Tying Knowledge to Action with kMail

David G. Schwartz and Dov Te'eni, Bar-Ilan University

THE KMAIL SYSTEM INTEGRATES E-MAIL WITH ORGANIZATIONAL MEMORIES TO DELIVER KNOWLEDGE ACROSS THE INTERNET IN A TIMELY, RELEVANT MANNER. KMAIL OPERATES BY PROVIDING CONTEXT THROUGH METAKNOWLEDGE-BASED MEMORY–CONCEPT ASSOCIATIONS AND BY DETERMINING THE APPROPRIATE MEMORY ITEMS THROUGH THE CREATION OF OM VIEWS.

YOU'VE GOT THIS MASSIVE, ALL-encompassing organizational memory. Now what? To make effective use of the OM's knowledge, you've got to tie it to action. But tying knowledge to action can only occur if you are aware of the knowledge and can identify it at the time of action and if the system can deliver it to that point of action. For that, you need tools.

Knowledge-enhanced e-mail, kMail for short, is just such a tool. This new class of e-mail application intertwines e-mail with knowledge management, thus representing a fundamental shift in the way we can distribute and maintain organizational knowledge.

In this article, we emphasize knowledge distribution and tie it into organizational action. Knowledge links up with action by a process of contextualization. To make these ideas work in practice, we rely on the Internet and e-mail as a transport layer for knowledge dissemination. We also describe the core components and techniques that help decide when to contextualize messages and with what knowledge.

Knowledge management: beyond time and space

> The knowledge of man is as the waters, some descending from above, and some springing from beneath.
> —Francis Bacon (1561–1626)

Knowledge management in general, and OMs in particular, are an attempt by the organization to transcend time and space in learning. In building knowledge-management systems (KMSs), we attempt to assimilate knowledge available within the organization and disseminate it to people connected to the organization.

Challenges associated with knowledge management fall into three general categories: acquisition, organization, and distribution.

Knowledge acquisition deals with the issues that surround knowledge extraction in its various forms—from the organization's knowledge bases, databases, printed resources, and people. In knowledge acquisition, we must deal with questions such as "Who is the authority on a specific area?" "How is knowledge in that area currently stored?" and "How can I get my hands on the knowledge and make it machine readable?"

Knowledge organization deals with the issues surrounding how to best store knowledge so that it can be retrieved when relevant. In knowledge organization, we must deal with questions such as "In what forms should we store this knowledge?" "How should we index this knowledge?" and "How is the user going to ask for this knowledge?"

Knowledge distribution, or dissemination, must tackle the problem of getting the right knowledge to the right place at the right time. It requires three conditions:

- *Awareness*: The user must be aware that there is relevant knowledge available. Does the manager know that a best-practices database exists? Does he or she know that there may be something in there to help resolve the current crisis?

1094-7167/00/$10.00 © 2000 IEEE

- *Identification*: The user must be able to readily identify that knowledge. Can the manager effectively use the best-practices database to find something that will help him or her?
- *Delivery*: The knowledge must be delivered to the point of need, in a timely manner.

Managers in an organization do not have the time or inclination to actively seek organizational knowledge. It would be far more effective if the knowledge could find them.

Consider the use of software help systems. No one reads through the help files of a software system. The biggest advance in help systems of the past decade has been context-sensitive help, which is successful because it ties knowledge to action. The use of organizational knowledge must go through the same transition.

We are enabling knowledge distribution for action. The development of kMail is a direct response to the need to tie knowledge to action, after considering the awareness, identification, and delivery aspects of knowledge distribution.

Moving knowledge management to the Internet. There are three major classes of application dominating the Internet today, all of which are relevant to knowledge-enhanced e-mail:

- Web sites that store inordinate amounts of occasionally useful information;
- Pop–client or browser-based e-mail services; and
- Web sites that provide extensive indexing, cataloging, and search services.

Put the three of them together in an organizational setting and a transformation occurs:

- The Web site becomes an HTML-based knowledge repository that can be viewed through an organizational lens.
- The e-mail service becomes a transparent interface to manage knowledge distribution.
- The indexing and search incorporates organizational metaknowledge and user profiles.

The end result is a kMail system.

