

Message Management Systems: Concepts, Motivations, and Strategic Effects*

Steven O. Kimbrough
University of Pennsylvania

Scott A. Moore
University of Michigan Business School

May 26, 1992

Abstract

This paper motivates the need for system-level message management software. It begins by considering information flows in the work place as a source of potential gains in efficiency. We next investigate workflow automation and EDI as indicative of current technologies applied to work processes and message management. Having described current technology and our vision of work processes, we propose an alternative, general purpose, software technology for supporting application-to-application communication. Problems of EDI, of process-to-process communication, and of describing information items are discussed in terms of the communication problems they present. We then justify the need for this kind of software and lay out the criteria (or plausibility conditions) for evaluating a proposal for this sort of system software. The use of a formal communication language is proposed as a common solution to these problems. This proposal is examined in the context of the EDI problem, so as to demonstrate how the proposal might work in practice. Practical benefits of the proposal are discussed that highlight the impact such a technology might have on business practices. The proposed solution is measured against the plausibility conditions presented earlier in the paper; it is found to be

*Thanks to David C. Blair, Eric K. Clemons, Michael Gordon, John Holland, Ronald M. Lee, and Paul Thagard for pertinent discussions and illuminating insights, but nothing here is to be blamed on them. Thanks also to the referees for helping us focus and tighten our paper. The research was supported in part by U.S. Coast Guard contract number DTICG39-86-C-E92204, Steven O. Kimbrough principal investigator. Published in *Journal of Management Information Systems*, 9:2 (Fall 1992), 29-52.

sufficient in some cases and in need of further investigation in others. We then discuss the industrial-organizational implications of having such a technology available. It is hypothesized to affect the number and form of cooperative business relationships as well as their scope and depth. It is also hypothesized to provide advantages to those firms that quickly adopt the technology.

Biographies

Steven O. Kimbrough is an Associate Professor at the Department of Decision Sciences, The Wharton School, University of Pennsylvania. He is currently the Davidson Visiting Professor at the University of Michigan, School of Business Administration. His main research interests are in the fields of logic modeling, and decision support and expert systems. His active research areas include: computational approaches to belief revision and nonmonotonic reasoning, formal languages for business communication, and knowledge-based decision support systems. Since 1985, he has been principal investigator for the U. S. Coast Guard's KSS (knowledge-based decision support systems) project.

Scott A. Moore is a Ph.D. student in the Decision Sciences Department of The Wharton School of the University of Pennsylvania. His current research interests include formal languages for business communication, logic programming, decision support systems, fleet mix, model management, and document retrieval. Much of his time is spent designing and implementing decision support systems for the U.S. Coast Guard. His education includes a B.S. in mathematics from Furman University, and an M.S.M. from the Georgia Institute of Technology. Mr. Moore has also been a manager of information services at a human resources association.

1 Introduction

The future of computer networking will, it is widely agreed, be a situation in which “anything, anytime, anywhere” may be communicated easily and at low cost [28]. Were anything approximating this situation to obtain, then it is obvious that there would be profound strategic consequences for very many organizations. While it is clear *that* there will be such consequences, it is far from clear *what* those consequences will be. In fact, the strategic consequences of “anything, anytime, anywhere” (AAA) communications will not be deeply understood—either in the long run or in the interim—without a similarly deep (we do not say “detailed” or “picayune”) understanding of the capabilities of the technology that will deliver such communications systems.

Much remains to be learned, however, before AAA communications systems will be both technically possible and economically attractive. The following passage, from a recent special issue of *Scientific Amer-*

ican on “Communications, Computers and Networks,” nicely categorizes the kinds of knowledge that must be acquired.

To escape the present chaos and to fashion our computers and networks into a true information infrastructure, we must endow networks with three key capabilities: flexible information transport capabilities, common services and common communications conventions. [16, page 65]

Until “the present chaos” has been overcome and “our computers and networks [fashioned] into a true information infrastructure,” it will be difficult to foresee the strategic and industrial-organizational (IO) consequences of such advanced communications systems. We may confidently forecast that flexible transport capabilities will be achieved, and will be widely available, very fast, and comparatively cheap. Further, it is plain that this will permit a wide variety of common services to be offered. Much else, however, is far from clear. What will those common services will be? What will tend to be offered publicly, what privately? What, if any, opportunities will arise for creating barriers to entry and exit to markets, and for other economic relationships? Will there be capital or knowledge-based barriers to entry? Will there be switching costs or first mover effects? How will this affect the ways organizations conduct business—both within the firm and with other firms? All of this is at present highly uncertain, as is the coming nature of the “common communications conventions.” And they—these and related questions, and the communications conventions—are highly interdependent.

Our aim in this paper is twofold:

1. We wish to shed some light upon the issue of the “common communications conventions” that will be required for AAA communications (or anything like it). Specifically, our thesis is that, for the sake of effectively establishing common communications conventions, a strong case can be made in favor of general, system-level software for managing messages flowing between applications. We call such software a *message management system* (MMS). Broadly, the case for it is similar to the case for database management systems. Further, we believe a strong case can be made that messages should be expressed in general-purpose protocols, which we call *formal languages for business communication* (FLBCs). Again, the justification for FLBCs is broadly analogous to that for database query languages. In both cases there is much to be gained from having a well-designed, publicly-defined, and widely-available standard approach to a problem. (In §4 we refer to this as the *systematization of application functionality*.) We hope to show that FLBC messages controlled by an MMS is a plausible technological basis for AAA communications. We shall

refer to this as our *technical point*.

2. We wish to advance certain remarks on the strategic and industrial-organizational (IO) consequences of the sort of AAA communications that will be possible with widespread deployment of MMSs. These remarks do not constitute a complete assessment of what will happen. Instead, they are offered as plausible, testable hypotheses.

We are broadly in agreement with the “move to the middle” hypothesis of Clemons and his co-workers [11, 13]. They argue that improvements in information technology (e.g., AAA communications systems) will both reduce communication and production costs (thereby leading to more cooperative relationships) *and* reduce risks of opportunism among contractors (thereby reducing transaction risk, thus leading to more outsourcing). They call this combination a “move to the middle.” We have two additional, complementary hypotheses. First, these advanced communication systems will often change the nature and dynamics of switching costs to all contracting parties, leading to much more symbiotic economic relationships that extend and evolve over time. Second, because substantial investment will be required to convert existing systems, first-mover opportunities and barriers to entry dependent upon the cost of converting will be created. We shall refer to this as our *strategic point*.

The remainder of the paper is organized as follows. In §2 we present the concept of information events. Our purpose in doing this is to focus the discussion on certain information- and communication-intensive activities that are common to many types of work. In §3 we discuss and compare workflow automation systems (based on document image processing) and EDI (electronic data interchange) applications in light of the information events concept presented in §2. Our aim here is to display the benefits that can be, and are being, achieved with existing applications of communications technology. Further, we discuss the limitations of these systems and present information to be used subsequently, in §8, when we discuss the strategic and industrial-organizational meaning of having advanced (AAA) communications.

