

SUPPORTING INFORMAL COMMUNICATION VIA EPHEMERAL INTEREST GROUPS

Laurence Brothers,
Jim Hollan, Jakob Nielsen, Scott Stornetta,
Steve Abney, George Furnas, and Michael Littman

Bellcore
445 South St.
Morristown, NJ 07962-1910
{quasar, hollan, nielsen, stornetta, abney, gwf, mlittman}@bellcore.com

ABSTRACT

In this paper, we introduce ephemeral interest groups for supporting informal communication. Ephemeral interest groups are electronic discussion groups that, in contrast to bulletin boards and the like, are short-lived and ad hoc. They are designed as a medium for informal discussions of items broadcast to a wider community. We have implemented a prototype system to explore ephemeral interest groups. We discuss the goals of the system, characterize its evolution over the last ten months of deployment, and sketch our plans for future developments.

KEYWORDS

Ephemeral interest groups, informal communication, discussions, hypertext, bulletin boards, email, information filters, critical mass of users, measures of informality.

EPHEMERAL INTEREST GROUPS

In a recent paper [3], we view support for informal communication in terms of needs, media, and mechanisms. There are underlying needs served by communication, and different media enable different mechanisms for meeting those needs. Our approach is to exploit the distinctive mechanisms of the computational medium, as opposed to imitating the mechanisms of traditional interaction. In our earlier paper we sketched a series of projects to elaborate this position. In this paper we provide empirical details of one of those efforts: supporting informal communication via *ephemeral interest groups*. Ephemeral interest groups are

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

short-lived, ad hoc electronic discussion groups spawned from items from a more widely-broadcast information source.

Electronic mail and other forms of electronic communication are increasingly being used for informal discussion. Most of these discussions are mediated by e-mail mailing lists or by bulletin boards. At Bellcore a simple bulletin board provides a forum for discussion of issues of interest to the general community. The majority of users access it through a Unix^{®*} program, *msgs*. While originally intended for announcements about the system that are “short pieces of information suitable to be read once by most users of the system,” it has come to be used more broadly to exchange information, make queries and requests, and discuss topics of general interest.

Based on our earlier proposal for supporting informal communication [3], we began to explore changes to the *msgs* system that would make it more useful for the community. We decided that any changes we made should be non-obtrusive, so that people who so desired could continue to use the system exactly as they had before. We understood that this greatly constrained the changes we could make, but we wanted to avoid any disruption to our colleagues.

A common complaint about *msgs* is that it contains too much irrelevant information. One might argue that a bulletin board consisting of multiple interest groups would better support users, because people could subscribe to or browse only those groups that interest them. However, we question this alternative for a number of reasons. First, there are many benefits of having a channel of communication that is read by virtually everyone in our community. In addition, it is difficult to generate a good set of categories for articles,

* Unix is a registered trademark of Unix System Laboratories, Inc.

and the burden of selecting an appropriate category is placed on the sender of each article. Studies of hypertext [1][2][4][5] have demonstrated that usability problems arise if systems require users to structure information prematurely. For these reasons, we felt that a hierarchical bulletin board system would not be useful for managing the kind of wide-ranging and short-lived discussions that are typical of *msgs*.

This led us to propose ephemeral interest groups. Unlike bulletin boards, netnews groups, and mailing lists, ephemeral interest groups can be created at virtually no transaction cost to users or to system administrators. The groups can be thought of as disposable, intended to last only days or hours.

We created a system, called *topics*, which almost invisibly supports ephemeral interest groups through *msgs*. In keeping with our design constraint, *msgs* continues on the surface as it has always been. But after reading an article, users now have the alternative of *joining* the article before going on to the next one. In the *topics* model, every article represents a potential topic of discussion, of interest to some subset of the community, and by joining an article, the user expresses interest in the topic. Joining an article adds the user to the ephemeral interest group associated with the article, creating it automatically, if necessary. The user receives a copy of the article in his or her regular e-mail along with a copy of discussion that has already taken place. To post to the ephemeral interest group, users need only reply within their normal mail reading program to any of the articles in the discussion. Thus, *topics* is similar to a traditional mailing list, but the procedure for creating the list is automatic and instantaneous.