The kMail paradigm grew from a research project we undertook to investigate ways to improve intercultural communication in multinational corporations.[1,2] That research, resulting in the Hypermail system, showed that there is a well-defined set of threats to communication. Our investigations found that there are two fundamental influences on the effectiveness of electronic communication: cultural and organizational norms and sender–receiver distance, the latter depending on personal background, experience, company role, and other context-forming criteria.

The starting point of the communication process is the goals of organizational communication. Based on Jurgen Habermas's work,[3] the categories of such goals would be

1. Commanding a specific action;
2. Managing collective action: thinking collectively; monitoring communication, command, and control; and setting work procedures and rules;
3. Influencing (persuasion, leadership, lobbying, propositions);
4. Providing information for future action (knowledge dissemination); and
5. Seeking information for future action (knowledge acquisition);

MANAGERS IN AN ORGANIZATION DO NOT HAVE THE TIME OR INCLINATION TO ACTIVELY SEEK ORGANIZATIONAL KNOWLEDGE. IT WOULD BE FAR MORE EFFECTIVE IF THE KNOWLEDGE COULD FIND THEM.

We can make each of these action-oriented goals more attainable by tighter integration of knowledge.

Why e-mail? When an organization's managers communicate by e-mail, in most cases, they are trying to achieve some action or responding to a request for some action. Those actions' effectiveness depends largely on the correct understanding of what the best action should be. This, in turn, depends on knowing the proper context in which that action is requested. The five goals of communication we've listed are equally valid for electronic communication.

Daniel O'Leary[4] discusses three significant limitations to knowledge management initiatives that focus on knowledge as an end in itself.

- Knowledge does not necessarily result in action to create value: organizational knowledge-management initiatives must be action-oriented and create value if they are to acquire the management priority and heavy resources needed to fund their development.
- Knowledge processes are dynamic: KMSs all suffer from one significant barrier—creating the knowledge and keeping it current.
- Knowledge cannot be a surrogate for creativity: knowledge should inform and guide creativity. For this to happen, knowledge must be available during the creative processes followed in day-to-day business activities.

The question we posed to ourselves was "How can we best keep knowledge dynamic, use it in action-oriented situations, and make it the backdrop for creativity?" The answer is through e-mail, the quintessential Internet application. Consider the following:

- Every organization, without exception, will have an e-mail infrastructure before it reaches the stage of developing an OM.
- E-mail communication in a modern organization is over 78% action-oriented, according to a recent study.[2] Organizations must converge to action, and communication is perhaps the foundation for most organizational action.[5]
- Managers, and knowledge workers of all kinds, interact with their e-mail systems on a daily basis—it is a standard operating procedure. This means that using e-mail as the window into an OM gives us the smallest delta for change in an organization's daily activities.
- Managers are motivated to achieve successful communication. They want their instructions understood and their answers to queries to be effective.

There are many ways to connect people to knowledge through the use of intelligent agents, push technology, or search engines. But connecting people with knowledge should be transparent. A knowledge-management tool is nearly worthless if managers must learn specialized techniques or complex interfaces to get at that knowledge. Best-

practices database usage for every new case is a far cry from the ubiquitous use of organizational knowledge that businesses are looking for.

Contextualization and organizational knowledge. To use knowledge in an action-oriented manner, the e-mail messages must be able to access that knowledge in its correct context. Contextualization must rely on organizational knowledge for two components: knowledge to provide the additional context layers around action and knowledge to identify the conditions in which to contextualize messages. We assume here that organizational knowledge is explicit and available, which is extremely naive and outright wrong in many cases. However, new technologies such as expert support systems, intranets, and document (including e-mail document) management are making it feasible to store knowledge in an accessible fashion. Our starting point, therefore, assumes an existing URL-based OM, which could enable linking the knowledge to ongoing communication.

Core components

Retrieving knowledge in context is not a trivial task. To do so, kMail combines core components—message parsing to identify concepts in an e-mail message; search and indexing to create the metaknowledge overview of the OM; and profile matching and linking concepts—to memory items in the OM to minimize sender–receiver distance.

Accessing the OM through metaknowledge. Metaknowledge plays an essential role in connecting e-mail with OMs. The identification of the sender and recipient of an e-mail message is of paramount importance. Knowing this lets us access a wealth of user-specific metaknowledge for use in focusing access to the OM. Of course, both the sender and recipient identifications are readily available in an e-mail application.

