer process easier.

- (2) The designer and the manufacturer can negotiate what the design should be, for example, by iterating the design process or in joint meetings. A variety of electronic meeting support and communication tools could help this process and could make it more desirable relative to alternative ways of managing the same interdependencies.
- (3) Sometimes the need for explicit negotiation can be eliminated by moving some of the knowledge about the constraints of either task from one engineer to another. For instance:
 - (a) Some of the manufacturer's knowledge (the knowledge about the manufacturing constraints, not about how to do the manufacturing) can be made available to the designer, for example, by training the designer in methodologies such as design for manufacturing or by embodying the knowledge in an intelligent CAD system.
 - (b) Some of the designer's knowledge can be transferred to the manufacturer. For example, if a system like gIBIS [Conk88] is used to capture more of the designer's intent as well as the details of the part, the manufacturing engineer might be able to change some details of the design to make the parts easier to build while preserving the intent.
- (4) A third party, such as a common superior, may be able to resolve problems as they arise or to give enough initial direction that problems do not arise.

This analysis seems to be easily transferred to other domains. For example, a bank and a potential borrower have to agree on a common object, a loan. The typical approach seems to be case (1) above: the bank offers a loan with its standard terms and a person who wants the loan takes it or leaves it. In some cases, the bank and the borrower negotiate the details of the loan, case (2) above. Finally, one can imagine transferring some of the bank's knowledge about making loans, for example, to a computer program that a potential borrower could run to explore possible loan conditions (case (3a)), or to a third party who would suggest which bank would be best for a given applicant (case (4)).

Processes underlying coordination

In attempting to characterize more precisely different coordination processes, we have found it useful to describe them in terms of successively deeper levels of underlying processes, each of which depends on the levels below it. Table 3 shows a preliminary diagram of the levels we have used. For instance, most of the coordination processes listed in the last column of Table 2 require that some decision be made and accepted by a group (e.g., what goal will be selected or which actors will perform which activities). Group decisions, in turn, require members of the group to communicate in some form about the goals to be achieved, the alternatives being considered, the evaluations of these alternatives, and the choices that are made. This communication requires that some form of "messages" be transported from senders to receivers in a language that is understandable to both. Finally, the establishment of this common language and the transportation of messages depends, ultimately, on the ability of actors to perceive common objects such as physical objects in a shared situation or information in a shared database (e.g., see [Such87]). These layers are analogous to abstraction levels in other systems, such as protocol layers for network communications.

Even though the strongest dependencies appear to be downward through these layers, there are also times when one layer will use processes from the layers above it. For instance, a group may sometimes use decision-making processes to extend the common language it uses to communicate (e.g., see [Lee90b]), or a group may use coordination processes to assign decision-making activities to actors.

| <i>Process Level</i> | <i>Components</i> | <i>Examples of Generic Processes</i> |
|------------------------------|---|--|
| Coordination | goals, activities, actors, resources, interdependencies | identifying goals, ordering activities, assigning activities to actors, allocating resources, synchronizing activities |
| Group decision-making | goals, actors, alternatives, evaluations, choices | proposing alternatives, evaluating alternatives, making choices (e.g., by authority, consensus, or voting) |
| Communication | senders, receivers, messages, languages | establishing common languages, selecting receiver (routing), transporting message (delivering) |
| Perception of common objects | actors, objects | seeing same physical objects, accessing shared databases |

Table 3. Processes underlying coordination.

Example : Selecting actors to perform activities

To see how this framework can be used to analyze coordination processes, let us consider the part of the activity assignment process that involves selecting which actors will perform which activities. For this example, we will analyze one particular method that can be used for this process: a competitive bidding scheme like that used in many kinds of markets. Our analysis will draw upon the version of this process formalized by Smith and Davis [Smit81] and extended by Malone ([Malo87b], [Malo88]).

