

Democratic Participation in the Discursive Management of Usenet

John C. Paolillo, SLIS and Informatics
and David Heald, Computer Science
Indiana University, Bloomington
<paolillo@indiana.edu> <dheald@indiana.edu>

Abstract

Internet voting, sometimes proposed as a means of enhancing democratic participation, is partly inspired by the democratic process of newsgroup creation on Usenet. To better understand how online voting might influence democratic participation more generally, we conducted an empirical investigation into the voting activity on newsgroups in the comp hierarchy of Usenet since 1989. Counter to expectation, participation does not appear to be organized into factions or interest groups, but rather there are distinct, individualized patterns of voting. At a coarser level of analysis, some interest-based patterns do emerge, but these appear to correspond to frequent individual voters instead of coherent groups of voters. Noting that the Usenet voting protocol is designed to function principally as a gauge of participant interest, rather than as a genuine plebiscite, we conclude that the design of the Usenet voting system may not adequately gauge the electorate's will in an electronic democracy where voter turnout and democratic participation are chief concerns.

1. Introduction

Of all the claims about the Internet, among the most attractive and simultaneously the most contentious, is the notion that networked communication leads to greater democracy, whether through the open airing of differing opinions in public fora, the dissemination of information while bypassing politically repressive governments, or the promotion of members' collective consciousness. Optimists see a further potential for technology to transform the mechanisms of government itself in the introduction of Internet-based voting technology [1, 14]. Already, the existence of new communications media such as the World-Wide Web has radically affected the discourse around the political process, as participants from political parties and candidates to political action and public interest groups, to social activists, multi-national corporations, news media outlets and independent polling organizations, and even individual citizens all attempt to take part in the political discourse by disseminating their views and organizing citizen participation on the [1, 2, 7]. As more and more people get online, the argument goes,

voters can take advantage of the increased information available to them, and the greater freedom (i.e., the ability to "vote in one's pajamas") and lower costs [11] afforded by online voting will thereby ameliorate past problems of potentially uninformed, apathetic electorates and limited access to the polls on account of location, disability, race, etc. These forecasts raise the expectation of higher election turnouts, including among traditionally disenfranchised groups.

In much this spirit of optimism, the Spring 2000 Arizona Democratic Party primary offered voters the option of casting online and mail-in ballots in addition to more traditional voting options. This experiment was not entirely successful, however, and it was challenged early on by voting-rights groups who feared that problems of differential access to the Internet could lead to de-facto disenfranchisement of minority voters, and that lack of sufficient security measures could compromise the online vote entirely [19]. Interest in online voting is nonetheless high, and given the hairsplitting outcome of the 2000 US presidential election, both in Florida and elsewhere, we can expect increasing public demand for technological guarantees to the exercise of democratic rights.

The appeal to technology as a remedy for democratic ills owes much to the apparently democratic practices found in many online communities. Formal decision-making procedures are found in electronic communities of any sort, from discussion groups and mailing lists [22, 23] to MUDs and MOOs [4, 13]; plebiscites, polls and referenda are found that deal with everything from direct policy-making to elected representation. And while the discursive manifestations of online democratic participation have been probed and critiqued [5, 10, 12] the formal democratic processes of voting have received comparatively little attention. Thus, we have little direct information from which to judge whether the electronic discharge of civic responsibilities will affect the democratic process favorably or unfavorably.

It is with this in mind that we undertook our present study of voting in Usenet newsgroup creation. Rather than recapitulating earlier studies of democratic process, or lack of it, in public discussion on newsgroups, we seek here to study the formal procedure for creating, removing and re-organizing newsgroups. The entire process has many democratic elements, including

formal proposals, public debate, but we will focus primarily on the formal polling of users, an aspect of the process well-documented in public records. Verification tallies of Usenet newsgroup creation votes since 1989 may be obtained from publicly available archives, for over a thousand newsgroups, with millions of votes cast by many thousands of voters. The Usenet voting model is a widely emulated model of online democracy,¹ and Usenet itself is a well-studied form of computer-mediated communication, with input from social science as well as design perspectives [3, 6, 12, 17, 20]. For these reasons, we chose to study voting on Usenet to gain greater insight into the potential and practice of online democracy.

2. Usenet News

Usenet news is a system for supporting persistent, asynchronous discussions which now carries several gigabytes of message traffic per day, around the world. Being available in over 200 countries, it is one of the most common uses of the Internet, and accounts for nearly one third of its data traffic [23]. Usenet is a network of loosely connected server sites with largely informal agreements to share message traffic. The basic means of communication on Usenet is through messages posted to a server, which individually are much like email messages in form and content. Each message is posted on one or more topic-oriented newsgroups, where it can be read by a user with a client program, called a newsreader. There are literally tens of thousands of Usenet newsgroups, hierarchically arranged. Many top-level hierarchies gather together newsgroups that are under common administrative control, e.g. the microsoft hierarchy includes the newsgroups administered by the news administrators of Microsoft Corporation, who decide what groups exist in that hierarchy, and which ones are available to news servers outside their corporate network. At each news server site, news administrators choose which of the available newsgroups to propagate on their own sites, and thus, make available to other sites which depend on them for newsfeed.