The metaknowledge kMail uses consists of two main components—user profile information and explicit memory–concept association (MCA) information. Both reside in a highly structured relational database. The (formal) metaknowledge serves as the link between the (informal) e-mail communications and the (semiformal) HTML-based organizational knowledge base. Figure 1 shows a metalevel architecture.[1,6]

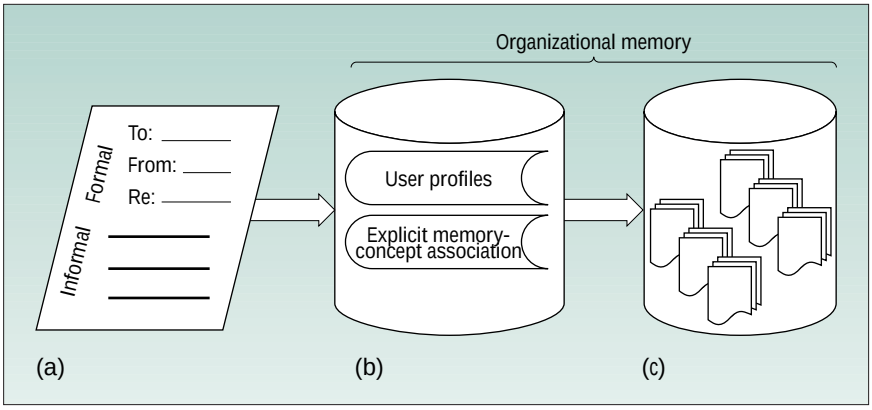

Figure 1. Three-tier architecture tying e-mail to organizational memory through metaknowledge: (a) application level—with concepts; (b) metalevel—formal metaknowledge; and (c) object level—semi-formal memory items.

Structuring the metaknowledge with MCAs. The OM metaknowledge relies on the use of MCAs. An MCA exists when there is either an explicit or derived association between a given concept and a memory item in the OM. A concept can exist without being associated with a memory item, and a memory item can exist without being associated with a concept.

For example, we may associate the concept "chocolate" with any number of memory items: chocolate cookies, chocolate milk, or a chocolate color swatch. These associations range in strength from user-specific, to project-specific, to department-specific, and on to organization-specific associations. Determining the correct association between general concepts and specific memory items is a function of OM views.

Contextualization through metaknowledge. Mark Ackerman[7] and others have addressed the use of contextualization in organizational memories. Ackerman refers to context information about the specifics of the sender–receiver attributes as well as the task characteristics prevailing at the time of the knowledge generation. We call this type of context information *situational context*. He argues that contextual information must often be dropped when building an OM to generalize the information. In establishing the "correct" level of context, we would argue that although context can be removed from OMs, it should not be discarded. On the contrary, by distilling contextual information and retaining it in the form of metaknowledge, we can achieve Ackerman's goal of generalizable OMs, without losing our ability to recontextualize that knowledge when the application so demands. kMail is just such an application. Having removed context to add longevity to the knowledge, we must then recontextualize that knowledge to add meaning to communications and tie knowledge to action.

Examples of situational context regarding a given item in the OM include attributes such as

- time or date of creation,
- time or date of modification,
- name of author,
- title of author,
- current project of author,
- number of years author is with the company,
- number of years author is in current position,
- department of author, and
- department details (country or mandate, for example).

Thus many context-free memory items can exist in the OM, only to have the context reintroduced when a user needs to link a memory item with a specific call to action. This approach is consistent with that of Ackerman and others who contend that if we want knowledge to be a useful long-term resource, we should distill it into a context-free form. The context itself is a part of the metaknowledge, and we reintroduce it to the OM when we use that knowledge to perform an action.

OM views

Beauty might be in the eye of the beholder, but meaning is in the eye of the e-mail author. If we want to properly bring an OM to bear on the concepts used in a message, we must focus on the author's intention to use those concepts. This could only be possible if we view the OM through the author's eyes. Although the degree of context required depends on the e-mail recipient,[8] it is the author who must determine the context that the message provides, as well as its depth. There is a limit to the recipient's input, at

this stage, for controlling whether to access the contextual knowledge or just let it slide by transparently. In this respect, kMail begins to address Buckingham Shum's context paradox,[8] with different views provided to different recipients and with the embedded links' nonintrusive nature letting the information overload regulate itself.

