Toward Intelligent Message Routing Systems

Thomas W. Malone
Stephen A. Brobst
Kenneth R. Grant
Michael D. Cohen

August 1985

CISR WP No. 129
Sloan WP No. 1709-85
90s WP No. 85-010

Center for Information Systems Research
Sloan School of Management
Massachusetts Institute of Technology

Abstract

As it becomes both technically and economically feasible to send electronic messages to a large number of possible recipients, the problem of deciding who should receive a particular message will become increasingly important. It will be essential to provide tools for extracting interesting and relevant messages from a potentially large pool of less useful information. This paper focuses on the application of techniques from artificial intelligence and organizational science to the development of intelligent message routing systems. Our discussion is illustrated with examples from a series of informal studies of how people share information in organizations, and descriptions of a system under development in our research group at MIT.

Toward Intelligent Message Routing Systems

Thomas W. Malone, Stephen A. Brobst, Kenneth R. Grant, and Michael D. Cohen

1. Introduction

Much of the work that has been done on electronic message systems has focused on technical capabilities and standards for *transporting* messages. In this paper we propose to shift this focus to a previously unappreciated problem that we will refer to as message *routing*. In short, the message routing problem consists of disseminating information via an electronic message system such that it reaches all those persons to whom it is valuable and yet does not interfere with the information processing of individuals who will find no value in its contents.

We argue that this problem will become increasingly important with advances in communication technology and we will suggest three fundamental approaches to its solution. Our characterization of these approaches will be illustrated with (1) examples from a series of informal studies that we have conducted on how people share information in organizations, and (2) descriptions of an intelligent message routing system currently under development in our research group at M.I.T.

2. Background

In recent years a series of standards for message handling systems has been developed based on the IFIP model defined in 1981. The establishment of such standards opens up exciting opportunities for the creation of a worldwide network of interconnected message systems. However, it is already a common experience in mature electronic message communities for people to be flooded with large quantities of electronic "junkmail", and the availability of such widespread communication ability has the potential to overwhelm people with even more messages that are of little or no value [Den82, Pal84, Wil84].

3. Previous Solutions

The most common approach taken to combat the information overload problem in electronic messaging environments is the establishment of some form of distribution list or conferencing structure based on interest groups. In such a scenario, users associate themselves with centralized distribution lists or conference topics which are used to disperse messages pertaining to particular subjects. In most cases, people may associate themselves with as many or as few distribution lists as desired.

Although distribution lists and conferences are topical in nature, their focus is generally much broader than the interests of a single person. Furthermore, since these methods of disseminating information are targeted for relatively large audiences, it is usually impossible for the content of all information distributed to be of interest to all recipients. The users implicitly specify a general area of interest by associating themselves with particular distribution lists or conferences, but they are ultimately responsible for deciding on a per message basis which information is of interest within this context.

Another method of dealing with the information overload problem is to use keyword filtering and other techniques developed for large text retrieval systems [Sal83]. Such techniques have almost never been applied to electronic messaging systems, and we believe they hold substantial promise. Traditional text retrieval systems, however, have substantial limitations [Bla85]. In this paper, we will suggest how more sophisticated techniques from the realms of artificial intelligence and organizational science can be used to design intelligent message routing systems that facilitate information sharing in on-line communities.

4. Studies of Information Sharing in Organizations

To help understand the kinds of automated aids for sharing that would be desirable, we have conducted several preliminary studies of how various kinds of information are shared in organizations. We have done some relatively unstructured interviews of 19 people about their information filtering experiences, needs, and desires. We have also performed more structured interviews of over 30 people which focused on different kinds of information filtering environments: (1) processing the contents of their in-box, (2) reading and dealing with electronic mail, (3) scanning electronic bulletin boards, and (4) examining a table of contents as a basis for deciding which articles to read in a journal. In each case, we asked the subjects to explain in detail why they made the filtering decisions they did. Then we used these "expert protocols" to identify the general processes

and specific kinds of knowledge and decision rules that people used for extracting valuable information from a large pool of available sources.

5. Approaches to Information Filtering

One very valuable outcome of the organizational studies was the refinement of our notions of three different approaches that could be incorporated into an automated message filtering system. We will refer to these techniques as the *cognitive*, *social*, and *economic* approaches to information filtering. In the following sections these approaches will be described, illustrated with examples, and extrapolated to implications for an intelligent message routing system.