The bulk of Usenet message traffic is carried in the “big eight” hierarchies (comp, humanities, misc, news, rec, sci, soc, talk), which were created in 1986 to replace an older, less organized naming scheme [22]. Initially, newsgroup creation, naming and removal in these hierarchies was regulated by the consensus of the “backbone cabal”, as it was called, a loose-knit group of administrators whose sites carried the most Usenet traffic, and who provided the feed for most other Usenet

sites. The establishment of the big seven hierarchies itself was an act undertaken by the backbone cabal. Perhaps somewhat predictably, large conflicts arose around that time over the newsgroup naming and creation policies of the cabal, particularly when a number of popular proposed newsgroups were turned down, such as rec.music.rock-n-roll, rec.drugs and soc.sex [22], for fear that carrying them would incite institutions to revoke support for the relatively vulnerable Usenet. These conflicts led first to the creation of the alt hierarchy, originally distributed through connections independent of the Usenet backbone, and later to the revamping of the newsgroup creation and naming policies in the big seven hierarchies.

3. The Newsgroup Creation Process

The revised newsgroup creation process, which remains in place today, formalized an earlier process where a proposal, a discussion and a popular vote preceded the creation of a new newsgroup. The process is initiated when a group of people, possibly communicating on an existing newsgroup, or on an outside mailing list, perceive a need to define a new discussion space. A motivated individual must then post a “Request For Discussion” (RFD) to the moderated newsgroup news.announce.newgroups. The RFD must specify the proposed name of the new newsgroup, the moderation policy, and a rationale for its creation. It should also include a charter, which is a formal description of the intended purpose of the newsgroup and the topics which are appropriate for it.

The RFD is typically cross-posted to other related newsgroups, and may be circulated on mailing lists where people have an interest in the proposed newsgroup. Following the RFD posting, a period of discussion lasting at least three weeks ensues. Discussion may take place on news.groups, other appropriate newsgroups etc., suggesting changes to the charter, the proposed newsgroup name, the moderation policy or other circumstances affecting the newsgroup’s creation. Following this period of discussion, a “Call For Votes” (CFV) is posted by the Usenet Volunteer Votetaker, an individual volunteer who oversees all aspects of the polling. Polling is generally conducted by email, with participants able to vote “yes”, “no”, or “abstain”. At the end of the polling period, the votes are tallied by the volunteer votetaker, who posts a RESULT article to news.announce.newgroups, in which the votes of each participant are listed, by name and email address, for verification purposes. If the vote tally has at least 100 more “yes” votes than “no” votes, and there are more than two thirds “yes” votes, then the vote passes, and the newsgroup is created. If either of these conditions is not met, then the group is not created, and

¹ Other voting systems discussed in the CMC literature include those of certain MOO’s [4, 13], online newspaper polls [21], and Group Decision Support Systems [15, 25].

a six-month waiting period must pass before another proposal can be made on the same topic.

The formal process of newsgroup creation on the big seven hierarchies is widely emulated as a model for Usenet discussion management which encourages public participation. It places a substantial part of the burden of decision-making on those who actually use the affected newsgroups, thereby protecting Usenet administrators from undue criticism when changes are made. And it has distinctly democratic elements, with its public debate, formal vote, quorum requirement and two-thirds majority rule. However, there are sharp differences between this form of democracy and the more usual sort. The polling procedure is intended to involve only those who have an interest, whether positive or negative, in the potential newsgroup's creation. It is not expected that all Usenet netizens will inform themselves about and vote on each of the newsgroup proposals that arises. The intent is that "yes" votes should be cast by those who actually anticipate using a newsgroup, while "no" votes should be cast by people who object to its existence for some reason (not merely because they wouldn't use it).

Consequently, in any newsgroup creation, one expects a vast majority of users who have no interest, and hence do not participate in the vote. The function of the vote, therefore, is not so much one of a democratic referendum as it is a simple interest poll, with an opportunity for opponents of the change to voice their opinions. If the interest is weak or if opposition is substantial, then the change is blocked [22, p. 399-400]. All the same, the open nature of the process permits ordinary users to take an active role in the discursive management of Usenet, on more than just a single issue. Thus, a question arises: does users' participation in newsgroup creation votes reflect more the local nature of interest polls, or do Usenet users show evidence of becoming engaged netizens, organizing themselves around particular issues? Do newsgroups get created in isolation from one another, or do users attempt to actively direct, through participation in newsgroup votes, the development of the discursive space that is Usenet?