An OM view is similar in concept to a database view.[9] Where a database view takes a database structure and produces a logically filtered view suited to a given query, an OM view takes an OM's logical content and produces a filtered view suited to a given user in a certain situation. In both cases, the underlying database or OM does not change—just the way we see it. As with a database view, an OM view can be transient and has no physical existence.

Creating views on the fly. To implement kMail, each time a user authors an e-mail message, the system creates a new OM view. The view consists of OM concepts relevant to the concepts from the current e-mail, within the context of this user's activities. The system first presents the author with this view for the purpose of confirmation, validation, or modification and then sends the validated view to the kMail recipient along with the kMail message.

In fact, there is no significant data transmission required on top of a standard e-mail message. Segments of the OM view link up with the selected concepts in the kMail message. The OM itself stays in place. Only the relevant portion of the view travels across the Internet to the kMail recipient.

OM views for external consumption. Users can also tailor OM views in kMail for external consumption. An increasingly attractive scenario for the use of KMSs involves opening up part of that knowledge to customers or strategic partners. This raises a host of questions regarding access and security. OM views facilitate this type of activity. The sender can administer a secure OM-view creation mechanism in much the same way as a database scheme to determine what views he or she can make available to the outside world.

Creating OM views to minimize sender–receiver distance. To create an OM view and minimize sender–receiver distance, we perform a series of queries on the OM metaknowledge database. These queries result in forming the initial view presented to the sender.

Four views form the basis of the initial analysis that the kMail server (Figure 2) performs on the OM, each view being tested by a different query on the OM metaknowledge. The resulting user view concatenates and combines these queries. Alternatively, the system can present each view separately, as with the departmental view in our example below.

A series of tables represent the user profiles and explicit MCAs in the metaknowledge. The key metaknowledge that we must eventually arrive at consists of the entries in the memories table that best match the memory item's intended use. The memories table, however, is a context-free representation of the OM. It contains no indication of who has used this memory item and for what purpose. Its purpose is to provide an unbiased pointer into the collective OM that can then be mapped for appropriate use.

The memories table consists of four fields:

- MemoryID: a unique identifier automatically assigned to each new OM entry,
- MemoryURL: the URL where the memory item can be found,
- MemoryDescription: a brief description of the memory item, and
- MemoryType: a category code that indicates the type of contents in the memory item.

MemoryType is a significant piece of metaknowledge in that it lets us rank and sort OM entries based on intended usage. Table 1 lists MemoryType codes and their corresponding meanings.

Other metaknowledge tables required to implement the situational context attributes discussed earlier include

- People(PersonID, E-mail, Surname, Firstname),

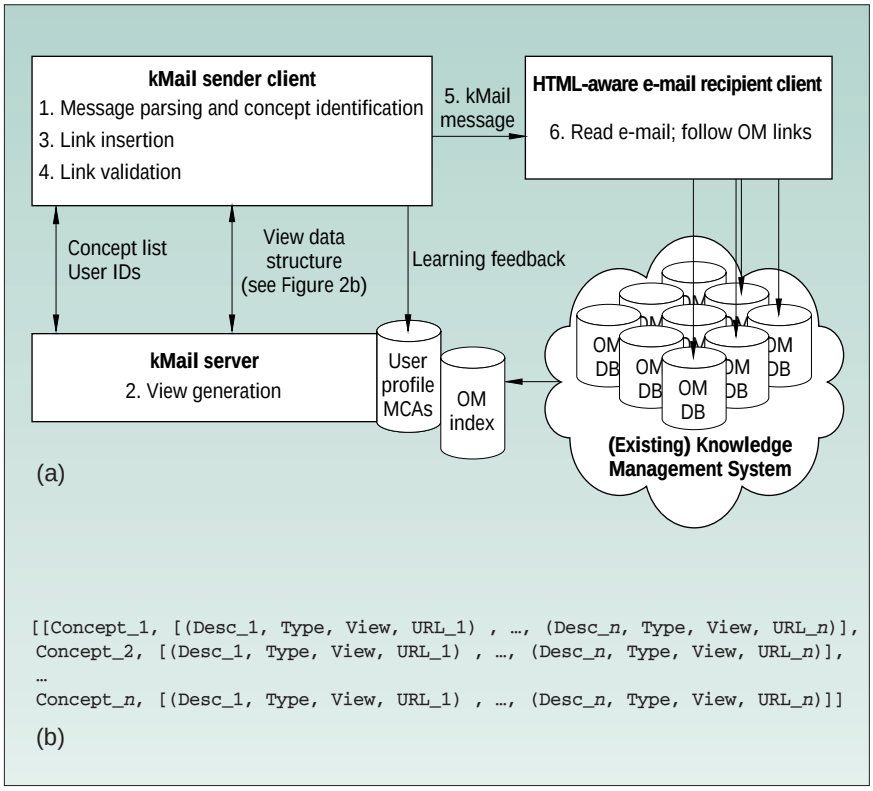

Figure 2. kMail (a) process and (b) the view the data structure returned to the kMail sender client.