An important consequence is that follow-ups can easily be sent to those people who have expressed an interest in a topic, rather than to the whole community. In this way, we expected to decrease follow-up activity on *msgs*. At the same time, we expected that there would be more discussions overall, and that discussions within ephemeral interest groups would have more follow-ups than they would have had on *msgs*, since participants know that they are addressing only those who have expressed interest in the topic, not the whole community. We also expected this awareness to influence the character of discussions; in particular, to make them more informal and personal.

Analysis of Prior Msgs System Usage

Before we made any changes to the *msgs* system, we analyzed the content of a collection of *msgs* articles and interviewed users of the *msgs* system. We classified about 700 articles from the period April 8 to May 24, 1991 into five mutually-exclusive categories:

- Information (249 articles, or 36%) The posting's main intent was to provide information to the community of *msgs* readers.

- Question (134, or 19%) The posting asked a question that could be answered by an e-mail response, or some other simple action from one's computer or telephone.
- Request (79 articles, or 11%) A request differs from a question in that responding to a request involves an action that cannot be handled from one's computer or phone; for example, giving a ride home from work.
- Offer (68 articles, or 10%) An offer differs from information in that taking advantage of an offer requires one to do more than just read it. An example of an offer is a posting that a book is available to the first person who picks it up.
- Scheduled Event (168 articles, or 24%) The posting announces a talk or other event taking place at a specific time.

We also marked articles with three optional codes, orthogonal to the five categories above:

- Leisure (119 articles, or 17%) Postings which seemed to have no direct bearing on work or work-related activities.
- Follow-up (194 articles, or 28%) Postings which were clearly following in the thread of a previous posting. Of these, 125 were classified as discussions, as opposed to repeated talk announcements.
- Groups (180, or 26%) Postings which were clearly intended for a definable subset of the *msgs* community.

From this analysis, we looked for opportunities to move traffic from *msgs* into ephemeral interest groups. For example, the analysis indicates that follow-ups to questions are often posted back to the entire community, presumably because the author believes people other than the requester are interested in the response, but has no way of determining precisely who might be interested. If an ephemeral interest group were attached to the query, and interested parties expressed their interest by joining the group, follow-ups could be posted to the ephemeral interest group rather than to the entire community. Similarly, many of the scheduled event announcements are second and third reminders. We conjectured that users would prefer to receive follow-up notices only if they had at least a slight interest in the talk to begin with, something that they could judge when they read the first talk announcement. Through ephemeral interest groups, a user could express interest in the first announcement, and follow-ups could be limited to interested parties. We also expected that some of the postings classified as groups would have been addressed to only the relevant individuals, if ephemeral interest groups had been available. In short, the postings that might not have appeared on *msgs*, had ephemeral interest groups been available, were the follow-up postings and group postings. Together, they make up 49% (343 postings) of the volume of traffic on *msgs*. The prospect of reducing *msgs* traffic by half seemed attractive.

We also did a series of informal interviews of a stratified sample of the user community. Those interviews confirmed our intuitions. *Msgs* is widely used in the community; it is

Message 10071:
From suri@polya.bellcore.com Mon Jan 13 13:28:47 1992
Subject: Lab 212 Seminar, Thursday 01/16/92
(16 lines) More? [ynq]

Figure 1 *Msgs* Prompt.

viewed as essential for obtaining a variety of work related information; and it is perceived to contain too much irrelevant information.

Modifications to *Msgs* to Support Ephemeral Interest Groups

Because we wanted to evaluate our ideas about ephemeral interest groups in the existing large community of *msgs* users, we had to work within the limits of the *msgs* system. *Msgs* users are predominantly Bellcore researchers, but also include some administrative and managerial staff. We recognized that, while many might be interested in using new features we created, we had to leave the system essentially unchanged for those who did not wish to participate. Because of the primitive nature of the *msgs* interface, this proved to be a significant constraint. However, we viewed the opportunity to experiment with a large user community as offsetting that limitation.