4. Voting on the comp Hierarchy: Method

To investigate these questions, we downloaded a corpus of "newsgroups" messages from the archive at <ftp://ftp.isc.org/pub/news.announce.newsgroups/>. These archives contain the charters for all newsgroups proposed in the big seven hierarchy from 1989 to the present date (April 2001, at the time of this research), as well as the RFD and CFV postings, and the final RESULTS for each new newsgroup referendum. To focus our inquiry, we selected the comp hierarchy, which hosts a broad range of computer-related discus-

sions. For this time period, 390 newsgroup polls were processed in the comp hierarchy, in which more than 88,000 users participated and more than 117,000 votes were cast.

We then extracted information about the final votes for two stages of analysis. First, we displayed the size, outcome and date of each vote graphically, so that overall trends in newsgroup creation on the comp hierarchy could be detected. We then constructed a database containing each of the votes in the corpus, so that a social network analysis of the voting patterns could be conducted. From the database, a table was constructed in which individual voters (identified by unique email addresses) were listed in rows, newsgroups were listed in columns, and the intersection of the two was recorded as 1 if a given voter voted in the poll for that newsgroup, and 0 if not. While our database records yes, no abstain and invalid votes distinctly, we have not yet been able to use that information on account of complications in the statistical analysis of the data; consequently, our analysis principally addresses the issue of voter turnout, and not the question of what political factions, if any, emerge from the polling data. Factor analysis, a procedure that identifies major dimensions of variation in a collection of variables, was performed to identify social networks of users voting on similar newsgroups and clusters of newsgroups having similar voter turnouts [16, 17, 18]. Eigen analysis was used to assist the identification of factors. In addition, a second factor analysis was performed in which newsgroups were gathered into categories of related topics; in this way, different views of the data were obtained, from which different patterns could emerge.

5. Results

Figure 1 displays the plot of newsgroup polls over time, from March 1991, the earliest vote on record in the comp hierarchy, to April 2001. In this plot, the vertical axis represents the overall number of votes; the largest poll consisted of 1502 votes. The shading of each data point represents the proportion of yes and no votes, with black representing all or predominantly yes votes, and white representing all or predominantly no votes. From this display, several facts become apparent. First, the overwhelming tendency is that a proposed newsgroup creation on the comp hierarchy results in passage of the proposal and creation of the newsgroup. Only in cases where the overall vote count is small do we find a predominance of no votes, and the failure of the newsgroup creation proposal. Moreover, we do not see a trend over time of larger or smaller newsgroup polls, or of increased or decreased probability of passing outcomes. Neither is there a trend for a larger number of votes taking place over time, as one might expect with the explosive increase in Usenet