Table 1. Memory type categories.

Code	Meaning
1	Definition
2	Graphic image (picture, schematic)
3	Policy statement or guideline
4	Specification (product, equipment)
5	Opinion

- Role(PersonID, StartDate, EndDate, Role),
- Supervisor(PersonID, StartDate, EndDate, SupervisorID),
- Project(PersonID, StartDate, EndDate, Project),
- Concept(ConceptID, ConceptName, ConceptDescription), and
- ConceptUsage(ConceptID, MemoryID, PersonID, Instance).

Four views of organizational knowledge

Given the OM views' basic structure, and the way the system uses them, kMail can apply multiple views in determining which memory items the concept in the current e-mail message should associate with. Four of these views are personal, supervisory, project-related, and role-related.

View 1—personal. A person will prefer to use memory–concept associations that he or she has used in the past. To determine the personal MCA view, we query the OM metaknowledge to determine if the sender, as identified by his unique PersonID (derived from the sender's e-mail address), has one or more memory items that he has explicitly associated with a concept.

For a given concept *C*, a person might have zero, one, or more memory items associated with that concept. This view presents a memory items list that the user has, at one time or another, associated with the given concept. Clearly, this cannot be a sufficiently tight restriction on the search space. It is, however, a viable starting point for further refinement of the result set. This hypothesis follows Eleanor Rorsch's theory of typicality in that it assumes a consistency of usage across situations for the same individual. [10]

View 2—supervisory. A memory–concept association of a person's immediate supervisor is preferable to a colleague's memory–concept association on the same project and less preferable to a person's own personal MCA.

For a given concept *C*, the supervisor can have zero, one, or more memory items associated with that concept. This view shows a list of memory items that the supervisor has, at one time or another, associated with the given concept.

Views 3 and 4—project-related and role-related. We can use a colleague's MCA with a different role but on the same project to effectively restore context. Alternatively, in View 4, a colleague's MCA with the same role on a different project, while less preferable than a same-project MCA, can be useful as well.

We use these first four views as a guide in determining the sort order of memory item results returned to the system for linking with e-mail concepts. After we execute each of the four queries, we merge and sort the results by MemoryType, ranking them within each memory item type according to the corresponding MCA. The system defaults to apply the highest rank to definition-type memory items, and continues down the list shown in Table 1. The actual order of presentation is a controllable parameter, considering that different users prefer alternative ranking orders. The latter three views are modeled after the theory of conceptual coherence.[11] We can also create additional views, such as departmental, in a similar manner, as described in the "Enhancing an e-mail message" sidebar.

Where do we go from here?

With kMail, the act of composing an e-mail message becomes integrated with the knowledge-management process unobtrusively. Whenever a manager links a knowledge item to a message concept, he or she is explicitly confirming that item's current validity. By tracking the usage (and nonusage) of memory items, we can effectively determine which OM parts are relevant and useful.

The editing controls (Figure B in the sidebar) give an e-mail author an intuitive interface for knowledge management. At any point in validating the memory items that the system selects for inclusion in the e-mail, the sender can create a new OM entry and modify or delete any of the entries that the system can retrieve. Modifications and deletions, however, occur only at the view level, so the underlying OM remains intact and we only update the metaknowledge to reflect the user's input. The exception to this immediate update is when the user adds a new memory item to the system. An entry of a new definition or a new URL inclusion might be relevant to the entire organization. A mechanism to manage this sort of knowledge maintenance is the subject of future research. Without an organizational-validation mechanism, the author in question must keep the new additions until such time as he or she can approve them for broader access.