In §§4–7, we present, discuss and provide an initial evaluation of our concept of an MMS. §8 is, as noted above, devoted to a discussion of the strategic significance of AAA communications systems, on the assumption that, in the near term, their capabilities and characteristics will broadly resemble those of the message management concepts presented in this paper. Finally, in §9 we summarize the findings of this paper.

2 Information Events

It is common and entirely natural to think of work processes as consisting of a series of tasks performed by servers. Materials and information flow between servers and within a server during the completion of a job. In this paper we concern ourselves with the processing of the *information events* related to the completion of a job.¹ Informally, we think of these information events as being activities such as retrieving documents, making decisions, sending notifications, etc. We can identify several types of information events:

1. **Transaction processing events.** Given the requisite data is at hand, these events, or processes, perform meaningful tasks for the work at hand by transforming the given information from one form to another. Examples include: calculating the outstanding principal on a loan, determining whether an application should be approved, and deciding where to forward a job for additional processing.
2. **Message formation events.** Processes communicate by sending messages of various sorts. The creation of a message for the purpose of communicating with another process is a message formation event. In terms of our framework, messages are formed for such purposes as requesting information, announcing decisions, and instructing a down-stream processor.
3. **Message transmission events.** Once messages are formed they must be transported to their destinations. This is done with transmission events. These are a peculiar hybrid of the material and the informational, and they fall under what was above referred to as “flexible information transport capabilities.”
4. **Message interpretation events.** These are the dual of message formation events; they are the mapping from a received message to the transaction processing software.

Just as in a manufacturing environment, where work can be made more efficient by looking at the flow of materials among and within servers, in an information-intensive environment work can be made more efficient by looking at the information events. We can ask *How is it possible to improve the processing of information events in order to hasten production, or to reduce errors and rework, or to reduce costs, or to enhance maintainability and flexibility, or to improve manageability of the system?* We believe that one very promising way to do this is to develop and deploy general-purpose software systems for

¹Thanks to Michael Gordon for the term, and to both David Blair and Michael Gordon for fruitful discussions on the subject.

handling message formation, message transmission, and message interpretation events—this is the MMS. Beginning in §4 we shall articulate this concept and its justification more precisely. First, however, we turn to a discussion of how less technically ambitious approaches are faring in today’s world. This discussion will serve to support both the functional promise and the need for MMSs.

3 Present Experience

The last decade saw firms invest heavily in increasing the efficiency of material flows and handling in manufacturing. Similar levels of activity are beginning to be seen in process management in information-intensive environments. These technologies promise increases in efficiency, quality, and response time for office work. In this section we discuss two examples of this technology: workflow automation and electronic data interchange (EDI). Both technologies deliver on their promises but both have shortcomings that are clearly seen in light of our work flow model and in comparison with each other. We conclude this section with a comparison of these two systems and a discussion of their shortcomings.

3.1 Workflow Automation

Much business activity involves storing and distributing documents. U.S. businesses have been estimated to generate one billion pieces of paper daily [20]. Ernst & Young reported that the insurance industry spends billions of dollars annually handling the paper containing up to 90% of the information insurance transactions require [29]. At least 70% of this cost is for salaries while another 20% is for storage costs [29]. Much of the cost is attributed to physically transferring the document from one person to another; one industry executive put this cost at up to \$50 per document [42]. To begin to control these costs firms are increasingly turning to document imaging processing systems (DIPSs) and associated workflow automation software.

A DIPS manages the creation, indexing, storage, and retrieval of electronic document images. The combination of this technology with current technology for mass storage and communications provides many employees with high-speed access to millions of document images. Incomplete or inaccurate document descriptions limit the system’s effectiveness (more on this later). Document image processing is, however, often an immense improvement over paper-based systems. There are well-known drawbacks to paper-based systems: documents are not filed consistently, are lost, and can only be used by one person at a time. A DIPS removes or reduces these drawbacks.

In addition, a DIPS provides a tool for managing the flow of paperwork within an office. This tool generally has been described as *workflow automation software*. Its main function is to route a document electronically through an office. The route of a particular document depends on the document type and decisions made by people who work with the document. In light of the enormous costs of paper-handling noted above, it is clear that significant cost savings and great efficiencies can be gained through intelligent application of this technology. The potential impact of such a system on the workings of a company are enormous:

- Procedures for handling documents can be automated, at least in part. The computer can be made to control the routes of documents through a business. This provides several benefits:
 - Consistent operational procedures. Example: Great Western Bank [49].
 - Auditing and management of document flows. Examples: Burroughs Wellcome Co. [8], Veterans Administration [19].
 - Easy location and retrieval of a document, no matter where it is in the office. Example: Lincoln National Corp. [51].
- Throughput can be increased and turnaround time reduced. The delay in passing a document from one person to another is effectively removed. Examples: Burroughs Wellcome Co. [8], PNC Financial Corp. [49], Lomas Mortgage USA [49], Pacific Mutual Life Insurance [29].
- Possibilities for documents getting “lost in the shuffle” can be minimized. When using a computerized DIPS, the possibility of losing a document is reduced dramatically since the electronic document is always available on the disk. Examples: British Petroleum Exploration [19], Lincoln National Corp. [51].

The result of all these changes is that customer service can be markedly improved. Companies have made the following claims relative to customer service improvements: more responsive to needs of clients [19], no complaints about a lost insurance claim [51], and reduction in response time [49].

The most significant problem with workflow automation software is that information in the documents is not machine processable. For example, a document may be identified as a purchase requisition and sent to purchasing. To process this document a clerk must view the requisition and determine what to do with it. Company procedures are followed to determine if the requisition should be approved thereby allowing bids to be gathered and a purchase order submitted. The problem? A clerk is still required to process the document. Though the computer knows where the document should be transferred, it does

not know about the contents of the document (e.g., what the requisition is for). The clerk must create the purchase order by reading the information off the requisition and bid before typing the information onto the purchase order. The workflow automation software improves the efficiency of the transfer of documents among workers and provides better access to information needed to work on a document; however, it does not do any of the processing on a document. In terms of our information event concept, a DIPS provides essentially no support for automating the handling of transaction processing events.

3.2 EDI

Electronic data interchange (EDI) can be characterized as application-to-application communication of business information in machine-readable (hence highly structured) form. In other words, “[t]he basic principle of EDI is that computer-generated trading documents, such as orders and invoices, are transmitted directly to a company’s trading partner’s computers across a telecommunications network.” [30, p. 6] Companies of all sizes are using EDI to streamline operations and reduce administrative costs. EDI has been credited with providing many benefits to many companies, including the following:

- Increase sales. Examples: trading partner of R.J. Reynolds [15, p. 41], Boise Cascade [45].
- Improve cash management. Example: General Electric [15, p. 41].
- Improve the efficiency of interorganizational shipments. Examples: Cressona Aluminum [1], Tesco [24, p. 104], R.J. Reynolds [15, p. 41], Rover [10], Boise Cascade [45].
- Decrease the number of times documents are processed by human operators. Example: Cressona Aluminum [1], Tesco [24, p. 104], Boise Cascade [45].

This in no way exhausts the benefits, the citations, or the companies involved but does provide a glimpse at the possibilities.