In this scheme, a client first broadcasts an announcement message to all potential contractors. This message includes a description of the activity to be performed and the qualifications required of potential contractors. The potential contractors then use this information to decide whether to submit a bid on the action. If they decide to bid, their bid message includes a description of their qualifications and their availability for performing the action. The client uses these bid messages to decide which contractor should perform the activity and then sends an award message to notify the contractor that has been selected.

In this case, the decision to be made is which contractor will perform a specific action. The choice results from a multi-stage process in which contractors decide whether to propose themselves as alternatives (by submitting bids) and clients decide which contractor to select based on their evaluations of the contractors' bids. The actors communicate by exchanging messages, and we can regard these messages as including representations of common objects (such as activities and bids) which both senders and receivers can perceive.

Viewing the activity assignment process in this way, immediately suggests other possibilities for how it can be performed. For instance, an authority-based decision-making process might be used in which a manager simply assigns activities to people who have implicitly already agreed to accept such assignments. This view also suggests how computer tools could be used to support a bidding process for task assignments in human organizations (e.g., see [Malo87a], [Turo83]).

CONCLUSIONS

In this paper, we have argued that many different disciplines can contribute to our understanding of coordination and that a better understanding of coordination can help us build useful cooperative work tools. In order to support these claims, we have shown examples of interdisciplinary transfers of ideas about coordination that have already provided useful insights for cooperative work tools, and we have sketched out the beginnings of a framework that can facilitate such interdisciplinary transfers and lead to the development of new general theories about coordination.

Clearly there is much left to be done. We hope, however, that the perspective we have suggested here will help build tools that enable people to work together more effectively and more enjoyably.

ACKNOWLEDGEMENTS

This work was supported, in part, by Digital Equipment Corporation, the National Science Foundation (Grant Nos. IRI-8805798 and IRI-8903034), the MIT International Financial Services Research Center, and General Motors/Electronic Data Systems.

Parts of this paper were included in a previous working paper [Malo88] and in proposals submitted to the National Science Foundation. We are especially grateful to Deborah Ancona, John Carroll, Michael Cohen, Randall Davis, John Little, and Wanda Orlikowski for comments on earlier versions of the paper, and to participants in numerous seminars and workshops at which these ideas have been presented.

APPENDIX: PREVIOUSLY SUGGESTED DEFINITIONS OF COORDINATION

“The operation of complex systems made up of components.” [NSF89]

“The emergent behavior of collections of individuals whose actions are based on complex decision processes.” [NSF89]

“Information processing within a system of communicating entities with distinct information states.” [NSF89]

“The joint efforts of independent communicating actors towards mutually defined goals.” [NSF89]

“Networks of human action and commitments that are enabled by computer and communications technologies.” [NSF89]

“Composing purposeful actions into larger purposeful wholes.” [Holt89]

“Activities required to maintain consistency within a work product or to manage dependencies within the workflow.” [Curt89]

“The additional information processing performed when multiple, connected actors pursue goals that a single actor pursuing the same goals would not perform.” [Malo88]