Learning. Whenever the user sends an e-mail message with an MCA link that he or she has explicitly verified, this information goes to the server to update the user profile. We consider explicit verification the simple act of viewing an MCA link to a given concept and leaving it in place. This "non-act" provides confirmation that the user accepts the default association that the system chooses. If the sender explicitly changes the MCA link, replacing it with an alternative, this information is transmitted as well. Not only can this information update the user view, it also provides important feedback as to whether the chosen OM views and the order in which they are applied are appropriate. By monitoring the ways in which knowledge links up with action in practice, kMail takes a step toward being a participatory[8] KMS, alleviating the danger of providing outdated materials.

Limitations. As in any KMS component, there are two nonsystemic functions that are vital to the system's success yet difficult to control. The first, common to all knowledge-management environments, is a sufficient level of user participation to update knowledge and keep it current. kMail takes a step toward promoting this participation by providing update functions to the e-mail author. But, when a manager is composing an action-oriented e-mail message, he wants that message to be clarified and sent—and that is not the most opportune time to be updating knowledge. In this respect, kMail suffers from the same shortcoming as any other KMS.

A second, kMail-specific limitation relates to appropriate views creation and use. As we have shown, the system's successful use depends on properly defining the views relevant to each user. This, too, is in the user's hands, and although we have based initial views on viable contextualization theories, there is clearly more work we could do here.[10,11]

Future directions. Updating the underlying OM is possible by trickling down modifications from the metaknowledge whenever a user has initiated a change. This requires establishing further validation mechanisms, because what the user is changing is his orga-

nizational knowledge view, not necessarily the underlying memory items in the OM.

Incorporating recipient feedback is another direction in which we could extend this work.

CONNECTING KNOWLEDGE TO ACtion through e-mail provides a two-way street for knowledge management on the Internet. It allows relevant OMs delivery directly to an action point, while it also creates a natural setting in which users can update and maintain the system as part of their daily e-mail interactions.

The use of memory–concept association, a metaknowledge-based associative memory, provides a series of useful, extensible associations between different users and the OM knowledge. Managing the MCAs through a straightforward view mechanism lets us allow external access to an organization's internal memory items with sufficiently fine levels of control.

The kMail development, with its roots in dealing with communication problems in multinational corporations, opens new, practical directions for knowledge management and dissemination across the Internet. ■

Acknowledgments

We thank Ehud Dayan for developing the kMail code and Delta Textiles Ltd. for cooperating with our field study. This work was supported by a grant from the Chase Manhattan Bank.

References

1. D.G. Schwartz, "When E-Mail Meets Organizational Memories," *Int'l J. Human-Computer Studies*, Vol. 51, No. 3, 1999, pp. 599–614.
2. D. Te'eni and D.G. Schwartz, "Contextualization in Computer-Mediated Communication," *Information Systems—The Next Generation*, L. Brooks and C. Kimble, eds., McGraw-Hill, New York, 1999, pp. 327–338.
3. J. Habermas, *The Theory of Communicative Action*, Beacon, Boston, 1981.
4. D.E. O'Leary, "Using AI in Knowledge Management: Knowledge Bases and Ontologies," *IEEE Intelligent Systems*, Vol. 13, No. 3, May/June 1998, pp. 34–39.
5. C.A. O'Reilly and L.R. Pondy, "Organizational Communication," *Organizational Behavior*, S. Kerr, ed., Grid, Columbus, Ohio, 1979, pp. 119–150.

Enhancing an e-mail message

Figure A shows a portion of the original e-mail message composed by a product line manager regarding the plans for the upcoming line of men's trunks.