It has been widely reported [3, 24, 46] that a consequence of the decision to commit to integrating EDI into a company’s operations is that the company changes the way it does business. EDI “is a discipline forcing businesses to eliminate imperfect, wasteful, resource-consuming complexity in working practices” [3, p. 121]. EDI is taken to be the method by which the company can tightly integrate its order entry, manufacturing, shipping and receiving, and accounts payable. This is a complex process but can yield great benefits to the company.

The market has shown that EDI is a good idea. Current market growth vindicates those who have recognized the benefits of this

technology. However, the technology has important shortcomings. We address these below (in §5.1) when we address the group of abstract tasks an MMS should be designed to address.

3.3 A Comparison

EDI and workflow automation are two technologies that are similar, complementary, and insufficient.

- *Similar.* Both are used to manage the processes of information sharing and transfer. In terms of our work flow modeling framework, both represent attempts at an MMS—one uses structured messages to transfer information contained in that structured message (EDI), the other uses structured messages to manage the routing of unstructured documents (workflow software). Both are most effective when combined with a redesign of business processes ([19, 41, 44, for workflow automation] and [1, 24, for EDI]).
- *Complementary.* EDI is only useful for well-defined business documents; workflow automation software has heretofore been attached to DIP systems and unstructured documents.
- *Insufficient.* Both systems are insufficient for the long-run communications needs of industry. EDI is only used for well-defined trading documents thereby ignoring a vast part of business operations. EDI handles all four types of information events but only for a small subset of business activity. Workflow automation software based on DIPS aids only in the transfer and coordination of document flows but not in the processing of the information. Workflow automation software provides moderate support for message formation, transmission, and interpretation events but almost no support for transaction processing events. Also, the way in which the documents are indexed means that these systems work best on documents that can be easily and accurately indexed.

A compromise between these two solutions is needed that allows 1) non-trading documents to be shared between companies, 2) the computer to process the transferred information, and 3) business process re-engineering to more easily integrate disparate parts of a company and industry. The benefits of such a system is clear in light of the discussion in §§3.1–3.2. Many types of documents could be routed automatically through a company using such a technology. Once the document arrives at its destination, the computer acts on the document since its contents are known to the system. Basic loans can be approved, meetings tentatively scheduled, inventory can be checked, requests for information can be processed—the possibilities are almost

endless. People are not taken out of the processing loop but can be relieved of mundane processing and paper-shuffling chores. In the next four sections we propose a solution that has these features.

4 Background for Message Management

One form of progress in application software is in what we call (for lack of a better and more established term) the *systematizing of application functionality*. This progress occurs in identifying functionality common to many applications, generalizing the requirements for such functionality, and implementing code that meets the generalized requirements and that is largely application-independent. What was once done by application software is now done by system software, or at least by shells and other reusable tools.

Prominent and successful examples of such systematization abound. Very many applications need to access structured data in files, and need to do so in a reliable, reasonably efficient, maintainable way, consistent with good management practices. This functionality has been systematized in the form of database management systems. Report writers, communications packages and protocols, forms programs, DSS generators, expert system shells, XWindows and Motif, spreadsheet packages, model management systems, and—at the low end—subroutine libraries can all be seen as instances of progress in systematizing application software functionality.

Our aim in what follows is to propose and discuss what we call *message management systems* (MMSs). It is our claim that such systems—properly conceived and implemented—would be a promising addition to the list of successful examples of systematizing application software functionality. What if ordering supplies from a vendor were—from the user’s point of view—as simple as saving a file? What if negotiating delivery of goods could be automated and, again, were as simple as saving a file? What if sending a message to a commercially-acquired software module (e.g., a user interface management system) were as simple as retrieving a record from a database system? The purpose of an MMS, as we conceive it, is in large part to provide the means for realizing these sorts of improvements.

In order that any such proposal for a new system be taken as plausible, several conditions must be met.

1. The basic concept of the system must be articulated in a clear and operationalizable manner.
2. The functionality that is to be systematized must be identified

and shown to be sufficiently general to make the systematization plausibly worth the cost and effort. The benefits of successful systematization should be identified.

3. A collection of general requirements for such systems must be identified.
4. A general approach to designing, implementing, and maintaining such systems should be articulated.
5. The computational costs associated with the proposed design approach should be shown to be acceptable.

In what follows we aim mainly to present the motivations for, and the concept of, an MMS (in §5 and §6, respectively). Thus, what we have to say here focuses mostly on conditions 1 and 2, above. Briefly and in passing, we shall have something to say about items 3, 4 and 5, but detailed discussion of them, even in beginning fashion, must be left for another paper.

5 Families of Problems

It will be helpful, in introducing our concept of an MMS, to take an abstract, general look at categories of problems pertaining to the handling of information events. Thus, we begin our broader motivation for an MMS by identifying three related groups of problems: EDI problems, module-to-module communication problems, and data description problems.

5.1 EDI Problems

The concept of EDI has both a broad and a (comparatively) narrow sense. In its narrow sense, as we have been using and shall use the term, EDI refers to the actual ongoing replacement of standard business documents with computer-to-computer messages expressed in established protocols. EDI in the broad sense—called here and elsewhere *AAA-EDI*²—is even more ambitious, both technically and in the level of its potential effects on practice.

We can understand the technical concept of narrow-sense EDI simply by examining the products of existing standards organizations and of the firms that use the standards. No such precision is possible with EDI understood in the broad sense, *AAA-EDI*. Here it is only possible to note that the basic concept is that EDI standards and practice should encompass a much broader range of documents and information

²For “anything, anytime, anywhere” EDI [28, 34, 43].

than is currently covered. These include such documents as case histories, insurance claim information, hospital cost reports, credit checking information, enrollee listings, earnings and benefits reports, license applications, and regulatory submissions. With AAA-EDI, then, it is widely envisioned that an era of *electronic commerce* will be possible [53].

Understood either broadly or narrowly, EDI faces a number of recurring challenges and problems. Briefly, we would characterize them as follows (see also [34]):

1. **Existing EDI standards are not expressively powerful.**

EDI standards are now organized as record-passing protocols (see [33] for the distinction between record-passing protocols and languages used for communication) for particular *transaction sets*, e.g., invoices, advance shipping notices, bills of lading and so forth. There is a different record structure for each transaction set, so that software that interprets messages from one protocol is largely useless for handling messages from any other transaction set. This arrangement severely limits the extensibility, and hence the expressive power, of the protocols. Also, the existing protocols do not allow representation of message operators (e.g., boolean combinations of messages; see below for a longer list of such operators), and it is difficult to see how this could be done effectively with a record-passing protocol approach. The existing standards are limited in philosophy to representing a typical invoice, purchase order, etc. There is no underlying theory of what needs to be said, nor is there an underlying technical approach that can plausibly accommodate what will need to be said in the future.

For an example of the type of difficulty we envision, consider the following. We have worked with a company using EDI that wanted to say “The shipment we said we shipped did not actually get shipped.” This type of message was in fact frequently required, but the companies involved had to handle it with telephone calls because their EDI system could not represent this message. Eventually, an ad hoc solution was developed whereby several fields were used in a non-standard way to send this message. This would be a fine solution to this problem if it could be guaranteed that many more such messages would not be discovered causing other such ad hoc solutions to be tacked-on to the system. This guarantee cannot be provided.