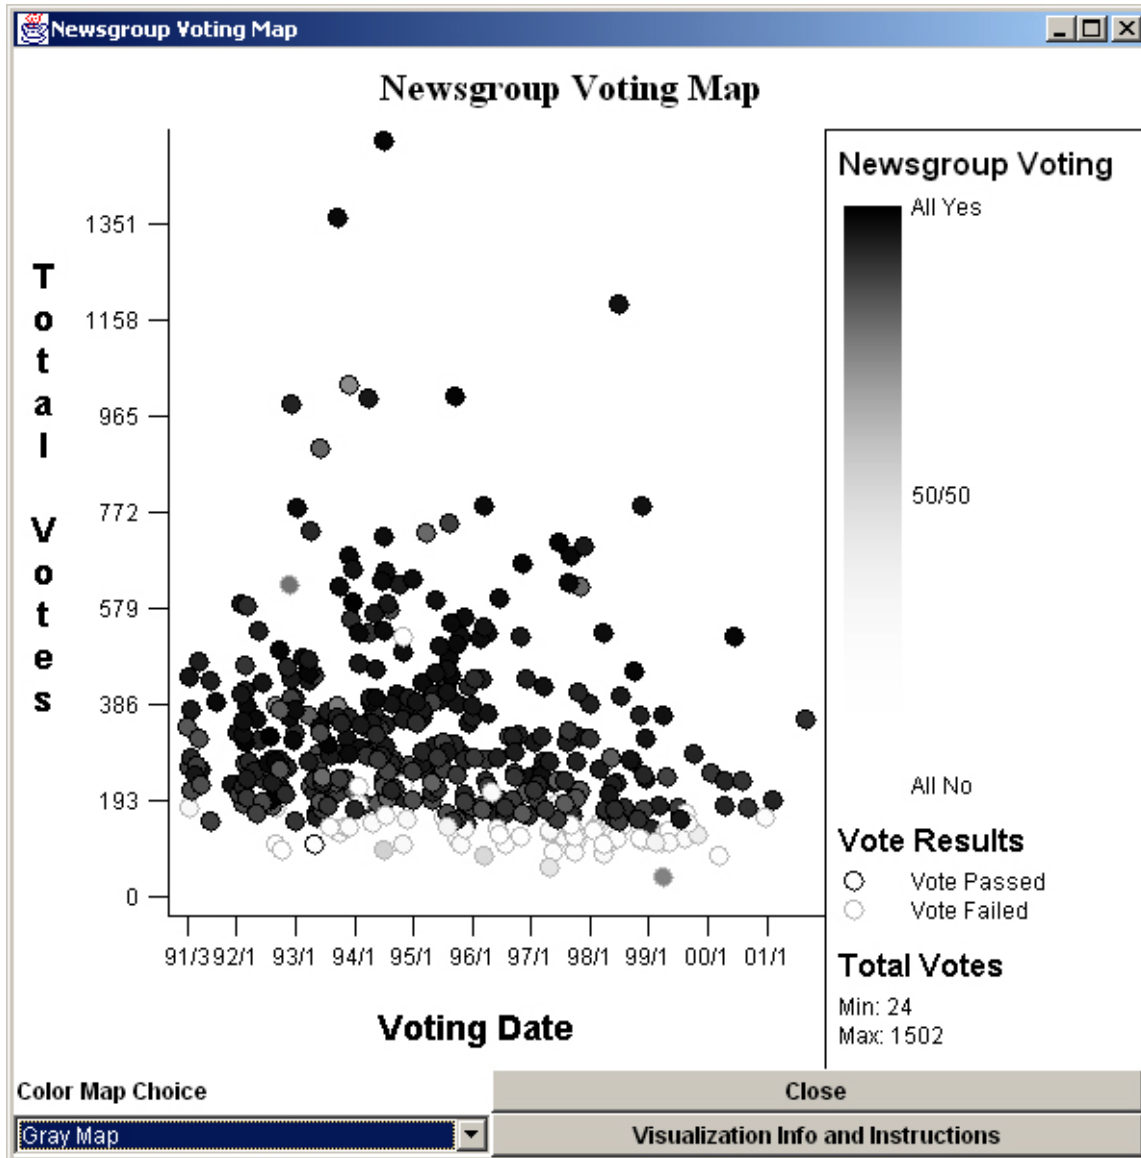


Figure 1. Outcomes of newsgroup creation polls in the comp hierarchy since March 1991.

readership over the same period [23]. Thus, newsgroup creation in the comp hierarchy seems to be a relatively stable process.

Further analysis reveals that single-issue voting predominates, with only a very small number of voters participating in a substantial number of newsgroup creations. As can be seen from Figure 2, where the number of participants who vote in any given number of newsgroup creations is plotted, the overwhelming majority of voters (76,735) participates in one and only one newsgroup creation vote. Those participating in more than one vote on the comp hierarchy from 1989 to 2001 comprise only one third of the voters in that

period. Voters involved in more than one tenth of the votes (i.e. 47 or more) comprise only 1.5% of the voting population.

A social network analysis was conducted using Factor Analysis [24], a technique for reducing the dimensionality of large data sets. The factor analysis was used to identify collections of newsgroups which shared populations of voters. Newsgroups whose voting lists correlate with one another should emerge as underlying "factors" in the analysis. Since factor analysis essentially models the original data as a combination of these abstract factors, one must determine a number of factors to extract that provides suitable economy of




Figure 2. Number of participants in newsgroup creation votes on the comp hierarchy by voting frequency (log-log plot).

description without changing the character of the data. This was accomplished by conducting an eigen analysis of the correlation matrix, and employing Bartlett's chi-square test to select the number of factors [24, p. 280]. To simplify the procedure, the table was constrained to the 46 newsgroups with the largest votes (used in two sets of 23 newsgroups each), and voters voting four times or more. In both factor analyses, 20 factors needed to be retained: almost every newsgroup loaded on its own distinct factor. In other words, the correlations among any of the newsgroup votes are very low, so that each newsgroup creation vote is largely a distinct event, with minimal overlap among participants in any two newsgroup votes.

It is possible that in this procedure, the presence of a large number of low-frequency voters may have swamped out observable correlations in the turnout of higher-frequency voters, who might be argued to constitute a different population entirely. To test this hypothesis, a series of similar analyses were conducted, with each new analysis removing one or more ranks of low-frequency voters from the data pool prior to computing the correlation matrix. If high-frequency voters have significantly different turnout behaviors from the lower-frequency voters and congregate into groups which attend distinct classes of newsgroup creation events, then at some point a markedly different pattern with a much smaller number of factors should emerge. To test this hypothesis, seven successive analyses were conducted on the 23 largest newsgroup votes, beginning with voters who voted 7 times or more, and ending with those voting 20 times or more.




Figure 3. Eigen values of the correlation matrix at seven different values of the cutoff of number of votes cast.

By way of illustration, the eigen values for each of these analyses are plotted by factor in Figure 3. Eigen values close to one represent factors that account for the variation of approximately one of the original variables in the analysis; here, the original variables are newsgroup votes, so an eigen value of one would correspond to the variance of approximately one newsgroup's voting pattern. Eigen values much larger than one suggest a degree of economy, so that an eigen value of five suggests that a given factor accounts for the variance in approximately five newsgroup votes. In Figure 3, the largest eigen value is approximately 4.5, which comes from factor 1 in the analysis of voters with more than 20 votes. This data set includes only 3.4% of the votes cast, and 117 of the original 88,394 voters. The apparent economy of this model is both modest on the one hand, since we must retain 16 factors out of an initial 23 variables under Bartlett's test, and costly on the other, because of the loss of 96.6% of the relevant data. Thus, the apparent economy in factors, which we intended to represent representing interest groups, is explained by the paucity of the remaining data.