5.1 Cognitive Filtering

This approach to message filtering is based on characterizing the contents of a message and the information needs of potential message recipients and then using these representations to intelligently match messages to receivers. Distribution lists and simple keyword matching are rudimentary forms of this approach.

When making use of the cognitive approach, people often identified the content of a message by matching against certain keywords or phrases in its text. From this assessment of the content of a message, a decision is made regarding its applicability to current information needs. Some examples of people using this filtering technique are given below.[1]

- *My work currently focuses on fault tolerance issues for a static dataflow machine. This journal article that addresses fault tolerant computing in a distributed processing environment is therefore of direct interest to me.-* **research assistant, table of contents study**

- *That seminar [pattern recognition] would have been of interest to me if it wasn't on a Tuesday afternoon.* - **research staff, bulletin board study**

- *I would do an exclusive scan looking for telecommunications articles except for voice mail. In [other subjects] I would do the opposite - I would search for a word that would trigger a hit.-* **corporate executive, table of contents study**

- *That article sounds too OR-ish, so I wouldn't be interested.-* **research administrator, table of contents study**

Notice that these decisions are generally based on much more sophisticated heuristics than a simple keyword search. Complex combinations of both positive and negative keyword filtering techniques were used.

Throughout our organizational studies, the cognitive approach was used with great efficiency by the people with whom we worked. Part of the reason for this success appears to be that there are many *classes* of messages, for which receivers have a pre-formed notion of what information was necessary to assess the applicability of the message. To represent these different classes of messages and the irrelationships and attributes, the intelligent message routing system we are developing uses knowledge representation techniques from the realm of Artificial Intelligence involving frame inheritance networks [Bra85].

For example, we expect the users of our intelligent mail filtering system to be able to conveniently compose messages using a network of different templates for different types of information (e.g., meeting announcements, bug reports). A *meeting announcement* template, for instance, would have fields for "time," "place," "organizer," and "topic." Its subtype, *seminar announcement*, would add a field for "speaker." Another subtype, *PC Committee Meeting Announcement*, for instance, could have default values for time, place, and organizer already filled in. Receivers of messages will be able to construct much more sophisticated filters than would be possible with simple keyword searches (*e.g.*, *"show me all the announcements for seminars at MIT organized by people in my department except for seminars which occur on Tuesdays "*).

Different groups can develop detailed structures to represent the information of specific concern to them. For example, a product design team might have an elaborate network of message types describing different aspects of the product (*e.g.*, market size estimates, response time specifications, alternative power supply vendors). Then, for instance, marketing specialists who believe that the critical factors determining potential market size for the product are cost and response time can devote most of their attention to the messages concerning these factors and ignore all the rest of the technical specifications for the product.

One characteristic of message classes that is critical in formulating filtering rules is the *purpose* of the message. We expect that an important part of the frame inheritance network for message types will be a taxonomy of the various *communication acts* that a message might embody [Sea75, Ked74]. For example, messages whose purpose is to *request information* should be routed to people who know about the topic of the message, while messages whose purpose is to *provide information* should be routed to people who are interested in the topic of the message.

5.2. Social Filtering

This approach does not rely on the ability of a message routing system to understand either the contents of messages or the interests of receivers. Instead, it focuses on supporting the personal and organizational interrelationships of individuals in the electronic message community. The social filtering mechanism complements the cognitive approach by judging the potential value of a message based not just on its topic - but also on the characteristics of its sender.

One important dimension used for filtering by participants in our study was the organizational relationship between the sender and the receiver. For example, one participant said: *This message is from my supervisor. That means it's probably of importance to me.* The use of organizational relationships for filtering can often be quite subtle. In one case, a person focused on a particular message in his electronic mailbox because it was sent by his group leader's secretary: *I know that she would not have sent me anything, so the message must be on [my group leader's] behalf.* These organizational relationships often interact in complex ways with personal relationships as well: *"I would have to respond to a memo based on several requirements...if it came from James Long[2] [chairman of the board]...I would react quickly - that would be an alert. If it came from Sarah Rowe...who is next on the hierarchy...that would be another alert flag. If it came from John Smith...who is my direct manager...that would be a real immediate alert flag - because...I deal with him [personally] - I work with him, assisting him on matters."* We plan to include facilities in our message routing system that will let receivers construct their own individual filtering rules based on organizational and other characteristics of senders.