2. **Standards keep changing.** A frequently-heard lament is “The nice thing about EDI standards is that there are so many to choose from.”
3. **Mapping is a serious problem.** Mapping is the translation be-

tween messages and the applications that use (or generate) them. Automated mapping is required for any reasonably advanced EDI system, certainly for any AAA-EDI system. No theory or general solution currently exists, however, for performing mapping, hence it is often an expensive and problem-ridden function.

4. **Dialog management is a serious problem.** Individual messages (invoices, bills of lading, etc.) are usually sent as part of a larger conversation, or dialog, between trading partners. The career of any particular dialog contains crucial information pertaining to mapping, validating messages, and determining responses. Although standards for EDI-based negotiations are beginning to be discussed [18], we are a long way from having practical and powerful dialog modeling and management capabilities.
5. **EDI applications tend not to generalize.** A number of innovative EDI applications have been reported, some of which have significant consequences for industrial organization [21, 34]. Reports from the field, however, are nearly unanimous in saying that there is a high fixed cost in setting up any innovative EDI application and that this investment must largely be repeated, even for similar applications of EDI.

Before beginning to discuss, however briefly and tentatively, solutions to these problems, we turn to a quick discussion of two related families of problems.

5.2 Process-to-Process Communication Problems

In a typical EDI system, processes running on separate computers (usually belonging to separate firms) send messages to one another of some special type (orders, invoices, etc.). This is just a special case of automated communication between (and among) separate processes, where in the general case the messages may serve many purposes and the processes may or may not be on the same machine. General, principled approaches to process-to-process communication have only begun to be explored, mainly as a facet of the area of distributed AI (DAI) [7, 22, 23, 25].

Just as it would be nice (with EDI) to have a workable, standard way to send orders, invoices, and other business documents, so it would be nice (with process-to-process communications) to have standard ways to send messages to software modules. With such standards in place software modularity would be greatly enhanced, as would the market for specialized procedures, such as user interface management

systems, model management systems [32], and so on. Existing “standards” such as PostScript, XWindows (including Motif), and Apple Events [2] can be seen as protocols for special sorts of process-to-process communication. Much the same can be said for messages in an object-oriented system.

Just as EDI protocols gain credibility from the fact that they are successfully in use, so do the existing “standards” just mentioned. But EDI—as we have seen—has its problems, and these problems similarly apply to the more general case of process-to-process communication.

5.3 Problems of Describing Information Items

Information items include everything that can be stored digitally, including documents, hyperdocuments (collections of linked documents whose links are managed by a hypertext system), models, data, still and moving pictures, and sounds. A *description* of an information item can include any type of information that a user deems to be important for understanding or retrieving the item, such as the source, the creation date, units of measurement, related passages, or related documents. All these types of items need to be described, retrieved, and reasoned about in computer applications.

The requirements of each of these processes are related though it is the description of the item that forms the basis of the problem that we are focusing on: if the item is not described effectively, then the retrieval mechanism will not retrieve the item properly or effectively. The problem of retrieval is significant, and it is known that existing systems and methods do not work as well as would be desired [5, 6, 50].

What is needed, then, is a standard way of describing information items that is powerful enough to express any statement and that is in a format that can be understood by arbitrary computer programs. These more powerful descriptions could be used to improve corporate memory [27], enable litigation to be more effectively defended against or carried out, enable software systems to be written less expensively and more efficiently [17], and generally enable companies to go about their business in a more efficient manner.

6 Solution Concept

What these three families of problems have in common is that they are all communication problems. Clearly this is so, at least in some basic sense, for EDI and for process-to-process communication. It is no less the case, we submit, for describing information items. There, the person describing the information is attempting to communicate

with whomever may be interested in that information at some future time.

Even if our three families of problems are all communication problems, there is more than one sense to the term communication, and we will not have said anything interesting about these problems—as communication problems—until we have presented a unifying and fruitful perspective, a framework, for treating them. To begin on this, it is useful to note a passage from one of the earliest papers on communication theory, by Warren Weaver [52, page 4]:

Relative to the broad subject of communication, there seem to be problems at three levels. Thus it seems reasonable to ask, serially:

Level A. How accurately can the symbols of communication be transmitted? (The technical problem.)

Level B. How precisely do the transmitted symbols convey the desired meaning? (The semantic problem.)

Level C. How effectively does the received meaning affect conduct in the desired way? (The effectiveness problem.)

“Information theory,” or the mathematical theory of communication, only deals with the technical problem. The theory is an engineering theory, built upon an engineering framework. That framework is foundational for understanding communication at all levels. Weaver describes the basic communication framework as follows [52]:


The *information source* selects a desired *message* out of a set of possible messages . . .

The *transmitter* changes this *message* into the *signal* which is actually sent over the *communication channel* from the transmitter to the *receiver*. . . .

The *receiver* is a sort of inverse transmitter, changing the transmitted signal back into a message, and handing this message on to the destination.

Rather obviously, our communication problems have to do with Weaver’s level B, and perhaps level C. What is not at issue (here) is whether a representation of, say, a purchase order can be transmitted effectively, i.e., whether the signal produced by the transmitter can be sent over the communications channel without unacceptable levels of error. Instead, the question is how to *represent* a purchase order (or, e.g., an expression for calling a software module, or a description of a document) so that once the representation is received the requisite meaning may be effectively extracted from it. In order for this to happen the representation has to carry the information in the first place, and the receiver has to have the ability to interpret the representation.

Figure 1 Schematic for Messaging with a Message Management System



With the (engineering) information-theoretic framework as background, we are now prepared to describe our solution concept. In doing so, we will continue to focus on the EDI problems, but this is without loss of generality. What we say here largely applies to the other two families of problems, for all three are communication problems; what is being discussed here is a solution concept to a general communication problem, not just the EDI problem. Also, our purpose here is to *present* the solution concept. Justification and details will have to await future papers.

Consider, now, a scenario for how an EDI message could be sent and received with an MMS, as we conceive it (figure 1). As in the case of standard EDI, the information source is a human or some procedures belonging to the organization's data processing system.

1. The information source creates a message. This message is an expression in some general-purpose FLBC, and is created by the application with the aid of software for mapping between the application and the FLBC in use (FLBC_S in figure 1).
2. The message is passed to the MMS, which conceptually belongs to the transmitter in the communication system and which is a general-purpose software module, being able to process a wide variety of messages. (These messages may perhaps be in a variety of FLBCs.) Further, the message from the information source may contain reference to information items (e.g., data, files, documents, knowledge bases) that are available to the MMS, but that are not actually part of the message itself. For example, the message may be a request to send a file. Instead of including the file in the message, instructions are included that allow the MMS to access the file. (Also, by prearrangement the MMS may have access to sender-specific information.)
3. After logging the received message, it is the job of the MMS

to take whatever actions are necessary to forward the message. This may involve translating (or mapping) the message to another format. (Thus, in Figure 1 we see that the sender-side MMS produces a message in the FLBC_T language. ‘T’ is for transmission.) For example, if the message is an offer to purchase a certain number of widgets under certain conditions, the MMS might determine that the addressee requires that such messages be sent in the format of a particular release of the X.12 transaction set for purchase orders. If so, then the MMS translates the message (plus referenced information as appropriate) into the relevant X.12 protocol expression, and the message sending story proceeds as in the case of standard EDI, above.

