Managing Large Scale On-line Discussions: Secrets of the Open Meeting*

Roger Hurwitz and John Mallery

The Artificial Intelligence Laboratory
Massachusetts Institute of Technology

Abstract. The possibility of communities steering themselves by large scale, online meetings are limited by support system capabilities for managing their discourse. An asynchronous conferencing system the Open Meeting supported a meeting in which thousands of U.S. government workers discussed reengineering government services. It enabled users to learn and share opinions about proposed changes, because its design focused on decomposing the inflows of comments, structuring multi-lateral conversations and maintaining civility. Access was provided over SMTP and HTTP to a topically differentiated hypertext synthesized from an object database, which was extended by users' comments. The comments composed virtual conversations, structured by a discourse grammar that constrained what types of comments could be attached in specific contexts. Civility benefitted from use of moderators and from participants' having similar bureaucratic culture backgrounds. The results suggest that cultural diversity is a greater challenge than numbers to enabling effective on-line democratic action.

Can the Internet invigorate the democratic practices of large communities? Will it become a new public sphere[1], a space where people gather to rationally discuss issues in the hopes of reaching consensus and influencing policy makers? Expectations for that have been raised by a) the diffusion of on-line community forums and newsgroups, b) the power of the Internet to quickly and cheaply gather dispersed individuals around common interests, and c) its use by governments and political activists to distribute information. Parties, candidates and even rebellions now have web sites and email. As expressed in the ugly buzzword *disintermediation*, digital networks can create direct links between people and government, bypassing the mass media, interest groups and other gate keepers.

* The on-line meeting described herein was a collaboration of the Artificial Intelligence Laboratory, the National Performance Review, Lawrence Livermore National Laboratory and Mitre Corporation. Jonathan Gill, Thomas Kalil, Randy Katz and Howard Shrobe provided more support. Research was partly supported by the Defense Advanced Research Projects Agency of the Department of Defense under contract number MDA972-93-1-003N7.

[1] The term public sphere was brought to contemporary attention by Habermas [1], who links the decline of public spaces for political discussion to the rise of mass media.

Some argue that the Internet will not nurture a public sphere, which, in any case, is an inappropriate model for democratic practice. Its earlier manifestations in Berlin salons and American civic associations were unrepresentative and exclusivist; its rhetoric, abstractions and self-evident truths assume an underlying homogeneous society, steered by white, literate and propertied males[2]. On this critical view, characteristics of on-line communications as well as the growing diversity of voices on the Internet prevent the realization of a public sphere, even one accessible to all and where difference is acknowledged. These include a bias toward speedy, unreflected responses, the absence of social context cues and hyperlinking which, pointing a reader elsewhere, dissolves the authority of the text at hand. On-line political discussions will be situated, personal, and chaotic; more likely to immobilize a community than help form its collective will.

This paper joins the debate by examining how the Open Meeting an experimental network-based information system brought together many, dispersed people to discuss government policies, which affected their work lives. The system's architecture and its theory of knowledge representation have been described elsewhere[3]; the present focus is on how the system organized the large scale discussions, particularly with regard to managing inflows, structuring interactions and maintaining civility. The achievements and shortcomings of the experiment are examined for their relevance to issues of on-line governance. Clearly just broadening people's access to one another and to information does not generate participation, order discourse or produce group decisions. So it is important to show how it is possible to get some of these results.

1 Overview of the System

The Open Meeting system was initially designed and implemented in fall, 1994, to support an on-line meeting, in which several thousand United States government workers conducted virtual discussions of recommended changes in government operations and performance standards. At the system's core is a persistent object database, which coupled to a task engine, supports email and web servers. These present messages and pages, with content generated on-the-fly by the database, in response to users' inputs through queries, remote commands or interactive forms. Inputs can include users' comments which are added to the database and reconfigured as threaded conversations or annotation overlays. The recommendations for discussion had been prepared by the National Performance Review (NPR),[4] a unit in the Office of United States Vice-President Al Gore, with the mission of "reinventing government." It believed a large scale on-line meeting would

- Increase the visibility and support for these recommendations;
- Enhance their prestige through association with new technology;
- Gather reports on promising practices in line with them;
- Bypass organizational boundaries that keep employees in different organization from discussing common work problems;

– Overcome geographical barriers that confine discussions on government to the capital;
– Demonstrate the possibility of on-line/ electronic governance.