A very widely used method of social filtering relies not just on the characteristics of the author of a message but also on the references and recommendations of other people. Perhaps the best example of this phenomenon is the everyday experience of most researchers who attempt to "keep up with the literature" in their field. It is our experience that almost no one actually scans all the journals that are potentially relevant to their interests. Instead, they read the articles which their colleagues most often refer to and recommend.

A rudimentary kind of support for this process already exists in most electronic messaging systems in the form of facilities for forwarding messages to other people who may find them interesting. A more elaborate kind of support would allow receivers to create lists of people whose opinions they value on various topics. When these people read messages, they may choose to "endorse" them. Messages can then be prioritized for a given receiver based on the number of endorsements they receive from people

on that receiver's endorsement list. An obvious extension would allow assigning different weights to different endorsers and letting endorsers indicate the strength of their endorsements. For high volume topics, this process can be formalized still furtherby having some people perform the role of an "editor" who selects messages on the basis of both relevance and quality.

5.3. Economic Filtering

This approach focuses on how to provide both positive and negative incentives to control the quantity and quality of information flows. The widespread use of distribution lists and electronic bulletin boards facilitates the dissemination of information to an extremely large audience with very little effort. The danger associated with this is that the very low marginal cost of sending additional copies of messages [Cra82] may encourage the proliferation of electronic messages that have little informational value to the majority of their recipients. It is the goal of the economic filtering mechanism to minimize this phenomena, while encouraging people to disseminate valuable information they might otherwise have neglected to send.

In our organizational information sharing studies we found that persons oftenhad to make a cost versus value decision to determine whether or not to process a particular piece of information. The length of a message was the primary factor used by a recipient to estimate its cost. For example, one person described his filtering mechanism as follows: *I am extremely tight on time this week...I don't have the spare moments to read any messages with more than a screenful of text unless they are extremely important.*

A more subtle economic filtering technique employed by some of the subjects in our studies was to use the cost of the message to its sender as a, consideration in evaluating the potential information gain from a message. For example, a secretary to the president of a small company explains her filtering process as follows: *He [the president] is too busy to look at all the mail that comes into his office. I try to weed out anything that looks like a form letter or has a bulk mailing label on it.)* If the number of recipients of some piece of information is very high (*i.e.*, form letters and bulk mailings), then the cost per person of formulating the information will be fairly small to the sender. However, the optimal piece of information to a reader is one that is customized to his needs - at great expense to the sender of the information. This is a good explanation for the fact that a person will [almost] always read mail exclusively addressed to himself. The information content in these messages has [presumably] been chosen especially for the particular needs and interests of its recipient.

There is, however, an important tradeoff that takes place between quality and personalization. Even though published articles have a large number of recipients, the lack of personalization is often implicitly outweighed by higher quality. In fact, publications with a larger audience seem to have more appeal, at least in part, because they can afford to publish higher quality information. By spreading the cost over a larger number of subscribers, a large volume publication can afford to absorb the higher cost of providing extremely well formulated and carefully selected information to its readers. Insuring this high level of quality is, of course, one of the primary goals of an editor.

One feature of most current message systems that is highlighted by an economic perspective is that the cost of sending multiple copies of messages is borne primarily by the receivers who must process them. The economic approach to information filtering implies that it might sometimes be in the best interest of an electronic mail community to shift these incentives so that senders of messages are charged in a way that reflects some of the costs to receivers [Den82]. For example, the senders of unsolicited messages can be charged (in actual money or some point system) in proportion to the value of the time people will spend reading the messages, that is, more for long messages, more for messages sent to many people, and more for messages to highly paid recipients. In an even more extreme use of this idea, people who receive undesirable "junk mail" can indicate that fact and the sender will then be penalized by an additional surcharge. This scheme thus provides incentives to both reduce "junk mail" and to make messages more concise.

A somewhat subtle way to reduce "junk mail" is suggested by the observation above about how receivers use the cost of a message to its sender (*e.g.*, bulk rate versus first class postage) as a consideration in filtering. An electronic messaging system can use this approach by letting senders spend resources to signal receivers that a message deserves more than the usual priority. Some receivers might then have "asking prices" on their mailboxes that screen out all messages below a certain cost [Den82]. Other receivers might construct rules that use the cost of the sender as one of a number of factors in filtering and prioritizing messages.