(For many reasons, the choice of an X.12 protocol expression for the FLBC_T message would be an undesirable restriction. It would be much better to use a more general and powerful FLBC, where practicable. Our purpose in using the X.12 example is to show that the message management concept fits with and adds value to the standard EDI protocols. Further, we note that these protocols include fields in which free text can be inserted. These fields could be used for sending expressions in a more general FLBC. This might be useful, e.g., as an interim measure while testing the usefulness of FLBC ideas.)

4. The sender-side MMS then forwards the FLBC_T message, using an appropriate telecommunications system.
5. Upon receipt of the signal (an FLBC_T expression, say an X.12 purchase order), the receiver logs the message and translates the signal into an FLBC_R expression (the FLBC for the receiver of the message). The translation may rely on general knowledge available to the MMS, as well as particular knowledge of the intended recipient.
6. The resulting FLBC_R expression is forwarded to the appropriate application module.
7. The application at the receiving end receives the FLBC_R message, maps it to the proper, application-specific form, and processes the received information for consumption by humans or their surrogates.

We have deliberately kept simple this scenario for employment of an MMS. Many variations are possible, some of which—we believe—can produce valuable features in a cost-effective and maintainable way. For the present, we will confine ourselves to pointing out a few basic advantages of the message management concept, as here articulated, on the assumption that the FLBCs are defined in a flexible and expressively powerful fashion.

1. With an MMS in place, EDI standards changes (including enhancements and expansions) can become largely transparent to particular applications. If a standard changes in a way that does not substantially increase expressive power (or requirements), then the problem of accommodating changes is the problem of revising the rules for translating between expressions in the FLBC and expressions in the changed standard. If a new standard is created or an old standard is greatly enhanced, the problem is still fairly simple. Within the MMS, the translation rules can be expanded and this would normally be straightforward. Among the particular applications serving as message sources, those for which there was a need to exploit the new standard would be changed (or created) so as to be able to send an expanded list of possible messages. (Notice, too, that this benefit may accrue in the case of module-to-module communication. For example, by placing an MMS between, say, applications and a user interface management system it would be possible to swap in a different user interface management system without undue disruption to the applications.)
2. With an MMS in place, the mapping problem for EDI can be ameliorated. Translating between expressions in an FLBC and expressions in a communications protocol is not a particularly difficult problem. What is difficult is mapping between an application process and *any* other system of formalized expression. If the FLBC can be made sufficiently expressive and flexible, however, this problem is ameliorated, for once applications are revised to accommodate messaging with an FLBC, the incremental cost of doing more such messaging is greatly reduced.
3. With an MMS in place, tracking and management of communications can be enhanced in a cost-effective manner. In principle the matter is simple. Descriptions of all significant events are recorded as FLBC expressions. Then, general-purpose software for reasoning about FLBC messages can be used to filter, search, and interpret the log files. Common, extensive use of an FLBC creates something of a network externality here. Once everyone is using FLBC expressions for logging, writing (and maintaining) the general-purpose software for extracting information from a series of FLBC expressions becomes much easier.
4. With an MMS in place, dialog, procedure, and negotiation management can more easily be systematized. What we have in mind is this. For a dialog (procedure or negotiation) to be managed there must be some model present for determining proper and permitted dialogs (steps in the procedure, moves in the negotiation). In negotiating a purchase, for example, one cannot accept

an offer until it has been made. The natural way to model a dialog (procedure, negotiation) is with a grammar of some sort, which will specify proper sequences of messages. Having the messages in a common FLBC will facilitate design and use of message grammars. It is not difficult to imagine how a negotiating position might be stated by placing a preference ordering on permitted responses in a negotiation or dialog grammar.

5. With an MMS in place, particular EDI applications may be more easily generalized to work with other business partners. By centralizing the locus of change it may be hoped that the cost of adding new communicants can be greatly reduced. Further, the use of an FLBC (a *formal* language for business communication) allows us to add creative procedures for transforming the messages, without having to modify the applications that created them. Ron Lee has provided an impressive illustration of this with his multilingual “CASE-EDI” tool [37]. Because his system relies on a formal language for communication among business partners, and because text generation technology is fairly mature, Lee’s system can automatically and transparently translate messages from one language to another. This is a natural exploitation of the sort of message management architecture we have been describing.

7 Summary: Message Management

Having sketched our concept of an MMS and how it might work, we return now to the five plausibility conditions laid out in §4.

1. *The basic concept of the message management system must be articulated in a clear and operationalizable manner.*

Briefly, our concept is that an MMS is a software system for handling module-to-module communication, for EDI and for other purposes as well. Modules communicate with the MMS using an FLBC having great flexibility and expressive power. Further, messages in the FLBC may indicate other information items (e.g., files and knowledge bases) that the recipient of the message (either the MMS or a module) may use in processing a message. Further, the MMS may exploit its own knowledge resources in performing its functions. The FLBC should be powerful enough to be used as a general indexing language for information items, particularly documents.

2. *The functionality that is to be systematized must be identified and shown to be sufficiently general to make the systematization*

plausibly worth the cost and effort. The benefits of successful systematization should be identified.

In §5 we presented three families of communication problems for which we believe a well-constructed MMS is at least a partial solution. Much of our discussion has focused on EDI, where the case for message management is perhaps clearest. We explicitly discussed the high-level benefits for EDI in §6. Further, we believe that the existence of a well-constructed MMS will create a situation with positive network externalities, since the incremental cost will be low of extending such a system to handle new functionality. Once such a system is in place and works well for some aspects of EDI, there will be compelling reasons (low costs for high benefits) to use it for all of an organization's EDI. Once a filtering and retrieval module for messages in the FLBC is in place in the MMS, it just makes plain sense to use the apparatus for indexing and retrieval of information items generally. Finally, once the MMS is working well for EDI generally, extending it for use in module-to-module communication is only natural, especially when the communication environment needs to be carefully managed (e.g., for transaction monitoring) or is highly fluid (e.g., when software is to be ported and modules must be made to work with similar but different communicants).

3. *A collection of general requirements for such systems must be identified.*

We have not said much directly and in detail on this subject. Briefly, we see two sorts of requirements: those for the FLBC and those for the MMS proper. The FLBC must be expressively powerful and gracefully open to extension. These requirements are rigorous, and in our view essentially mandate use of a recursively-defined formal language. For the sake of expressive power, such an FLBC would have to employ quantification, various sentential operators (boolean operators certainly but also modal operators, temporal operators, and deontic operators), propositional attitude operators (e.g., those required by speech act theory, such as asserts, promises, requests, questions, and declares [35, 33]), and defeasible assertion and implication [31].

The requirements for a general MMS are also extensive. Among other things, it must perform mapping between FLBC messages and an array of specialized protocols. In general, it must be able to use a knowledge base of FLBC messages in performing various filtering, retrieval, and inferential tasks, including constructing error messages, log messages, acknowledgments, and (for negotiations, procedures, and dialogs) determine the next message in a sequence of events and messages.