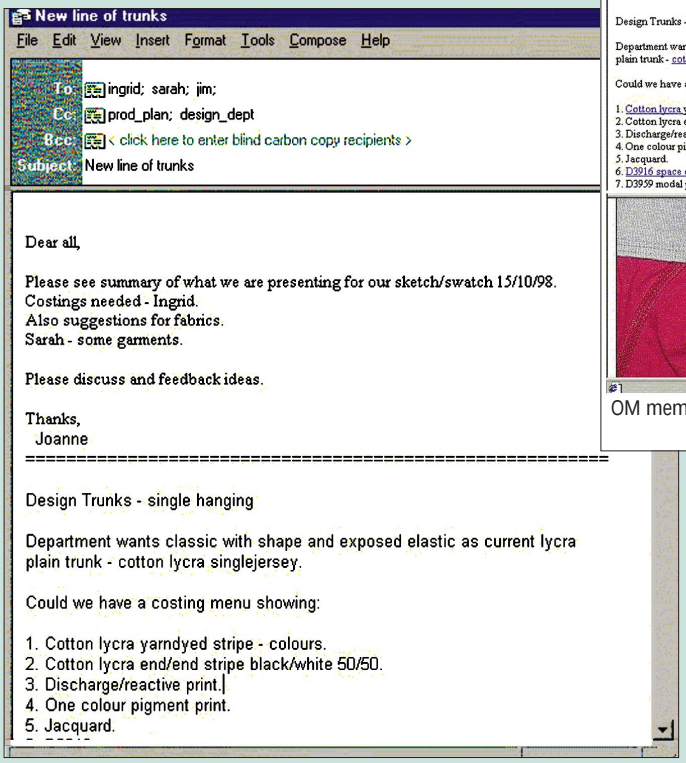

Figure A. Original e-mail message.

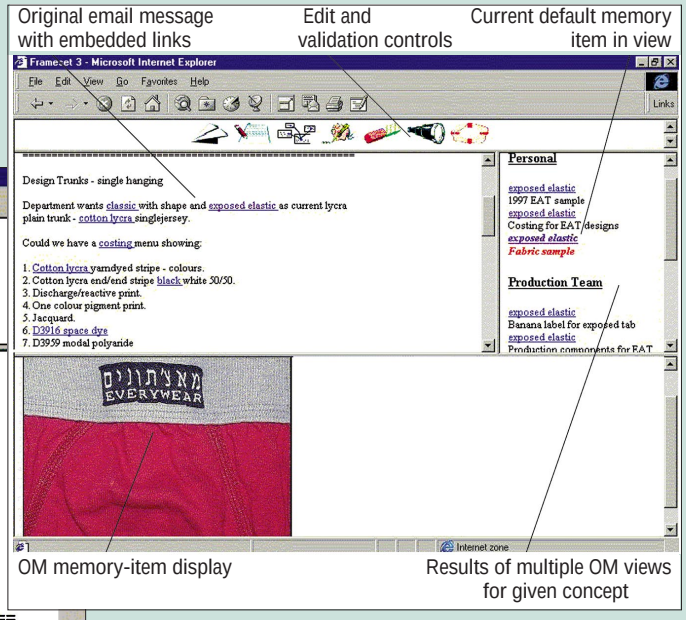

Figure B. Default personal view, exposed elastic tab trunks.

embedded links and related potential links. Figure B shows how a departmental-analysis view affects the "exposed elastic tab" knowledge. In this instance, the view presents memory items based on department of origin (Production, Design, Sales, Finance). The initial link chosen is the fabric sample, indicated in red, as the preferred organizational memory of the e-mail author. Here the system presents the results grouped by department, following the parameters that the sender sets. The

6. A. Abecker et al., "Toward a Technology for Organizational Memories," *IEEE Intelligent Systems*, Vol. 13, No. 3, 1998, pp. 40–48.

7. M.S. Ackerman, "Definitional and Contextual Issues in Organizational and Group Memories," *Proc. 27th Hawaii Int'l Conf. System Sciences*, IEEE Computer Soc. Press, Los Alamitos, Calif., 1994, pp. 203–224.

8. S. Buckingham Shum, "Negotiating the Construction and Reconstruction of Organizational Memories," *J. Universal Computer Science*, Vol. 3, No. 8, 1997, pp. 899–928.

9. C.J. Date and H. Darwen, *A Guide to the SQL Standard*, Addison-Wesley, Reading, Mass., 1993.

10. E. Rorsch, "Principles of Categorization," *Cognition and Categorization*, E. Rorsch and B.B. Lloyd, eds., Lawrence Erlbaum Assoc., Hillsdale, N.J., 1978.

11. G.L. Murphy and D.L. Medin, "The Role of Theories in Conceptual Coherence," *Psychological Rev.*, Vol. 92, No. 3, 1985, pp. 289–316.

David G. Schwartz is an assistant professor at the Graduate School of Business Administration of Bar-Ilan University, Israel, where he serves as head of the Information Systems Division. His research interests include Internet business applications, knowledge management, distributed AI, and the use of management models for intelligent agent architectures. He serves as editor of the *Journal of Internet Research* and is co-editor of *Internet-Based Knowledge Management and Organizational Memories*. He is also a founding partner of Apropos IT Ventures, a venture capital fund focusing on early-stage information technology companies in Israel. Contact him at the Center for Global Knowledge Management, Graduate School of Business Admin., Bar-Ilan Univ., Ramat-Gan, 52900, Israel; dschwar@mail.biu.ac.il or david@aproposit.com.