In any system where senders pay some of the costs of communication (including current systems where the cost is only time) there is a disincentive for people with potentially valuable information to disseminate it. To remedy this deficiency, the recipients of valuable information can sometimes compensate the sender. For example, [Tur83] describes a system that supports an internal "free market" for information and services within an organization. Certain kinds of information are highly desirable and employees who know (or have skills to obtain) this information can sell the information or "contract out" some of their time to high bidders anywhere in the organization.

5.4. Combined Approaches

It is important to realize that the most useful message routing systems are likely to be combinations of the approaches we have described. For example, the pricing schemes are likely to work much more effectively when the topic structuring mechanisms described earlier are used to indicate which topics receivers are willing to pay for and which they would have to be paid to read about. One problem with this approach, however, is that senders might try to receive extra payments by classifying their messages into highly valued topic areas - for which the messages were actually irrelevant. This problem can be ameliorated by using social filtering mechanisms. For example, "editors" can be paid for filtering information on certain topics. Editors who filter well will be able to charge more for their services and those who don't will find their "subscribers" dwindling away.

6. Conclusions

As we move into an information age in which it is both technically and economically feasible for almost anyone to communicate large volumes of information almost instantaneously across long distances, message recipients are experiencing a desperate need of support in their information processing task. In this paper we have addressed this problem by developing a framework for intelligent information sharing systems. Three approaches to information filtering have been distinguished: cognitive, social, and economic. Several implications of these approaches for the design of an intelligent message routing system are summarized in Table 1.

Traditional user interfaces for computer systems have emphasized the support of separate problem-solving by isolated individuals. In designing electronic messaging systems, however, it is at least as important to design good "organizational interfaces" for supporting cooperative problem solving by groups of people [Mal85]. We believe that such organizational interfaces, like the intelligent information sharing systems we have described here, will become increasingly important as computers continue to gain widespread use.

7. Acknowledgments

Table 1:
Sample Implications of Different
Filtering Mechanisms for Designing
Intelligent Message Routing Systems

I. Cognitive Approach

- Frame inheritance network for message classes

II. Social Approach

- Encode sender characteristics

- Endorsements

III. Economic Approach

- Pricing structures for controlling information flows

References

[Bla85]   Blair, D. and Maron, M. "An Evaluation of Retrieval Effectiveness for a Full-Text Document-Retrieval System". Communications of the ACM, Vol. 28, No. 3, March, 1985. pp. 289-299.

[Bra85]   Brachman, R. and Schmolze, J. "An overview of the KL-ONE Knowledge Representation System", Cognitive Science, in press.

[Cra82]   Crawford, A. "Corporate Electronic Mail - A Communication Intensive Application of Information Technology". Quarterly, September, 1982, pp. 1-14.

[Den82]   Denning, P. "Electronic Junk". Communications of the ACM, Vol. 23, No. 3, March, 1982. pp. 163-165.

[Ked84]   Kedzierski, B. "Knowledge-Based Project Management and Communication Support in a System Development Environment". Fourth Jerusalem Conference on Information Technology. Jerusalem, Israel, 1984.

[Mal85]   Malone, T. "Designing Organizational Interfaces". Proceedings of the CHI '85 Conference on Human Factors in Computing Systems. San Francisco, California, 1985.

[Pal84]   Palme, J. "You Have 134 Unread Mail Do You Want to Read Them Now?". IFIP Conference on Computer Based Message Services. Nottingham University, 1984.

[Sal83]   Salton, G. and McGill, M. Introduction to Modern Information Retrieval. New York: McGraw-Hill, 1983.

[Sea75]   Searle, J. "A Taxonomy of Illocutionary Acts" in K. Gunderson (ed.), Minnesota Studies in the Philosophy of Language. Minneapolis: University of Minnesota Press, 1975.

[Sta83]   Stallman, R. "ZMail Manual". M.I.T. Artificial Intelligence Laboratory, 1983.

[Tur83]   Turoff, M. "Information, Value and the Internal Marketplace". New Jersey Institute of Technology, Newark, New Jersey, 1983.

[Wil84]   Wilson, P., Maude, T., Marshall, C., and Heaton, N. "The Active Mailbox - your on-line Secretary". IFIP Conference on Computer Based Message Services. Nottingham University, 1984.

Footnotes

[1]Italicized typeface is used to represent remarks of our subjects closely paraphrased from interview notes while quotations are reserved for verbatim transcriptions from tape recorded interviews.

[2]The names of all those persons referred to by the interviewees have been changed.