Most users read *msgs* articles using the *msgs* program, but a minority use a reader in GNU Emacs, the Andrew Mail System, or various netnews readers. Though we were not able to modify the netnews readers, by modifying *msgs*, and making minor changes to the Emacs and Andrew readers, we were able to provide the ephemeral interest group functionality to the large majority of users, with fairly minor modifications to the three systems.

Before settling on our final approach, we considered two other mechanisms for exploring ephemeral interest groups. First, we considered creating a special program to attach ephemeral interest groups only to announcements of talks, but rejected that because its success required all users to change their established behavior for announcing talks. Next, we implemented the system using an early version of dynamic mail objects. Several trial users complained of disruption to their normal mail activity. Finally we decided to implement ephemeral interest groups through ordinary electronic mail, making only minimal changes to the *msgs* program. Only interest group creation and subscription would be handled through *msgs*. Figure 1 depicts the basic prompt of the existing *msgs* bulletin board system. The user is presented with the sender, subject, date, and size of an article, and has the choice of reading it, skipping to the next article, or quitting the application. There are a number of additional features, though most users are unaware of them. These include saving the article to a file, entering a mailer, or jumping to a specific article. The system has no on-line command-lists or help.

The only visible modification we made was the addition of a fourth alternative to the command prompt: instead of

[ynq], the new prompt is [ynqj]. The effect of the 'j' command is to join the user to the ephemeral interest group associated with the current *msgs* article.

Most of the ephemeral interest group functionality is provided by a separate program, called *topics*. When a user types 'j' in response to a *msgs* article, the article is sent by electronic mail to the *topics* program. *Topics* looks in a central archive for an ephemeral interest group corresponding to the given *msgs* article, creating a new group if none already exists. It adds the user to the subscriber list for the group, and sends a digest of current discussion to the user via electronic mail. If the user replies to the digest, or to any other mail distributed by *topics*, the reply goes to *topics*, not *msgs*. *Topics* redistributes such replies only to subscribers of the ephemeral interest group, not to the entire *msgs* community.

In addition to handling requests generated by the 'j' key, and redistributing subsequent ephemeral interest group postings, the *topics* program responds to explicit commands sent via electronic mail. For a given group, users can request a list of subscribers or a digest of all postings. They can unsubscribe from the group, subscribe to other existing groups, or create new groups. A 'help' command is also provided.

EPHEMERAL INTEREST GROUP USAGE

In this section, we examine how ephemeral interest groups have been used, particularly, whether they have been used as we expected, and what weaknesses should be addressed in the next iteration of the system. The major results are: that ephemeral interest groups require a fairly large user community to function well; that ephemeral interest groups are good information filters; that discussions in ephemeral interest groups are indeed more informal and personal than discussions on *msgs*; and that ephemeral interest groups have been perhaps too unobtrusive, resulting in less usage than there might have been.avenues

Trends in Ephemeral Interest Group Formation

Use of the *topics* system began in August 1991. As of July, 1992 1279 ephemeral interest groups have been formed, and there have been 713 follow-up postings to groups. 337 people have joined at least one group; 172 have joined at least three groups.

The number of functioning ephemeral interest groups created in each of the first 46 weeks after the introduction of the 'j' command is shown in Figure 2. For the purposes of this figure, a *functioning* ephemeral interest group is one




Figure 2 Number of functioning ephemeral interest groups for the first 46 weeks. Only groups with at least three postings and at least three subscribers are counted.

with at least three postings (the original *msgs* article and at least two follow-ups) and at least three subscribers. Obviously, there are major random fluctuations in interest group formation from week to week. The long-term trends do show some interesting patterns, however. After an initial flurry of activity, as users tried out the new feature, interest group creation declined to about two new functioning groups per week in early October 1991. Usage increased slowly until mid-November 1991, when a fairly steady level of about four new functioning groups per week was achieved. This result confirms that new technology is often slow to achieve full penetration of a large user environment.