REFERENCES

- [Amer81] American Heritage Dictionary. Boston: Houghton Mifflin, 1981.
- [Bali86] Baligh, H. H. Decision rules and transactions, organizations and markets. *Management Science*, 32, 1480–1491, 1986.
- [Bali81] Baligh, H. H., Burton, R. M. Describing and designing organizational structures and processes. *International Journal of Policy Analysis and Information Systems*, 5, 251–266, 1981.
- [Barn64] Barnard, C. I. *The Functions of the Executive*. Cambridge, MA: Harvard University, 1964.
- [Beer67] Beer, S. *Cybernetics and Management* (2nd ed.). London: English Universities Press, 1967.
- [Boul56] Boulding, K. E. *The Image*. Ann Arbor, MI: University of Michigan, 1956.
- [Cibo87] Ciborra, C. U. Reframing the role of computers in organizations: The transaction costs approach. *Office Technology and People*, 3, 17-38, 1987.
- [Conk88] Conklin, J., Begeman, M. L. gIBIS: A hypertext tooling for exploratory policy discussion. In Tatar, D. (Ed.), *Proceedings of the 2nd Conference on Computer-supported Cooperative Work* (pp. 140–152). New York: ACM, 1988.
- [Crow90] Crowston, K. *Modeling Coordination in Organizations*. Ph.D. Dissertation, Sloan School of Management, Massachusetts Institute of Technology, Forthcoming (1990).
- [Curt89] Curtis, B. Modeling coordination from field experiments. In *Organizational Computing, Coordination and Collaboration: Theories and Technologies for Computer-Supported Work*. Austin, TX, 1989.
- [Debr59] Debreu, G. *Theory of value: An axiomatic analysis of economic equilibrium*. New York: Wiley, 1959.
- [Emer69] Emery, J. C. *Organizational Planning and Control Systems: Theory and Technology*. New York: MacMillan, 1969.
- [Erma80] Erman, L. D., Hayes-Roth, F., Lesser, V. R., Reddy, D. R. The HEARSAY-II speech understanding system: Integrating knowledge to resolve uncertainty. *Computing Surveys*, 12(2), 213–253, 1980.
- [Flor88] Flores, F., Graves, M., Hartfield, B., Winograd, T. Computer systems and the design of organizational interaction. *ACM Transactions on Office Information Systems*, 6(2), 153–172, 1988.
- [Forr80] Forrester, J. W. *Systems dynamics*. New York: North-Holland, 1980.
- [Fox81] Fox, M. S. An organizational view of distributed systems. *IEEE Transactions on Systems, Man and Cybernetics*, 11(1), 70–79, 1981.
- [Galb73] Galbraith, J. R. *Designing Complex Organizations*. Reading, MA: Addison-Wesley, 1973.

- [Gard85] Gardner, D. *The Mind's New Science: A History of the Cognitive Revolution*. New York: Basic, 1985.
- [Hart90] Hart, P. & Estrin, D. Inter-organization computer networks: Indications of shifts in interdependence. *Proceedings of the ACM Conference on Office Information Systems*, Cambridge, MA, April 25-27, 1990.
- [Hewi86] Hewitt, C. Offices are open systems. *ACM Transactions on Office Systems*, 4(3), 271-287, 1986.
- [Holt89] Holt, A. *Personal communication.*, 1989.
- [Holt88] Holt, A. W. Diplans: A new language for the study and implementation of coordination. *ACM Transactions on Office Information Systems*, 6(2), 109-125, 1988.
- [Hube88] Huberman, B. A. (Eds.). *The Ecology of Computation*. Amsterdam: North-Holland, 1988.
- [Korn81] Kornfeld, W. A., Hewitt, C. The scientific community metaphor. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-11, 24-33, 1981.
- [Lai88] Lai, K. Y., Malone, T., Yu, K.-C. Object Lens: A spreadsheet for cooperative work. *ACM Transactions on Office Information Systems*, (Oct), 1988.
- [Lawr67] Lawrence, P. R. & Lorsch, J. W. *Organization and Environment*. Boston: Graduate School of Business Administration, Harvard University, 1967.
- [Lee90a] Lee, J. Sibyl: A qualitative decision management system. In Winston, P. (Ed.), *Artificial Intelligence at MIT: Expanding Frontiers* Cambridge, MA: MIT Press, 1990.
- [Lee90b] Lee, J., Malone, T. W. Partially Shared Views: A scheme for communicating among groups that use different type hierarchies. *ACM Transactions on Information Systems*, 8, 1-26, 1990.
- [Malo87a] Malone, T. W. Computer support for organizations: Towards an organizational science. In Carroll, J. (Ed.), *Interfacing Thought: Cognitive Aspects of Human Computer Interactions* Cambridge, MA: MIT Press, 1987.
- [Malo87b] Malone, T. W. Modeling coordination in organizations and markets. *Management Science*, 33, 1317-1332, 1987.
- [Malo88a] Malone, T. W. *What is coordination theory?* (Working paper #2051-88). Cambridge, MA: MIT Sloan School of Management, 1988.
- [Malo90] Malone, T. W. Organizing information processing systems: Parallels between organizations and computer systems. In Zachary, W., Robertson, S., Black, J. (Ed.), *Cognition, Computation, and Cooperation* (pp. 56-83). Norwood, NJ: Ablex, 1990.
- [Malo87c] Malone, T. W., Grant, K. R., Turbak, F. A., Brobst, S. A., Cohen, M. D. Intelligent information-sharing systems. *Communications of the ACM*, 30, 390-402, 1987.