We conclude from this that there is no reasonable cutoff between high and low frequency voters which reveals the hypothesized interest groups. Our initial finding is thus confirmed: newsgroup creation events are individual events, which engage only those members of the Usenet readership who are interested in the outcome, and there is no significant overlap among the turnouts for any two creation events.

Category	# groups	Category	# groups
sys	55	binaries	8
lang	44	emulators	7
other	37	ai	7
graphics	30	arch	7
os	29	mail	6
soft-sys	22	security	6
databases	22	unix	6
dcom	21	org	6
infosystems	15	programming	5
bbs	12	society	4
cad	11	compilers	3
windows	11	groupware	3
software	10	object	3

Table 1. aggregated newsgroup categories

One possible reason for this outcome might be a fault in the analysis procedure. Newsgroup votes represent small samples of the relatively large number of voters. Consequently, the variance of these samples would be rather large, and correlations may not emerge unless something is done to “smooth” over the gaps in the data caused by low sample sizes. The easiest way to accomplish this smoothing is by aggregating the votes from different newsgroup creations into a smaller number of sets that may be expected to behave differently from one another. This move, while it allows the search for meaningful patterns to proceed, also changes the nature of the research question addressed. Instead of asking whether individuals have a common interest in voting on specific newsgroup creations, we now only ask if they have a common interest in certain general types of newsgroups.

The aggregation was accomplished in two steps. First, we identified the sub-hierarchies in the names of each of the 390 newsgroups, to arrive at an initial set of 52 categories based on the sub-hierarchies alone. Next, newsgroups in categories with 3 and fewer members were re-classified into other categories, based on the full names of the newsgroups therein. This procedure resulted in 25 named categories plus “other”, for a total of 26 categories of newsgroups, listed in Table 1.

As before, a table with columns for the newsgroup categories and rows for individual voters was created. Each cell recorded the number of times a voter voted on a newsgroup creation in a given category. Voters who voted seven times or more were included in this data set, as the resulting matrix could be expected to be both manageable in size, as well as representative of the complete patterns of the data, as determined by the prior analysis. This matrix was then subjected to an eigen analysis and a factor analysis, much as the prior data set was.

The eigen analysis for this data set proved to be markedly different from the prior analysis. Whereas the

1. Applications (217)	5. Implementation (10)
cad	compiler
database	emulator [C]
soft-sys	
binaries	6. Mail (6)
os	mail
bbs [7,8]	
sys [8]	7. Unix (6)
graphics [2]	unix
dcom [3]	8. Groupware (3)
arch [5,C]	groupware
2. Programming (49)	9. Windows (26)
programming	windows
lang [1,4]	infosystems [1,6]
3. Security (12)	A. Society (4)
security	society
org [1,2]	
	B. AI (7)
4. Object-Oriented (13)	ai
object	
software [1,6]	C. Other (37)
	other [1,7]

Table 2. Factors of newsgroup aggregates

first eigen value in the comparable earlier analysis was approximately 2.5, suggesting a very modest level of economy to be attained, in the aggregated analysis, it is greater than 10, suggesting a more useful level of economy. Twelve factors are significant according to Bartlett’s test, meaning that simplification of the aggregates, i.e. dropping about half of the original variables, is desirable. The corresponding factor analysis produced the twelve factors identified as (1) through (9) and (A) through (C) in Table 2. Many aggregates showed significant correlations with more than one factor. Hence, each aggregate is listed under the factor for which it had the highest loading, while aggregates having loadings greater than 0.3 on another factor have that factor listed in brackets alongside them. Each factor is headed by an interpretive label, which we have assigned to characterize the factor as a whole.

A number of the aggregates show little correlation with any of the others, loading on their own distinct factors, such as Mail, Unix, Groupware, Society, AI and Other. Of these, only AI and Society entirely lack correlations with the other aggregates (the number of newsgroups in these aggregates is small, which may be the reason they show no strong correlations with other aggregates). The remaining aggregates are all directly or

indirectly inter-correlated to some degree. In terms of voting practices, this means that voters voting in one category show a tendency to vote in all or any of the others. At the same time, some inter-correlations are stronger than others, thereby structuring the inter-relations among the aggregates. The structure of the inter-correlations becomes apparent when they are mapped graphically, as in Figure 4, where the factors are denoted by ellipses, and the aggregates that are correlated with more than one factor are indicated a dot in the factor on which they load highest, and a line extending to the other factor(s) they load on.