One dramatic effect in the figure is the severe reduction of interest group creation in the last half of December. In fact, the laboratory was only closed for a few days around Christmas, but a fairly large number of users were gone for extended periods of time throughout December. The elimination of functioning interest group creation during the holiday season can be ascribed to the need for a critical mass of users. Our laboratory may not be large enough to sustain ephemeral interest groups unless almost everybody is present.

In addition to the downturn in interest group creation during the holidays, the figure shows that interest group formation took almost the whole month of January to regain its former level of about four new functioning groups per week. *Topics* usage was probably so slow to recover because users had effectively forgotten about it, and the 'j' in the *msgs* prompt was not enough to remind them of its availability. In short, our attempt to be unobtrusive may have been *too* successful. We return to this point below.

Usage remained fairly steady in February and March, with between two and four new functioning ephemeral interest groups created each week. Starting with the week of April 1st, however, usage declined almost to zero and has remained very low ever since. The week of April 1st saw an extensive discussion carried out on the main *msgs* board with a large number of practical, humorous, political, and "flaming" postings, all on the topic of nesting geese in the parking lot. We hypothesize that this high-visibility discussion made it acceptable again to post follow-up messages to the main *msgs* system which reemerged as the forum for electronic discussions in our lab.

Ephemeral Interest Groups as Information Filters

To assess how effective ephemeral interest groups were at directing information only to interested parties, we conducted a user survey in late February and early March of 1992. This period corresponded to the 25th through 27th week of the use of the *topics* system. We surveyed two different kinds of *msgs* users: those who subscribed to ephemeral interest groups, and those who did not. We counted as subscribers only those users that subscribed to at least three groups, of which at least two were within the last 90 days. People who subscribed to fewer groups were considered not to be *topics* users.

We collected every follow-up article posted to either the general *msgs* bulletin board or to a discussion group during the survey period. We sent each follow-up article to a randomly chosen *topics* user and to a randomly chosen non-user. We also sent each *topics* follow-up posting to a random member of the interest group it had been posted to. We classified interest group participants as being either the originator of the *msgs* article that spawned the group, or as

	No. Queried	Respondents	Resp. Rate	Relevance
Originators rating follow-ups to the ephemeral interest group they started	11	7	64%	4.9
Subscribers rating follow-ups to a subscribed ephemeral interest group	14	14	100%	3.9
Subscribers to ephemeral interest groups rating follow-ups to <i>other</i> groups	36	24	67%	2.7
Subscribers to ephemeral interest groups rating follow-ups posted to <i>msgs</i>	31	23	74%	3.0
Non-subscribers rating follow-ups to an ephemeral interest group	43	21	49%	3.0
Non-subscribers rating follow-ups posted to the <i>msgs</i> board	40	19	48%	2.7

Table 1 Rated relevance of follow-up articles posted either to an ephemeral interest group or to the general *msgs* bulletin board. The difference in ratings given by the originators and subscribers to the group to which a follow-up article was posted is significant at the $p < .05$ level. The difference in ratings given by subscribers to the group to which a follow-up article was posted and subscribers to other groups is significant at the $p < .01$ level. The differences in ratings between the bottom four lines in the table are not statistically significant. The difference in response rates between the subscribers and non-subscribers is significant at the $p < .01$ level. Relevance was rated on a 1–5 rating scale, with 1 indicating a completely irrelevant article and 5 indicating a very relevant article.

being a regular subscriber. For postings to interest groups, every posting after the initial seed article was considered to be a follow-up article. For postings to *msgs*, we made a subjective decision as to whether the article was a follow-up to a previous *msgs* article.

Each survey respondent was asked to assess how relevant the specific follow-up article was to him or her personally. The concept of “relevance” was described to the respondents as including both job-related need-to-know and more personal interest in the contents of the article. Relevance was rated on a 1–5 scale, with 1 indicating a completely irrelevant article and 5 indicating a very relevant article.