4. *A general approach to designing, implementing, and maintaining such systems should be articulated.*

We have little to say about this condition, partly because little is known about the problem and partly because quite a lot is known. On the little side, we note that not much experience has yet been had in building and testing MMSs (but see condition 5, below). On the a lot side, we assume that declarative programming and knowledge representation techniques, presumably using an embedded languages approach (see [4]), will be used to construct prototype MMSs. Quite a lot is known about these programming techniques, and it is reasonable to assume that with iterative development and testing the proper architecture will emerge.

5. *The computational costs associated with the proposed design approach should be shown to be acceptable.*

It is much too early to say much with confidence regarding the computational practicability of this approach. Nevertheless, some experience has been had and it has been positive and propitious. Lee's system [36, 37, 38] uses an FLBC for messaging related to international purchasing. Kimbrough has experimented with an FLBC for office communications [35, 33]. The DSS shell Max uses an FLBC for internal module-to-module communication [32]. Scott Moore has implemented document retrieval prototypes relying on document indexing with expressions in an FLBC. The PUP6 prototype system (for interactive concept formation) uses a simple formal language for messaging. The language is able to describe agents in a standard way [40]. Finally, the Contract Net system used a specialized formal language, called the Contract Net Protocol, for communication among multiple agents attempting to allocate tasks jointly [48].³

In sum, the concept of an MMS is a reasonably definite idea. It is motivated by a series of important problems and it would appear plausible that a good implementation would contribute significantly to solving these problems. Finally, there is some evidence that the concept can be made to work in practical settings. Thus, we find the prospects exciting and hope others will also be motivated to explore them.

8 Discussion

Recall the aims of this paper. First, we want to the issues related to the “common communications conventions” required for AAA communications—

³Both PUP6 and Contract Net are much discussed in the distributed AI literature. For an introduction to that literature, see [7, 22, 23, 25].

this is our technical point. In this section we work on the second aim: to put forth and justify some hypotheses on the strategic and industrial-organizational (IO) consequences of AAA communications made possible by widespread use of MMSs.

We have attempted to build a plausibility case for the claims that broad, robust communication capabilities are needed and that systems able to provide these services are feasible. In this section we assume the existence of these systems and explore their implications.

8.1 A Move to the Middle

As stated in §1, we are broadly in agreement with the hypothesis that investment in information technology can lead to more cooperative relationships among firms (what Clemons and associates call a “move to the middle” [11, 13]). They hypothesize that information technology encourages and enables this cooperation by reducing the costs of integration and the transaction risk. If this is the case (and we agree with their hypothesis), then communications systems based on the concepts presented in this paper would strengthen this tendency. Why is this so? Consider the decomposition of transactions cost as presented in [11]:

$$\text{transactions cost} = \text{coordination cost} + \text{operations risk} + \text{opportunism risk}$$

The costs of each of these components would be reduced by using the concepts presented in this paper:

- *Coordination cost*: “the costs incurred by the firm in coordinating with the unit producing the product” [11]. Two components of this cost are the cost of exchanging information and the cost of delays in the communication channel. The cost of exchanging information is reduced because an MMS/FLBC communication system allows a broader range of information to be expressed electronically. Since computers are more efficient and less expensive than humans in handling information, the total cost of exchanging information would be reduced. The reduction of cost of delays in the communication channel also follows from the expanded scope of the electronic communication channel.
- *Operations risk*: risk arising from the process of generating output, including risk concerning product quality, required delivery times, output quantity and capacity, and available flexibility. The only information technology-related component of operations risk is the cost of the mechanism that coordinates the daily operations of firms (both within one firm and among several firms). Again, the expanded scope of allowable communications with an FLBC

allows more types of activities to be controlled and monitored electronically. This reduces the amount of human intervention required thereby reducing the cost of operations risk. Also, with a MMS/FLBC communication system, it does not matter technically whether the operations are within or without the firm. The other components of this risk would be unaffected by this technology.

- *Opportunism risk*: the risk of being open to opportunistic behavior of business partners. One of the components of this cost is the amount of idiosyncratic investment required for coordination. Current EDI standards only allow for a minimum of coordination. At the same time the lack of standards in industry forces firms to make significant investments in communication technology before they can communicate. On the other hand, investment in developing an MMS and an appropriate FLBC is designed to be highly transferable among firms and among industries, reducing the opportunity for exploitation and hence the cost of opportunism risk.

Thus, we have seen that if the “move to the middle” hypothesis is correct, then we should expect to see an even greater increase in cooperation if communication systems based on our concepts are implemented.

8.2 Symbiotic Relationships

We have two additional hypotheses. The first relates to the nature of business relationships: these advanced communication systems will often change the nature and dynamics of switching costs of all contracting parties, leading to much more symbiotic economic relationships that extend and evolve over time.⁴ Clemons and Row hypothesize that increased levels of this information technology-supported cooperation might lead to more vertical quasi-integration, outsourcing, cooperation for economies of scope, and cooperation for economies of scale [13]. We agree with this and can further explicate reasons for both the supplying firm and the host company to form highly cooperative relationships: 1) The supplying firm desires to exploit enhanced communication systems to provide differentiated service in order to build tight bonds with all the host companies with which it does business. The supplying firm will attempt to leverage the system capabilities to make it less attractive for a host company to switch suppliers even though the systems themselves make it technically easier to do so. 2) The host companies desire to gain greater efficiencies by taking advantage of both the ex-

⁴Thanks to John Holland for provocative and stimulating remarks, made in a seminar, on the origin and evolution of symbiotic relationships in complex adaptive systems.

panded bonds with its suppliers and the specialized knowledge of the supplying firms.

The MMS and FLBC concepts encourage and enable change by providing a flexible and expressive method of communication. Companies may no longer have to change their applications to inter-operate with a different firm's applications. Each firm communicates through its own MMS, which handles the process of conversion to the appropriate language. The skills and knowledge gained in adapting one's applications and processes to work with the MMS of one company are largely transferrable to making them work with another company. The expressive power of the FLBC allows more expressions to be included in a well-defined protocol (i.e., the FLBC) and not written in an *ad hoc* adaptation of an existing protocol.⁵ This disciplined expansion of the vocabulary makes it much easier to describe the system to a new user and integrate the functionality of a system with any new system. Both sides of the economic transaction have this increased ability to change and the knowledge that the other side has this increased ability to change. This decreases the switching costs for both firms, thus lowering the opportunism risk of the transaction. This sort of technological improvement is critical to the continued growth of electronic communications.

We have now justified our claim that the MMS design makes it easier for firms to switch suppliers. As stated above, we do not think this will lead to a market in which services and goods merely are purchased as commodities. Why is this? An MMS/FLBC communication system enables evolutionary changes in a business relationship that will allow close and extensive bonds to be built between a company and its suppliers. We hypothesize the benefits from developing, sustaining, and extending these bonds will increase a company's efficiency and profitability more than the benefits from bidding for the same services on an open market. Two claims need to be justified here. First, we support the claim that an MMS/FLBC system enables evolutionary changes in a relationship. Standards changes are limited to the MMS since other applications communicate with the MMS in its language. The MMS is the only application that communicates with the external environment. Further, an FLBC has been designed to allow many kinds of statements to be added to the language without affecting other parts of the language. Each of these enables firms to extend the scope and scale of a relationship. We now support our claim that the benefits from these kinds of relationships will often be extensive and profitable.