As the last point suggests, in the NPR view, government workers comprised a community whose opinion needed to be articulated and mobilized. The NPR project team[2] for the meeting was guided by the metaphor of a New England town meeting, where any member can of the community can speak to issues of common concern. It envisioned an on-line environment analogous to a town hall, that would include chat rooms for an assembly hall, committee rooms and water cooler.[3]

The trajectory toward a community activation system changed with the involvement of the Open Meeting developers, members of an MIT group, whose research concerned intelligent networks that route and enhance messages, based on understanding of their contents and the organizations through which they flow. At the time, the research group was planning a web based system that would manage public participation in government inquiries on proposed changes in regulations. The initial plan called for interested parties to attach their views as annotations to the proposals, differentiated according to the types of comments being made, e.g., question, statement of support. Officials could then retrieve these comments by target and type from a textbase of review documents, using the system, in effect, as an annotation server. Alternatively, the comments on a document could be routed directed to officials whose profiles indicated their in that document. By enabling the public and officials to receive and reply to each other's comments, the system would support virtual discussions composed of threads of typed comments. Through the prism of this plan, the federal workers appeared like individual stakeholders responding to inquiries on proposed changes in regulations affecting their individual interests, work rules and performance standards. They were primarily sources of information about the benefits, feasibility and acceptability of the proposals.

To encourage broad participation and knowledge sharing, regardless of organizational differences, we consequently advised that the meeting focus on eleven NPR reports, each of which concerned reinventing an operating system, like human program development or information management, found in all departments and agencies and subject to the same rules and standards. These could provide common grounds for discussions across organizational boundaries, but recommendations for a specific department were unlikely to draw workers from other departments into a discussion. We also advised dropping the town hall format for several reasons:

[2] Led by Larry Koskinen and Andy Campbell

[3] The Clinton-Gore team used face-to-face "townhall meetings" 1992 presidential campaign and later in the Clinton I administration to get public attention for its issues and positions, while bypassing the national mass media. This populist strategy was highly successful with citizens and voters, but soured the team's relationships with the journalists, editors and owners of the media.

– a reluctance to assume the participants already recognized themselves as a community;
– the desired dispersion of participants over time zones made asynchronous communications more feasible;
– chatrooms, MUDs and MOOs would not scale for the desired number of participants;
– email would not support the GUIs which could suggest a town hall setting, and email was mandatory because relatively few workers then had web browsers.

The group proposed instead a system similar to the public comment system sketched above, as an application layer for Comlink, an enhanced document distribution system, developed primarily by Mallery and used for the indexing, publication and retrieval of White House electronic releases on the Internet.[5]. The principal features and functionalities of the resulting system included

– *An extensible, persistent object database* whose instances include texts and users. The text instances include hypertext links and indices, and so support extension of an initial textbase configurable as hypertext; arbitrarily constrained, on-the-fly construction of local hypertext; and views of the hypertext, restricted by arbitrary indices, like quality ratings, security classifications, etc. The user objects can include slots for interests and preferences, enabling subscription to and routing of texts.
– *Dynamically generated web pages* which facilitate representation and navigation of the online discussions and collaborations. Regions of the resulting hypertext can be isolated; pointers to comments and annotations descending from an initial text can be displayed in an outline or tree, each with an icon for the type of the comment. This permits both overviews of the link structure and direct access to those comments which interest the user.
– *Interactive form processing* to support user queries and comments, hypertext traversal, and search interfaces for retrieval by indices.
– *Message threading* that constrains input according to conversation or discourse grammars. Grammar means a set of (context sensitive) rules that specify the types of texts or comments that can be linked to another text or comment, according to its type. By type, we mean remarks that constitute recognizable moves or roles in discourse or communicative interaction, e.g., remark types in a debate include claim, challenge, rebuttal, question, answer, etc. A grammar thus represents a procedural order, often with possible branching.
– *Subscription to arbitrary node:* A retrieved text was wrapped in a form which could be used to request emailing of comments that are later attached to the text or any of its descendants. Since the particular text as a database object is automatically identified in the form, subscription is transparent: The user does not have to specify the node.
– *Moderator tools* enabling submissions to be handled before being posted to public view. These include a) a moderator's view, b) virtual queues for routing submissions on a topic to its moderator, c) forms for attaching ratings

and view restrictions, c) form letters accepting and rejecting submissions. With these, the Open Meeting can support moderated as well as unmoderated online conferences.