The results of the survey are reported in Table 1. One major result is that those who subscribed to a given discussion group considered follow-up articles posted to that group quite relevant to their interests (average score 3.9 on the 1–5 scale), in contrast to non-subscribers, who rated the same articles as less relevant to their interests (about 3 on the 1–5 scale). Discussion group originators considered follow-ups posted to their discussion groups as even more relevant to their interests (4.9). The higher rating of discussion group follow-ups by discussion group members cannot be ascribed to discussion group follow-ups being of higher “quality” than *msgs* follow-ups. The table shows that those who did not subscribe to a given group considered postings to that group to be no more or less relevant than postings to *msgs*. (None of the differences in the last four lines of the table are statistically significant.)

We conclude that ephemeral interest groups can successfully direct information to those who view it as most relevant to their interests. Given the information overload faced by most of our users, this filtering function is an important benefit of ephemeral interest groups.

The Tone of Discussions on *Msgs* and *Topics*

In order to get an idea whether *topics* discussions were more informal and personal than *msgs* discussions, we compared the frequency of personal pronoun use in each. Our assumption is that personal pronouns, specifically “I” and “you”,

are used more extensively in personal communications than in formal communications.

We analyzed a six week period (10/10/91–11/20/91) in detail, for both *msgs* and *topics*. During that period, there were 659 *msgs* postings, of which about a fifth (128) were “joined” to create topics. About a third of these (47) actually generated follow-up *topics* discussions, with a total of 107 follow-up articles (an average of 2.2 follow-ups each). During the same period *msgs* continued to be a forum for follow-ups as well, there being almost as many articles of a follow-up nature (95) posted there. We note that the average audience for discussion in these two media is quite different. *Msgs* discussions are read by all 250+ readers, whereas discussions in the ephemeral interest groups are read only by those who expressed interest in the topic (8–9 on average).

Msgs articles were categorized into three groups:

- *Msgs* articles that generated no follow-up *msgs* discussions
- *Msgs* articles that did generate follow-up *msgs* discussions
- *Msgs* articles that were themselves follow-ups to previous *msgs* articles

An analogous classification of *topics* articles yielded:

- *Topics* articles that generated no follow-up postings to the interest group (although by definition, people had “joined” such an article in hopes of future discussion)
- *Topics* articles that did generate follow-up postings
- *Topics* articles that were posted as follow-ups to a previous posting to the interest group

Note that only the last one of these involves articles that were mailed directly to the interest group; the first two involve articles that originally appeared on *msgs*.

For calibration purposes, we also examined electronic mail. The electronic mail data was more recent and was taken from 12 users who volunteered to run the word-counting

Personal Pronouns in Messages

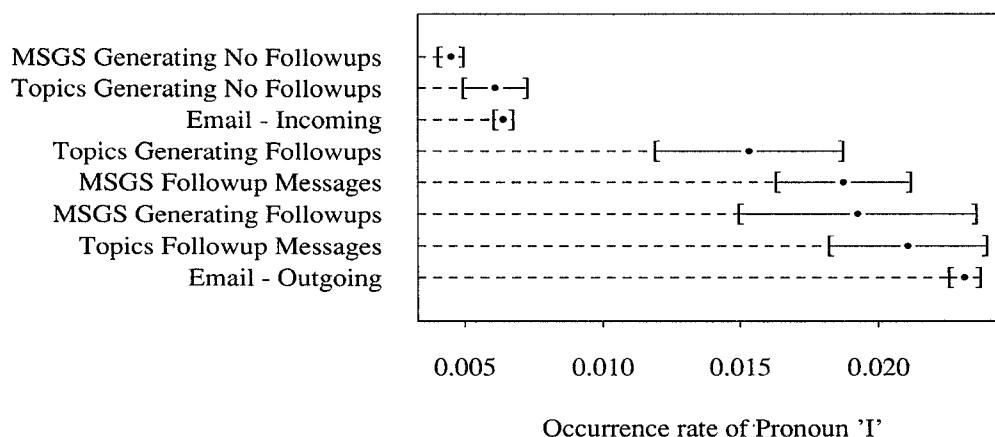


Figure 3 Analysis of the occurrence of the pronoun “I” in articles of different types. The plotted points indicate the estimated probability that a random word drawn from such messages is the pronoun “I”. Brackets show 95% confidence intervals on the estimates.