Dov Te'eni is the director of the Graduate School of Business Administration at Bar-Ilan University and professor of information systems. His research interests include human–computer interaction for decision making, decision-support systems, behavioral aspects of information systems, and software errors. Contact him at the Center for Global Knowledge Management, Graduate School of Business Admin., Bar-Ilan Univ., Ramat-Gan, 52900, Israel; teeni@mail.biu.ac.il.

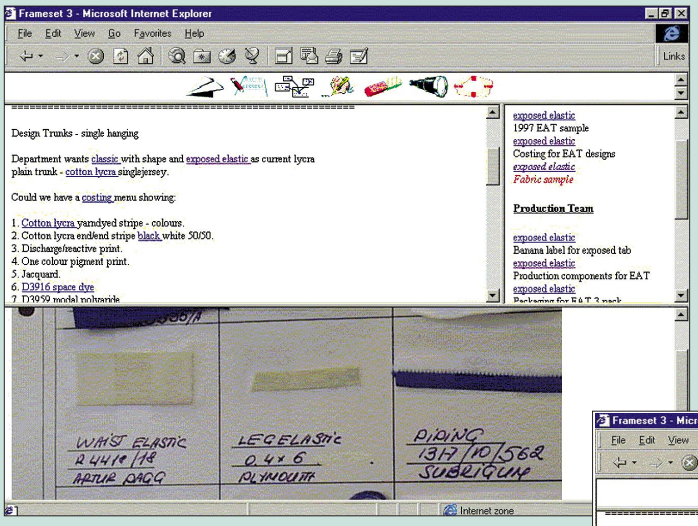

Figure C. Production view, exposed elastic tab components.

e-mail author can choose to link all, one, or a view combination results to the concept in the message. Figure C with production-related memory items and Figure D with design-related memory items show alternative OMs that the e-mail author can choose to include in the view to be sent.

Clearly, there are innumerable figuratively, not computationally ways to combine the user profile characteristics in determining relevancy. For example, a query for same role, same project might be an effective combination in situations where projects have multiple people filling the same role. There can be several alternate views to the departmental focus shown in our example.

The OM views can be a powerful tool in properly utilizing the OM. Domain-specific criteria help determine what the relevant views are in a given organization. Project-oriented companies will require a *same project* view; an organization aligned using functional management approach will require *functional views*; and so forth.

kMail is Internet-friendly in its use and resources allocation. By treating all knowledge resources as URL-accessible, we eliminate any need for replication and redundancy (aside from the use of common proxy caching to improve performance). Knowledge can be distributed strictly on an as-needed basis, and if we get particular about it, on a need-to-know basis as well. Need-to-know knowledge distribution is particularly relevant in a security-conscious or competitive environment.

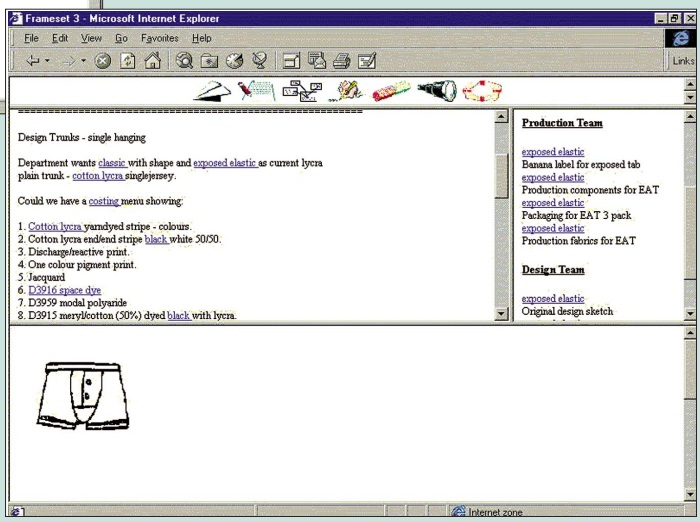

Figure D. Design view, exposed elastic tab original sketch.