The "move to the middle" will create many firms specializing in providing a service or good. Not only will more firms be capable

⁵We make these points by way of adding to the arguments offered by Clemons et al. [11, 13]

of managing relationships with these firms, more kinds of products will be secured through these cooperative relationships because the increased communication capabilities will allow monitoring and management of more complex relationships than previously possible. We say these are “cooperative” relationships—what do we mean by this? These relationships might resemble those between consulting firms and their clients. The difference between the familiar company–consulting firm relationship and the company–specialized firm relationship is that much of the coordination between the firms will be enabled by electronic communication systems of the future. The monitoring and management of the cooperating processes will be managed by computers, only involving humans when difficult or unforeseen circumstances arise. In this way a firm’s employees can spend their time solving difficult problems instead of shuffling papers and playing “telephone tag.” All information acts will be facilitated, bringing about greater efficiency.

We have attempted to make it plausible that firms will look to enter into these cooperative relationships with many suppliers. Each firm also knows that the technology empowers them to leave the relationship. This provides firms incentive to create other ways of strengthening their relationships. One way to do this is for each firm to learn more about the other firm’s operations and use this knowledge to tightly integrate the computer systems of the firms. An MMS/FLBC communication system enables many parts of the organizations to inter-operate and to grow more tightly integrated over time as each gradually learns more about the other. This tight integration with a supplier for a particular service will naturally lead to a reduction in suppliers to a firm for that service. Problems with coordinating and educating the people involved in providing and receiving the service will encourage firms to focus their energies on a particular supplier. Also, supplying firms will be more willing to commit assets to a relationship if it is known that a competing firm is not benefitting from these assets.

Unresolved though intriguing questions are:

- What industries and which business activities will see the formation of the kind of cooperation we have described?
- What are the industrial-organizational and competitive structure of these industries?

Effects similar to those described above are sometimes seen in firms that integrate EDI into their businesses and customer relationships. In [9, p. 62] the following observations were made about the banking industry and EDI:

- When a bank and its customer work together to integrate their EDI systems, “[t]he two organizations become more tightly interwoven.” This supports our claim that communication systems can be used to strengthen customer relationships.

- Banks are able to specialize and provide services for customers around the country that previously were only provided by local banks. This supports our claim that communication systems can support the creation of specialization where little existed.

This is not conclusive evidence but it illustrates and lends credence to some of our claims relative to the symbiotic relationship hypothesis.

8.3 Barriers to Entry

Our second hypothesis that we explore in this discussion, much less involved than the first, is the following: because substantial investment will be required to convert existing systems, first-mover opportunities and barriers to entry dependent upon the conversion costs will be created. Companies with communication systems in place have to convert the applications that use these systems in order to use them with an MMS/FLBC system. The conversion of this system may be facilitated by the corporate knowledge gained during the installation of previous communication systems. However, new companies and companies without systems to convert will have a cost advantage over firms that must pay conversion costs. This barrier to entry will provide these first-movers with some advantages:

- The first-mover will be able to begin moving up the learning curve (i.e., will gain economics of learning). The knowledge gained can be used to apply the new technology in unique ways.
- The first-mover may be able to offer new services other companies will be unable to offer. For these services the first-mover may be able to charge monopolistic rates or gain large parts of a market before other firms overcome the barrier to entry.

Benefits such as these have been reported of Rosenbluth Travel [12, 14] with current communication technology. As a new firm Rosenbluth applied information technology in a unique and advantageous way. Other travel agencies had already invested heavily in older technology and were not in a position to update their investment. Rosenbluth profited from this first application of technology and applied these profits to developing more advanced applications of technology that have still not been adapted by its competitors. Rosenbluth gained first-mover advantages through intelligent application of current communication technology. We hypothesize that these same types of benefits could be available for firms applying FLBC/MMS communication systems.

9 Summary and Conclusion

This is quite a long paper and much ground has been covered. Let us review what we have seen. We first tackled our *technical point*. Workflow automation and EDI systems were shown to be limited and imperfect approaches to the systematic, efficient handling of information events. We then attempted to make plausible our claim that a message management system using a formal language for business communication removes many of the imperfections. We also provided hints that indicate that an MMS/FLBC system is both technically and economically feasible.

The next part of the paper assumed the conclusions of the first part—namely, that if a system *can* and *should* be built, then it *will* be built. This part of the paper highlighted three possible consequences of such a move. The first is the hypothesis of Clemons et al. that says that an improvement in communication technology (in this case, MMS/FLBC) will bring about many and more cooperative outsourcing arrangements between hosts and vendors. Second, we claim the MMS/FLBC combination will enable these cooperative outsourcing arrangements to extend and evolve over time, benefitting both host and vendor. Finally, firms without an installed base to convert and firms easily able to bear the cost of converting to an MMS/FLBC system will, for a time, have a competitive advantage over other firms.

References

- [1] Allen, Linda G. (ed.), “Small-Town Company talks Big Savings, Better Service With EDI,” *Automation*, 38:7 (July 1991), p. 56.
- [2] Apple Computer, Inc., *Inside Macintosh, Volume VI*, Addison-Wesley Publishing Company, Reading, MA, 1991.
- [3] Baker, Carol, “EDI in Business,” *Accountancy*, 107:1172 (April 1991), pp. 121–124.
- [4] Bhargava, Hemant K., and Steven O. Kimbrough, “Model Management: An Embedded Languages Approach,” forthcoming in *Decision Support Systems*.
- [5] Blair, David C., *Language and Representation in Information Retrieval*, Elsevier Science Publishers, New York, New York, 1990.
- [6] Blair, David C., and M.E. Maron, “An evaluation of retrieval effectiveness for a full-text document-retrieval system,” *Communications of the ACM*, **28**, no. 3, pp. 289–299.
- [7] Bond, Alan H. and L. Gasser, eds., *Readings in Distributed Artificial Intelligence*, Morgan Kaufmann, San Mateo, CA, 1988.