scripts over their mail files. The two categories of mail examined were:

- incoming electronic mail – unfiltered collection of incoming e-mail articles received by the user
- outgoing electronic mail – unfiltered archives of outgoing e-mail articles sent by the user

We assume that outgoing mail is more personal on average than incoming mail, since the latter includes mass mailings. Put another way, outgoing mail counts articles before mail alias expansion multiplies their impact, whereas incoming mail is after alias expansion. Thus incoming mail includes proportionally more material written for bigger audiences. These two mail categories thus give some reference points for gauging the other categories of articles.

Figure 3 shows the proportion of the pronoun “I” in these different categories of articles. The first point to notice is that the proportions differ dramatically, with “I” being over four times more common in outgoing electronic mail than in articles that called forth no discussion (this difference is well over 10 standard errors).

There are several other notable patterns in these data. First articles that started discussions, either within *msgs* or *topics*, contained 2-3 times as many occurrences of “I” on average as those that did not; the follow-ups themselves had similar rates of occurrence of “I”. In all, the rates for *topics* discussions were almost 75% of the rate for outgoing electronic mail. This result seems to indicate that articles in an informal style are more likely to spark discussion, and further that the discussion is more likely to be informal. The same rule seems to distinguish even the two categories of articles without follow-ups: the *msgs* articles that at least became topics by being “joined” (even if no discussion ensued) use “I” significantly more often, on average, than *msgs* articles

that did not become topics. To summarize, ephemeral interest group postings are more informal than typical *msgs* articles, as we anticipated; but we did not find that ephemeral interest group discussions were significantly more informal than *discussions* taking place on *msgs*.

Increasing the Visibility of Ongoing Discussion Groups

One of our original expectations had been that the *topics* system would decrease follow-up traffic on *msgs*, while increasing the total number of follow-ups, by providing a more appropriate forum for informal discussion. When we compare the counts of follow-up articles in our study of pre-*topics* usage (April–May 1991) to the counts in the more recent study just described (October–November 1991), we see that the total amount of discussion did increase, but that there are only modest indications of a decrease in the follow-up traffic on *msgs* ($p \sim .05$). It appears that *topics*, though an attractive forum for many, is not attractive enough an alternative to draw a large proportion of follow-ups off of *msgs*.

We have already suggested part of an explanation: users had almost no reminder of the existence of topics. Indeed, a survey of *topics* users undertaken in mid-November 1991 showed that most of them were bothered by not knowing what interesting discussions they might be missing by not having joined the original posting on *msgs*. A week later, we conducted a small experiment to assess the extent to which a reminder facility would enable people to learn about important interest groups.

On November 25th an interest group was created to discuss issues related to a new set of management procedures at Bellcore. This group quickly gathered a large number of postings and subscribers, and thus seemed to be one that people might not have wanted to miss. Therefore, we posted

Message 10545:

From mlittman Fri Feb 14 08:52:25 1992

Subject: radio programming: coincidence or conspiracy?

(15 lines) #!topic:736 has 3 postings and 6 subscribers. [ynqj?]

Figure 4 Extension to *msgs*: Sample Prompt.

a reminder to the general *msgs* board the next day (November 26th) stating that a good discussion on this topic was underway and restating the instructions for how to join it. At the time of the posting of the reminder, the discussion group had 67 subscribers and 11 articles. Later we surveyed 41 of the users who joined the discussion group after the posting of this reminder to find out how they had joined the discussion. 23 users replied, giving a response rate of 56%. Of these 23 late joiners, only three had joined the discussion on the basis of the original article, indicating that the lifespan of a *msgs* posting for attracting new interest group subscribers can at times be less than a day. Twenty subscribers joined in response to the reminder, though one stated that he had intended to go back to the original posting to join, but came across the reminder article first, and joined it instead. Thus 19 of the 23 respondents joined solely because of the reminder article, corresponding to 83% of the responding late joiners. Given that there were 65 late joiners in total, we can estimate that 54 joined because of the reminder article. In total, the ephemeral interest group ended up having 132 subscribers, so the posting of the reminder increased the membership by about 69% over what it would otherwise have been.