2 Meeting the Challenges

This technology better met the NPR goals than chatrooms or list-serves, but it put more emphasis on the meeting as interactive knowledge building than the NPR team originally planned and had less regard for the meeting as community activation. Its deployment was shaped by our recognition of the major challenges that a large on-line meeting presented to the system's potential for organizing texts and building knowledge: Managing the fan-in of comments on the proposals, structuring the interactions of participants and maintaining civility among them.

2.1 Fan-in

In hierarchically structured communication flows, viz., 1-to-many/ many-to-one, a request for comments to a community can produce a stream of responses likely to overwhelm the source of the request, but miss other interested parties. Directing the response stream to all members of the community in a many-to-many pattern just compounds the problems of message volume (gain) and recipients' discrimination of messages that interest them. A strategy to meet this challenge is the decomposition of both the request and response streams, according to a taxonomy of the issue space. This was done for the on-line meeting by considering the NPR reports on the operating systems as subdomains and their respective proposals as categories in a reinventing government domain, and requiring that a comment be subsumed by a proposal. A comment either had to target a specific proposal or another comment descending from it. Instead of being asked "what are your views on reinventing government," a participant was in effect asked: "Should government offices be issued credit cards for making small purchases?" by the presentation of that proposal. Users could find proposals and later comments on them that interested them by drill down through the hypertext o rhrough a taxonomic listing. They could subscribe to receive comments added to such a proposal or, even more finely grained, to a node subsumed by it.

Decomposition of the domain was facilitated by NPR having prepared each of the reports according to the same strategic-plan like format: an Executive Summary described the problems in the respective system, a set of Recommendations proposed how to fix them, one or two Actions were described for effecting each Recommendation, and Appendices added some details of the implementations. These parts were easily reconfigured into a hypertext, using a standard node architecture, that replicated the structure across the main branches of the hypertext (the individual reports), simplifying implementation and providing a consistent user interface. Eleven such nodes, one for each operating system, were linked to a root document which presented the plan and procedures of the meeting. The standard node eventually had links to the parts of the report, including a link to each recommendation, and to other relevant documents, including a short Overview of the system problems, official examples of practices that satisfied a recommendation, and newsletters, which summarized the ongoing discussions of the proposals.

The decomposition is based on the idea that knowledge can be hierarchically localized in the sense that understanding in one branch or area of a domain requires knowledge of more general domain principles but not of fine grained propositions in other branches. The decomposition was intended to reduce the number of people in a conversation and the circulation of their comments, enhance their interest and focus and make the conversational threads easier to follow. The resulting topology was small circles of lateral communication, i.e., several dozen to several dozen, with isolation of conversations from one another being overcome in practice by users' visiting several conversations. Systematic, updated overviews of the discussions regarding an operating system were provided by the newsletters. These were hand-made summaries by the moderators, but automatic summarization or gisting which exploits the hypertext links is an important focus of work to extend the system.

2.2 Structuration

The challenge of devising an interaction structure for the meeting had two sides: First, the need to enable orderly, yet relatively unconstrained conversations, through which knowledge and information about the proposals could be built; second, the need to produce a coherent textbase which would index each comment according to its type of move in a conversation. The messages posted to the meeting were imagined as both pieces of ongoing conversations about particular proposals and parts of a self-extending, distributed hypertext, contributed by multiple sources, including NPR. However conventional messages threaders, such as subject, target ("in reply to"), submission time, were deemed insufficient to structure or represent a set of messages as a sequence of conversational or rhetorical.

Indexing the intentions of the moves is important: First, on the view that conversations are composed of interlocking moves or speech acts[6], on-line communications are problematic, because they have fewer inflection, non-verbal and setting cues which speakers in face-to-face communications exploit to make and recognize speech acts, their intentions, and the expected responses. When the move types of posted comments are indicated, a user can know what types are appropriate responses to a target she has chosen, and by identifying the type of her comment, she can indicate the type of response she expects. Second, when comments are threaded through their targets in hypertext display, the tagging clarifies the relationships between otherwise opaque texts (especially when the text itself is not displayed but only pointed at by a link), and provides an overview of the flow of intentions and expectations through the entire conversation. Finally indices are needed so for a given target, all replies of the same attitude, agreement for example, can be displayed together, retrieved together or more easily checked for redundancy.