- [Malo88b] Malone, T. W., Smith, S. A. Modeling the performance of organizational structures. *Operations Research*, 36(3), 421–436, 1988.
- [McGr84] McGrath, J. E. *Groups: Interaction and Performance*. Englewood Cliffs, NJ: Prentice-Hall, 1984.
- [Mill88] Miller, M. S., Drexler, K. E. Markets and computation: Agoric open systems. In Huberman, B. A. (Ed.), *The Ecology of Computation* (pp. 133–176). Amsterdam: North-Holland, 1988.
- [Mins87] Minsky, M. *The Society of the Mind*. New York: Simon and Schuster, 1987.
- [Mint79] Mintzberg, H. *The Structuring of Organizations*. Englewood Cliffs, NJ: Prentice-Hall, 1979.
- [Nii86] Nii, P. The blackboard model of problem solving. *The AI Magazine*, (Spring), 38–53, 1986.
- [Norm80] Norman, D. A. Twelve issues for cognitive science. *Cognitive Science*, 4, 1–32, 1980.
- [NSF89] NSF-IRIS. *A report by the NSF-IRIS Review Panel for Research on Coordination Theory and Technology*. Available from NSF Forms & Publications Unit, 1989.
- [Pfef78] Pfeffer, J. & Salancik, G. R. *The External Control of Organizations: A Resource Dependency Perspective*. New York: Harper & Row, 1978.
- [Sche60] Schelling, T. C. *Strategy of Conflict*. Cambridge: Harvard University Press, 1960.
- [Sear75] Searle, J. R. A taxonomy of illocutionary acts. In Gunderson, K. (Ed.), *Language, Mind and Knowledge* (pp. 344–369). Minneapolis, MN: University of Minnesota, 1975.
- [Rock89] Rockart, J. F. & Short, J. E. IT and the networked organization: Toward more effective management of interdependence. In M. S. Scott Morton (Ed.), *Management in the 1990s Research Program Final Report*. Cambridge, MA: Massachusetts Institute of Technology, 1989.
- [Smit81] Smith, R. G., Davis, R. Frameworks for cooperation in distributed problem solving. *IEEE Transactions on Systems, Man and Cybernetics*, 11(1), 61–70, 1981.
- [Such87] Suchman, L. A. *Plans and Situated Actions: The Problem of Human Machine Communication*. Cambridge: Cambridge University Press, 1987.
- [Thom67] Thompson, J. D. *Organizations in Action*. New York: McGraw-Hill, 1967.
- [Turo83] Turoff, M. *Information, value, and the internal marketplace* (Unpublished manuscript). New Jersey Institute of Technology, 1983.
- [vonB50] von Bertalanffy, L. The theory of open systems in physics and biology. *Science*, 111, 1950.

- [Wien61] Wiener, N. *Cybernetics: Or Control and Communication in the Animal and the Machine* (2nd ed.). Cambridge, MA: MIT Press, 1961.
- [Will85] Williamson, O. *The Economic Institutions of Capitalism*. New York: Free Press, 1985.
- [Wino87] Winograd, T. A language/action perspective on the design of cooperative work. *Human Computer Interaction*, 3, 3-30, 1987.
- [Wino86] Winograd, T., Flores, F. *Understanding computers and cognition: A new foundation for design*. Norwood, NJ: Ablex, 1986.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

© 1990 ACM 089791-402-3/90/0010/0370 \$1.50