This indicates that *topics* usage would be greater if users could easily find out about ongoing discussions. In fairness, a few of the late joiners stated that they had joined because they respected the opinions of the person posting the reminder, so we should not expect that increasing *topics* visibility will necessarily cause as dramatic an increase in subscription as we saw in our experiment -- but it is clear nonetheless that the invisibility of discussions is a shortcoming of the current *topics* system.

SUMMARY AND FUTURE DIRECTIONS

We have introduced the notion of *ephemeral interest groups* and described the results of a first experiment in making them available to a large user community. Our findings show that such groups can be supported with minimal modification to a simple bulletin board system; that they appear to require a moderately large user community to be viable; that they can serve as a novel and effective form of information filtering; and that ephemeral interest group discussions are significantly more informal and personal than typical postings to *msgs*. Judging by an overall increase in follow-up traffic, ephemeral interest groups appear to have facilitated discussions that would not have taken place on *msgs*. However, not as much follow-up traffic as we had hoped has migrated off *msgs* into ephemeral interest groups. We believe the invisibility of ephemeral interest groups within *msgs* is an important factor influencing their ability to draw

discussion traffic off of *msgs*, and one of our experiments supports that intuition.

We are currently experimenting with successors to *topics* that address the visibility problem and other shortcomings. One experimental system involves only slightly more obtrusive modifications to *msgs*. If a *msgs* article has already spawned an ephemeral interest group, that information is included in the *msgs* prompt (see Figure 4). In addition, users can directly access a list of current groups, their subscription lists, and their digests, directly from the *msgs* interface. These additional commands are not advertised in the prompt, but are described if the user type '?' for help. The modifications, though more substantial than our original modifications, still do not disturb the original *msgs* functionality for users uninterested in our innovations.

In a second experimental system, called Lynx, we abandon the *msgs* program altogether. A primary goal of Lynx is not only to address usability problems in the *msgs* program, but also to generalize the use of ephemeral interest groups to information sources other than *msgs*. Lynx provides a uniform, network-transparent browser to make information sources in our local environment more easily accessible, and to permit ephemeral interest groups to be spawned from items from any of those sources. In addition to the *msgs* information source, Lynx includes the Associated Press newswire, a calendar of events automatically extracted from *msgs*, and Unix man pages. Lynx provides facilities for browsing active and archived ephemeral interest groups. Finally, it also draws on local information sources to generate profiles of fellow group participants, including their pictures, locations, organizational affiliations, technical reports, etc., so that one can easily get to know other members of the community.

REFERENCES

1. Conklin, J., and Begeman, M. L. (1988). gIBIS: A hypertext tool for exploratory policy discussion. *ACM Trans. Office Information Systems* 6, 4 (October 1988), pp. 303-331. Also in *Proc. 2nd Conf. Computer-Supported Cooperative Work* (Portland, OR, 26-28 September), pp. 140-152.
2. Halasz, F. G. (1988). Reflections on NoteCards: Seven issues for the next generation of hypermedia systems. *Communications of the ACM* 31, 7 (July), pp. 836-852.
3. Hollan, J. D., and Stornetta, W. S. (1992). Beyond Being There. *ACM CHI'92 Proceedings*, pp. 119-125.
4. Monty, M. L. (1986). Temporal context and memory for notes stored in the computer. *ACM SIGCHI Bulletin* 18, 2 (October), pp. 50-51.
5. Nielsen, J. (1990). *Hypertext and Hypermedia*. Academic Press, San Diego, CA.