- [8] Booker, Ellis, "Work-flow system market grows," *Computer-World*, January 28, 1991, p. 29.
- [9] Canright, Collin, "Customization Comes to EDI," *Bank Management*, 67:10 (October 1991), pp. 59–62.
- [10] Classe, Alison, "Rover races ahead," *Accountancy*, 107:1172 (April 1991), pp. 122–123.
- [11] Clemons, Eric K. and Sashidhar P. Reddi, "A Move to the Middle: Information Technology and Outsourcing, Information Technology and Governance Structures," University of Pennsylvania, Decision Sciences Department, working paper 92–03–05, March 1992.
- [12] Clemons, Eric K. and Michael C. Row, "Information Technology at Rosenbluth Travel: Competitive Advantage in a Rapidly Growing Global Service Company," *Journal of Management Information Systems*, 8:2 (Fall 1991), pp. 53–79.
- [13] Clemons, Eric K. and Michael C. Row, "Information Technology and Industrial Cooperation," this issue, *Journal of Management Information Systems*.
- [14] Clemons, Eric K., Michael C. Row, and David B. Miller, "Rosenbluth International Alliance: Information Technology and the Global Virtual Corporation," *Proceedings of the Twenty-Fourth Hawaii International Conference on System Sciences, Volume IV*, IEEE Computer Press, pp. 678–686.
- [15] Clickunbroomer, Jeanette, "As Essential As A Telephone," *Global Trade*, 111:9 (September 1991), pp. 41, 50.
- [16] Dertouzos, Michael L., "Communications, Computers and Networks," *Scientific American*, **265**, no. 3 (September 1991), pp. 62–9.
- [17] Devanbu, Premkumar, Ronald J. Brachman, Peter G. Selfridge, and Bruce W. Ballard, "LaSSIE: A Knowledge-Based Software Information System," *Communications of the ACM*, **34**, no. 5, pp. 34–49.
- [18] Eckerson, Wayne, "Pioneering Users Moving to Faster Methods of EDI," *Network World*, **8**, no. 18, 161 Worcester Road, Framingham, MA 01701-9172, (508) 875-6400, p. 1.
- [19] Edelstein, Herbert A., "Imaging shifts emphasis to workflow management," *Software Magazine*, November 1991, pp. 96–105.
- [20] Egol, Les, "The "Less Paper" Office," *Chemical Engineering*, March 1991, pp. 199–204.
- [21] Feder, Barnaby J., "Moving the Pampers Faster Cuts Everyone's Costs," *New York Times*, business section, July 14, 1991, p. 5.

- [22] Gasser, Les, "Social Conceptions of Knowledge and Action: DAI Foundations and Open Systems Semantics," *Artificial Intelligence*, **47**, nos. 1-3 (1991), pp. 107–138.
- [23] Gasser, Les and Michael N. Huhns, eds., *Distributed Artificial Intelligence, Volume II*, Morgan Kaufmann Publishers, Inc., San Mateo, California, 1989.
- [24] Harris, Derek, "Fast orders help traders net profits," *Accountancy*, November 1991, pp. 104–105.
- [25] Hewitt, Carl, "Open Information Systems Semantics for Distributed Artificial Intelligence," *Artificial Intelligence*, **47**, nos. 1-3 (1991), pp. 79–106.
- [26] Huhns, Michael N., ed., *Distributed Artificial Intelligence*, Morgan Kaufmann Publishers, Inc., Los Altos, CA, 1987.
- [27] Huhns, Michael N., Uttam Mukhopadhyay, Larry M. Stephens, and Ronald D. Bonnell, "DAI for Document Retrieval: The MINDS Project," in [23], pp. 249–283.
- [28] Johnson, William, Jr., "Anything, Anytime, Anywhere: The Future of Networking," in [39, pages 150–175].
- [29] Jordahl, Gregory, "Paperless Tiger?," *Insurance Review*, March 1991, pp. 20–26.
- [30] Kimberley, Paul, *Electronic Data Interchange*, McGraw-Hill, Inc., New York, New York, 1991.
- [31] Kimbrough, Steven O. and Hua Hua, "On Modeling Nonmonotonic Reasoning with the Method of Sweeping Presumptions," forthcoming in *Minds and Machines*, (November 1991).
- [32] Kimbrough, Steven O., Clark W. Pritchett, Michael P. Bieber, and Hemant K. Bhargava, "The Coast Guard's KSS Project," *Interfaces*, **20**, no. 6 (November-December 1990), pp. 5–16.
- [33] Kimbrough, Steven O., Sashidhar Reddi, and Michael B. Thornburg, "On Messaging with Semantic Access in an Office Environment," University of Pennsylvania, Decision Sciences Department, working paper 89-09-06, December 1989.
- [34] Kimbrough, Steven O. and Cynthia A. Sherman, "On EDI and on beyond EDI," presentation at the Reginald H. Jones Center Information Systems, Telecommunications, and Business Strategy Project, the Wharton School, University of Pennsylvania, July 10, 1991.
- [35] Kimbrough, Steven O. and Michael J. Thornburg, "On Semantically-Accessible Messaging in an Office Environment," *Proceedings of the Twenty-Second Hawaii International Conference on System Sciences*, 1989.

- [36] Lee, R.M. and S.D. Dewitz, "Finding International Contracting Opportunities," *Proceedings of the Eleventh International Conference on Information Systems*, Dec. 90, pp. 1-13.
- [37] Lee, R.M., S.D. Dewitz, and K.T. Chen, "AI and Global EDI," *Proceedings of the Twenty-Fourth Hawaii International Conference on System Sciences, Volume IV*, IEEE Computer Press, pp. 182-191.
- [38] Lee, R.M. and S.D. Dewitz, "Facilitating International Contracting: AI Extensions to EDI," forthcoming, *International Information Systems*.
- [39] Leebaert, Derek, ed., *Technology 2001: The Future of Computing and Communications*, The MIT Press, Cambridge, MA, 1991.
- [40] Lenat, Douglas B., "BEINGS: Knowledge as Interacting Experts," in *Proceedings of the 1975 International Joint Conference on Artificial Intelligence*, (1975), pp. 126-133; reprinted in [7].
- [41] May, Thornton A., "Is Imaging for Real?," *ARMA Records Management Quarterly*, July 1991, pp. 3-8.
- [42] Nash, Jim, "Workman extends into work-flow management," *ComputerWorld*, December 2, 1991, p. 57.
- [43] Pollock, Andrew, "A Quirky Loner Goes Mainstream," *New York Times*, business section, July 14, 1991, p. 1.
- [44] Radding, Alan, "A Tale of Two Banks," *Bank Management*, November 1991, pp. 16-17.
- [45] Russ, Scott and Edward T. Hatchett, "Extending Enterprise Automation (H-5)," presentation at DISA Conference, Orlando, Florida, April 24-26, 1991.
- [46] Salamone, Salvatore, "EDI: Bottom-line booster or budget-breaker?," *Network World*, 8:13 (April 1, 1991), pp. 1ff.
- [47] Shannon, Claude E., and Warren Weaver, *The Mathematical Theory of Communication*, The University of Illinois Press, Urbana, Illinois, 1949.
- [48] Smith, Reid G., "The Contract Net Protocol: High-Level Communication and Control is a Distributed Problem Solver," *IEEE Transactions on Computers*, C-29, no. 12 (December 1980), pp. 1104-1113; reprinted in [7].
- [49] Smith, Ted, "Images of the Future," *Bank Management*, November 1991, pp. 12-19.
- [50] Swanson, Don, "Historical Note: Information Retrieval and the Future of an Illusion," *Journal of the American Society for Information Science*, **39**, no. 2, pp. 92-98.

- [51] Wallace, Bob, “Imaging net helps insurer cut red tape,” *Network World*, April 15, 1991, pp. 23–24.
- [52] Weaver, Warren, “Recent Contributions to the Mathematical Theory of Communication,” in [47, pages 1–28].
- [53] Weiss, Peter, “Electronic Data Interchange: Active Ingredient of Electronic Commerce,” from *A Five Year Plan for Meeting the Automatic Data Processing and Telecommunications Needs of the Federal Government*, (November 1990), available from the Government Printing Office.