The indices can be obtained by asking contributors to identify the types of moves they are making. This self-indexing is similar and not necessarily more inhibiting than the widespread practice of putting "smilies" and other cues into email messages to reduce their ambiguity.[4]. Self-indexing of utterances is also a standard linguistic practice, especially when there is a probability of misinterpretation or a particular type of move is privileged, in the sense of having a priority on getting attention, e.g., "I have a question?"[5]

We also imagined that the sequences of moves and the discussions they created would be constrained by a context sensitive discourse grammar, a set of moves and rules that restricts their possible combinations by specifying the moves that can be legally attached to each move. These rules abstract the notions of appropriateness and expectation that are associated with the move types in a sequence. Since they have institutional as well as logical bases, a grammar encapsulates a quasi normative order or a procedure employed for a social purpose, not just idealized linguistic phenomena[6] The sequences which we thought most likely to occur and best to support on-line were multi-lateral arguments for and against the proposals. Our model of argument was influenced by modern dialectical theories, which understand argument as an ordered, interactive process whose utterances are constrained to speech acts for making and supporting claims, challenges and rebuttal.[7, 9, 8]. The highly restricted move sets of the dialectical argument models were attractive because they barred the flames, digressions, emotional discharges and similar failings of newsgroups. However, we anticipated that discussions constrained by that grammar would be more conflictual and address fewer possibilities than if the grammar also permitted alternatives, examples, information seeking questions and answers. Because the Open Meeting was intended to stimulate interest and accumulate information about government reinvention, the teams selected moves that could lead to consensus while permitting differences of opinion: Agreement (reason for), Disagreement (reason against), Question, Answer, (propose an) Alternative, Qualification ("yes, but"), or (report a) Promising Practice. The root document explained these types, presented a distinctive icon for each, and asked users to frame comments according to them.

The grammar for the meeting was additionally shaped by pragmatic institutional as well as logical constraints. To avoid the scent of electronic plebiscite, polling, voting and simple endorsements were not implemented. Agreements and Disagreements had to provide reasons. Because the executive summaries, the

[4] Multi-media messages will not entirely eliminate the problem, especially when speakers do not know each other or are from different cultures.

[5] Self-indexing is often used in political discourse to emphasize the intention of certain utterances. One famous example is Zola's attack on the French government for its cover-up in the Dreyfus affair: *J'accuse.*

[6] People are often upset because discourse grammars work to mitigate the threats to face – people's power and dignity – that is inherent in social interactions. Violations of the grammar are consequently implicit attacks on face[10].

overviews based on them, the appendices and the example promising practices were approved by NPR leadership, they were in a sense above debate, at least by workers at the meeting and no comments could be attached to them. Finally, the attachment of alternatives to questions, answers or promising practices, and the attachment of anything but answer to question were considered illogical and excluded.

The discourse grammar is weakly enforced in the Open Meeting system by using it to drive the dynamic reconfiguration of input interfaces. Given the user's choice of target, the server presents an input interface which limits the choice of comment types to just those which can be "legally" attached to the target's type. Of course, this does not prevent a user from flouting the rules, by identifying a comment as a permitted move, but actually making one that is not. A common case of flouting in political meetings is the use of a turn in the question period to assert one's own position, sometimes eliciting the rebuke: "That's not a question?" In general, to prevent potentially disruptive flouting requires natural language processing capability that understands the intention of utterances – not just recognizes forms. Absent robust natural language processing technology for this, human moderators are needed to verify that the contributor has not flouted. As noted above, moderators for the meeting did verify that Agreements and Disagreements gave reasons, but verified no other move types, because of the low potential for disruptive flouting.

The resulting virtual discussions in the Open Meeting differed from face-to-face, multi-lateral discussions in not having an inherent linearity, that is, forward movement toward conclusions. In face-to-face conversations, discourse contexts or foci of attention are successively closed to further comments and are only reopened with difficulty[11]. A speaker might present a reason against a proposal, with the reason then becoming the context of subsequent debate among speakers. At some point, however, discussion returns from this stack to the top level and moves to another point, opening another context. There will be no backtracking to the first context, unless information or conclusions are premises that have been contradicted in the current context. In the virtual discussion on a proposal, all the contexts (and their contexts recursively) remained open to further comment, even if recent comments were in the same context. The discussion could therefore develop as a set of statement and reply sequences on different points, a long sequence elaborating one point (analogous to breadth versus depth search) or a mix of both approaches. This openness partly reflects the primary orientation of the meeting toward collecting information and building knowledge rather than moving a community toward a consensus.