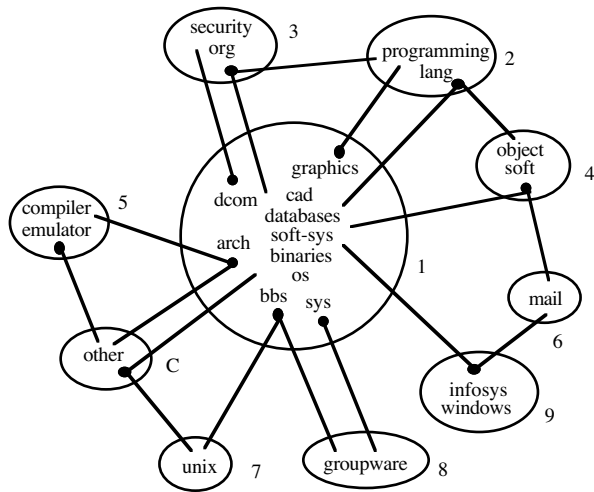


Figure 4. Inter-correlations among the factors of newsgroup aggregates.

Some of the inter-correlations among the aggregates appear to identify associative links of topical similarity. For example, the *bbs* aggregate links Groupware (8) and Unix (7) with Applications (1): BBS systems are applications that typically run on Unix systems and offer groupware functionality. Likewise graphics correlates more strongly with Applications (1), but still significantly with Programming (2); most of the graphics newsgroups voted on in this period are indeed intended for applications such as Photoshop, Freehand, and other major commercial titles. At the same time, there are votes on newsgroups for graphic algorithms, OpenGL and gnuplot, where more programming is involved. Other correlations, such as that between infosystems and windows (principally X-windows, not MS) on the one hand (factor (9)), or mail (6) on the other, are harder to explain, since there is no apparent overlap in the newsgroup topics of the different aggregates.

As the last step in the exploration of the newsgroup voting data, factor scores were calculated for newsgroup participants based on the first three factors, Application, Programming and Security, and these scores were plotted in three dimensions, so that clusters of voters sharing turnout patterns might be identified. The resulting plot is given in Figure 5.




Figure 5. Factor scores of 789 voters on the first three factors.

The data cloud in Figure 5. is a fairly uniform ellipsoid, with a small number of points in three tails, corresponding to each of the three main dimensions. Unfortunately, the number of voters in these tails is very small, and consequently no real clusters of voters can be said to exist. In addition, all of the voters in these tails are among the most active in the data; all of them have a total of 35 votes or more, that is they participated in 9% or more of the newsgroup creations under study. Also, it is not clear what groups the voters together in the tails. For example, the most extreme points in each of the lower left and upper right tails are actually points that correspond to the same individual. The most frequent voter used two different email addresses between 1989 and 2001, as can be discerned from the addresses (they are nearly identical). Assuming a relatively constant voting pattern, we should expect to find that these two addresses would behave similarly. However, they do not, and the fact that they fall in the extreme range of two different axes suggests that they are completely uncorrelated. One possible explanation for this anomaly would be if the one address had been used at an earlier time while the other was used later, and different types of newsgroup votes occurred in those two periods of time. If this were the case, then it could

happen that the two addresses, representing the same person's voting pattern, would appear to be uncorrelated. Alternatively this individual might have changed newsgroup voting tastes along with the change of email address, thereby making them uncorrelated. Apart from these two possibilities, we have no ready explanation for this observation.

To summarize, the social network analysis of newsgroup voting patterns has shown that the votes of individual newsgroups are largely uncorrelated with one another. Repetitions of the eigen analysis procedure in which smaller and smaller subsets of the data were analyzed failed to find a distinct pattern of vote turnout among more frequent voters, confirming this result. While this suggests that newsgroup votes are largely separate events, it is possible that the sample sizes are too small, and hence the variance of the samples too large, to meaningfully correlate them. Hence we explored an alternative strategy in which the newsgroup votes were first aggregated into 26 categories of newsgroups before conducting the factor analysis. The resulting analysis shows a greater degree of inter-correlation among the aggregated newsgroup votes, having a structure that resembles the topical relationships among the newsgroups in the aggregates. We then searched for clusters of voters that would turn out for newsgroup votes in common, but found only a very small number of such voters, all of whom were more active in newsgroup creation votes more generally.

6. Discussion

Viewed as a democratic exercise, voting on the creation of Usenet newsgroups might appear a disappointment. We do not find evidence for an engaged electorate; the vast majority of voters vote only once in the period under study. Other typical signs of democratic processes, such as the emergence of political factions or interest groups, are also not in evidence. Instead, it appears that individual voters attend to a relatively narrow set of issues. However, it is important to recognize that newsgroup creation votes serve a very limited purpose that is rather different from the nature of democratic votes. Spencer and Lawrence [22] describe the Usenet poll thusly:

“The main purpose of the poll was an interest vote. The ‘yes’ versus ‘no’ nature of the polling was partly intended to obfuscate the interest poll under the guise of a democratic vote and partly intended to give administrators the ability to refuse a group en masse if they found sufficient problems with the proposal, like with the name of the group.” (416)