Unfortunately, the outline form used on the web page to represent the links in the hypertext obscures the time dimension and possible focus or convergence in the discussion. This regime recursively indents the pointer to a comment with its move type icon under the pointer to its target, with comments at the same level listed in order of submission. As a result some comments at a deep level that were more recent than other comments at a more shallow level can higher on the page. Since all submissions are time stamped, it is trivial to present the hypertext in ways that could encourage or represent convergence, such as hiding or closing contexts that have received no comment for an arbitrary time and color coding contexts according to how recent are the last comments in each.

2.3 Civility and Moderation

The design included several means of maintaining civility and the quality of interactions. Users were prodded to self-reflection and some restraint on impulsive commenting by having to identify comment type and summarize the content in the a subject line. A user was also unable to disguise her identity, because had to submit the comment on a form that was mailed to her email address[7] In addition, moderators were used to minimize the posting of low quality, redundant and inappropriate comments.

The moderators' tools included virtual queues to allocate work, review forms, form letters to users, and a constraint-based view system. These were used to overview all submissions to the meeting, access unreviewed and pending submissions on a topic, rate a submission's quality and decide its status. Options for the last function included accepting the submission for exposure to users, rejecting it or returning it for revision with a form letter explaining the reason, and deferring a decision to another moderator. The moderators in the on-line meeting, volunteers recruited by NPR and assigned to specific nodes, were trained to reject submissions that were low quality, redundant, obscene, personal attacks, whistleblowing or commercial solicitations. By checking on the review form the reason for a rejection, a moderator automatically returned the submission to the user with the appropriate letter of explanation. Much of this labor can be further reduced by the development of rule and neural network based programs that can automatically detect some prohibited content types.

Moderation exploits the database support of views, since accepting a comment merely included it in the public view. As implemented, moderators could see all submitted documents with their review status and quality ratings, while the general user saw or received only those that had been accepted for public display. More generally, views are displays generated by constraints that determine what gets shown to whom, display of the textbase and can be apportioned by arbitrary criteria, e.g., security classifications, organizational roles. A role, like moderator, then can be defined as a capacity to make certain moves towards others, based on the information the role holder gets or sees. Information is thus the context for interactions and power relations in the group.

The discourse grammar and the moderator reviews were constraints on what was said during the meeting, but there were factors that diminished their potential for arbitrary suppression of voices by the moderators. First the grammar itself was fairly transparent and was consistent with the announced purpose of the meeting as a discussion. Second, the moderators for the meeting were peers of the users, shared an interest in a meaningful discussion of the recommendations, and had to give reasons for rejecting a submission. Nevertheless, it is worthwhile

[7] This method is no longer useful, due to the proliferation of free mail services.

comparing the risks that having rules of order and moderator authority poses for democratic practice with other means of control in on-line forums.

- *Unmoderated forums*, bulletin boards, mailing lists, etc., depend on the civility and common sense of their users. They are plagued by digressions and flame wars which drive away users, effectively suppressing diverse opinions. Efforts by some users as self-appointed "thread police" may be counterproductive. An option to improve control is collective moderation in which all participants rate the value, relevance and civility of the posted comments, with users receiving low scores being automatically barred. Certified registration could foil attempts by exiles to post under assumed names and also determine who is a participant.
- *Moderated forums* involve posted guidelines and moderators who usually preview (sometimes review) submissions to block (remove) from posting those violating the guidelines. The costs are a) possible abridgment of free speech rights[14], b) moderators seeing themselves as proprietors of the forum, c) moderators' interpretations of the guidelines are opaque regarding how and what they judge are violations. Users, after all, do not see what is not posted. The use of content analytic as filters on submissions would reduce the need for moderators and enable collective reviews of rules or examples of violations more fine-grained than the initial guidelines. This leaves open the questions of who writes the rules or chooses the egregious examples and how intelligent do the filters need be.
- *Adding Roles to Discourse Grammars* is a means to reduce the potential for aberrant, disruptive interpretations of the situation, that can be enacted within the constraints of the grammar. A cheap implementation of this idea is providing more directions to users in the introduction to the meeting, including the roles they might imagine themselves playing and there motivations in these roles. A more complete fix would be the support, similar to that given moderators, of roles that are either help the group pursue the discussion or manage itself. For example, a resident expert with access to specialized sources may be able to settle questions of facts that arise in the discussions. In contrast, a mediator could handle the problem of two participants using the forum as a private battleground by opening a back channel to them with suggestions to resolve their differences, or just remove their exchanges from the view of other participants. Increasing the organization of the meeting also increases disparities in power and social relations, since these are based on the information and moves available to the respective users.