“While in a real democracy it is the duty of every citizen to become informed about and vote on every matter, in a Usenet group creation poll it is your duty to vote only on proposals for groups that you intend to use or with which you have significant problems. To register sympathy support for a group that you do not intend to read distorts the measure of interest that the poll is primarily trying to identify.” (400)

In other words, the poll serves merely as a measure of interest, and excess participation is undesirable and not encouraged. This description fits our findings rather closely. At the same time, it is somewhat remarkable that the Usenet poll has been so widely emulated as a model of online democracy. If the ideal of democratic participation is the involvement of the public in all important matters of policy, then it would seem ill-advised to model the process of participation on an interest poll that is intended to garner the participation of only the select few who are already interested. Either there must be an error in the choice of such a model, or there is more democracy in the process than we are presently told. The question merits a closer look.

In one sense, there are aspects of Usenet which are “elemental” to democratic participation. Citizens of democratic societies typically enjoy freedoms of association and speech, in addition to the formal responsibilities of voting. Without these freedoms, it is often said, democracy could not exist. Usenet offers similar freedoms to an individual. One can associate relatively freely with others by choosing to read and participate in different newsgroups, and at the same time, one is typically free to speak in various ways, by posting to whatever newsgroups one wishes. Moderation policies, local access policies, etc. may affect the degree to which an individual can associate or speak in any instance, of course, but in general Usenet is open in these ways. By virtue of being open, Usenet encourages the participation of its users, and thus is democratic in that sense.

Democratic societies require more than handouts of empowerment to citizens, however; there are responsibilities that must be fulfilled for the society to function. In Usenet newsgroups, this means, among other things, the observance of discursive norms of behavior, so that one person's exercise of freedoms of association and speech do not trample those of others. These norms are sometimes codified as formal or informal codes of netiquette, moderation policies, topic appropriate posting guidelines, etc. Enforcement of these codes requires netizen participation, through collective cooperative action, through formal moderation, and in other ways. Again, Usenet must draw on those who use it for these functions to be carried out.

One important such function that must be fulfilled is the maintenance of the namespace of newsgroups. For users to exercise the rights of association and speech, there must be appropriate newsgroups for them to associate and speak on certain topics. If members of a group wish to participate on a topic that is not considered appropriate by others participating on the same newsgroup, then there must be a mechanism whereby a new newsgroup can be created for that topic, so that users can re-assert their rights to associate and speak without coming into conflict. Moreover, that process must be under the control of the users to some degree, so that they can directly guarantee the exercise of their rights to association and speech. It is this function that the newsgroup creation poll fulfills. The netizen-initiated Request for Discussion,² its specification of a charter for the new newsgroup, and the refinement of the proposal in public discussion are important steps that make the process democratic. The interest poll itself is merely the final phase of a process that is otherwise democratic.

The form of democracy encouraged by these practices is not the plebiscite form of democracy, but rather a Habermasian diverse plurality of democratic communities [8]. Consequently, Usenet probably cannot provide a suitable model for the design of online plebiscites. In particular, there is not the level or consistency of turnout in Usenet-style voting to provide the level of mandate required in a general plebiscite. Instead, the Usenet voting process provides a measure of interest that is useful for system administrators to decide whether to commit collective resources (disk space, a name, and the time required to perform the maintenance) to a group of users interests. Arguably, general elections also perform an interest-polling function; at the very least, the declining poll numbers in US elections are often interpreted that way. And whereas the plebiscite form of democracy runs the risk of imposing the will of a potentially tyrannical majority with views irrelevant to the policy in question [8], an interest-gauging system runs a risk of allowing a disinterested majority to abdicate responsibility for sharing in the determination of the public will, in favor of pursuing other interests. On Usenet, when a committed group of people desires a particular change, they can generally precipitate it, because of a lack of opposition from disinterested parties. The two models of democracy could hardly be more different.

² The RFD bears a strong resemblance to the RFC ("request for comments") process by which engineering standards were discussed and adopted on the Arpanet [9]. Usenet was originally a substitute for the Arpanet for sites that could not afford its access costs, and was influenced greatly by its models for self regulation.

7. Conclusions

The formal democratic practices of Usenet are designed in a way that accommodates its special purposes in the Usenet discursive and social context. These features of its design are reflected in the patterns of participation it encourages. Furthermore, the practices seem to work effectively in that context. For online democracy to be successful in other contexts, it seems that a similar level of consideration would need to be given to the purposes it is supposed to accomplish. Just as the relatively informal processes of public debate and political activism have moved increasingly onto the World-Wide Web and other electronic media, it seems inevitable that electronic versions of the formal processes of democratic responsibility will eventually become much more widespread than they are at present. If these mechanisms of voting and governance are to serve their intended functions, then it will be necessary to design them in ways that are compatible with those functions. In this design process we must take into account considerations of a nature that are entirely apart from the issues of access and security that dominate the debate around online voting today, but which are no less fundamental to the meaningful exercise of democratic rights.

References

- [1] Åstrom, J. 2001. Should democracy online be quick, strong or thin? *Communications of the ACM* 44.1. 49-51.
- [2] Becker, T. Rating the impact of new technologies on democracy. *Communications of the ACM*, 44.1. 39-43.
- [3] Baym, N. 1995. The emergence of community in computer-mediated communication. In Steven G. Jones, ed. *Cybersociety: Computer-Mediated Communication and Community*. Thousand Oaks, CA: Sage. 138-163.
- [4] Cherny, L. 1999. *Conversation and Community: Chat in a Virtual World*. Stanford, CA: CSLI Publications.
- [5] Coste, R.L. 2000. Fighting with speech: David Duke, the Anti-Defamation League, online bookstores, and hate filters. *Proceedings of the 33rd Hawaii International Conference on System Sciences*. Los Alamitos, CA: IEEE Computer Society.
- [6] Donath, J.; K. Karahalios, and F. Viegas. 1999. Visualizing Conversations. *Proceedings of the 32nd*

Hawaii International Conference on System Sciences. Los Alamitos, CA: IEEE Computer Society.

[7] Dutton, W.H., A. Elberse and M. Hale. 1999. A case study of a netizen's guide to elections. *Communications of the ACM* 42.12. 49-54.

[8] Ess, C.M. 1996. The political computer. In C. M. Ess, ed., *Philosophical Perspectives on Computer-Mediated Communication*. Albany, NY: SUNY Press. 197-230.

[9] Hafner, K. and M. Lyon. 1996. *Where Wizards Stay up Late*. New York, NY: Simon and Schuster.

[10] Herring, S. 1993. Gender and democracy in computer-mediated communication. *Electronic Journal of Communication* 3.2. Reprinted in R. Kling, ed., *Computerization and Controversy*, 2nd ed. New York: Academic Press (1996). 476-489.

[11] Larsen, K.R.T. 1999. Voting technology implementation. In *communications of the ACM*, 42.12. 55-57.

[12] MacKinnon, R.C. 1995. Searching for the Leviathan on Usenet. In Steven G. Jones, ed. *Cybersociety: Computer-Mediated Communication and Community*. Thousand Oaks, CA: Sage. 112-137.

[13] Mnookin, J.L. 1996 *Virtual(y) Law: The Emergence of Law in LambdaMOO*. *JCMC* 2.1 , <http://www.ascusc.org/jcmc/>

[14] Mohen, J. and J. Glidden. 2001. The case for Internet Voting. *Communications of the ACM*, 44.1. 72-85.

[15] Nunamaker, J.F.; R.O Briggs, D.D. Mittleman, D.R. Vogel, P.A. Balthazard. No date. *Lessons from a Dozen Years of Group Support Systems Research: A Discussion of Lab and Field Findings*. <http://mies.cs.depaul.edu/research/JMIS.html>

[16] Paolillo, J.C. 1999. The virtual speech community: Social network and language variation on IRC. *Proceedings of the 32nd Hawaii International Conference on System Sciences*. Los Alamitos, CA: IEEE Computer Society.

[17] Paolillo, J.C. 2000. Visualizing Usenet: A factor analytic approach. *Proceedings of the 33rd Hawaii International Conference on System Sciences*. Los Alamitos, CA: IEEE Computer Society.

[18] Paolillo, J.C. 2001. Language variation in the virtual speech community: A social network approach to Internet Relay Chat. *Journal of Sociolinguistics*. Oxford: Basil Blackwell.

[19] Phillips, D.M. and H. A. von Spavosky. 2001. Gauging the risks of Internet Elections. *Communications of the ACM*, 44.1. 73-85.

[20] Sack, W. 2000. Discourse diagrams: Interface design for very large-scale conversations. *Proceedings of the 33rd Hawaii International Conference on System Sciences*. Los Alamitos, CA: IEEE Computer Society.

[21] Schultz, T. 1999. Interactive Options in Online Journalism: A Content Analysis of 100 U.S. Newspapers. *JCMC* 5.1. <http://www.ascusc.org/jcmc/>

[22] Spencer, H. and D. Lawrence. 1998. *Managing Usenet*. Sebastopol, CA: O'Reilly and Associates.

[23] Smith, M.A. 1999. Invisible crowds in cyberspace: Mapping the social structure of the Usenet. In M.A. Smith and P. Kollock, eds., *Communities in Cyberspace*, 195-219. London: Routledge.

[24] Tacq, J. 1997. *Multivariate Analysis Techniques in Social Science Research: From Problem to Analysis*. Thousand Oaks, CA: Sage.

[25] Turoff, M.; S.R. Hiltz, M. Bieber, J. Fjermestad, and A. Rana. 1999. Collaborative Discourse Structures in Computer Mediated Group Communications. *JCMC* 4.4. <http://www.ascusc.org/jcmc/>