3 Discussion

3.1 Use Patterns

Announced over specialized mailing lists and bulletin boards for government workers, the meeting attracted the interest and participation of workers from a wide and geographically dispersed range of federal government offices and military installations. By its beginning in early December, 1994, over 4200 people had returned an on-line registration survey. 90% were government workers, with over 60% located outside the Washington, DC area, including in all fifty states and twenty countries besides the United States. They represented 40 executive branch agencies and departments, with the large and more technically oriented organizations, viz., Defense, NASA, Interior, tending to have the most representatives. They were typical of government workers with respect to age (average in early 40s) and experience, but significantly more senior, educated, and technically oriented (60% vs. 25% in supervisory capacities; 47% vs. 15% with MA or more; 66% vs. 11% in information systems or engineering). Women were under-represented (30% vs. 42%), but their rate of registration was double the estimate percent female of Internet users at the time (15%). These distributions indicated that interest and network access were key factors for participation. The registration rate of women, despite their lower levels of employment in technical positions, and the large registration at Housing and Urban Development were due to special outreach efforts by NPR personnel to these groups. In contrast, few Treasury and National Security Agency personnel registered because these offices had firewalls blocking access to outside networks.

Over 3000 users accessed the meeting during its two week, but only about 2000 went beyond the introductory page and welcoming pages to a standard node and only 1000 attended the discussion threads. Several factors accounted for the attrition: End-of-the-year and holiday chores left workers little time for the meeting; they found email traversal of the hypertext too cumbersome and the texts themselves overwhelming, and may user were interested only in information and not discussions.

The distribution of web accesses by file types shows two patterns among users accessing a standard system node. About two-thirds of these users surfed the meeting material in several reinvention topics, looking primarily at the overviews and newsletters when available. Users in the second group tracked online discussions on several proposals apiece, in one or two reinvention topics, through both retrieval from the hypertext interface (16 comments read on average), and subscription to mailings of comments posted to the thread (as indicated by the concentration of subscriptions on those topics and proposals which were most accessed). Those who commented were necessarily in this group and similarly focused in their comments; with few exceptions, they posted to just one or two discussions. There was relatively little use of the tools in the meeting environment for search directly by reinvention topic, comment type, government agency referenced and other indices. On the whole, these patterns suggest an orientation toward the meeting as a source of general, rather than specific information, and an interest on the part of the in-depth users in what their peers had to say on a few specific issues.

290 different individuals contributed 1013 comments that were accepted for posting. The ratio of one contributor for five actual lurkers at the discussions is much higher than the 1 to 10 or 20 ratio variously estimated for newsgroups.

Half the comments were identified as Agreement and 15% were Disagreements. The relatively large number of Questions (167), Alternatives (106) and Promising Practices (72) indicates the contributors' willingness to seek and share new information at the meeting. The near absence of the cognitively complex Qualification (3) supports the contention that cyberspace does not nurture the reflection desired in a public sphere, although a sampling of comments found that some qualifications were misidentified as agreements. Possibly these users did not want their comments to be construed as unsupportive. As a group, the commenters were positive toward the proposals and tried to present constructive ideas, some of which NPR leadership culled for further examination.

Web users were generally satisfied with the system, email users complained about its clumsiness for traversing hypertext, and both groups complained about an overwhelming amount of text to read. These responses, particularly the last, were anticipated. We believed that, consistent with their bureaucratic training, people attending the meeting would want to read the "manual" (the reports) before commenting on the proposals, so we wanted to reduce the texts to much shorter pieces. Because, the NPR team lacked the authority and personnel to do the editing, the overviews were produced as a means for users to get acquainted quickly with reinvention topics.

3.2 Discourse Management

The on-line meeting demonstrated that virtual discussions localized to specific topics in a domain could attract people across organizational boundaries, remain coherent and produce possibly valuable ideas. What factors account for the success? Can they be replicated in other venues and with other groups?

The meeting provided individual self-esteem incentives for access and participation, felicitous levels of discussion and effective discourse management. Users knew they could acquire, share and be seen to have information or ideas. Any positive motivation which they as information workers had toward these activities was strengthened by the salience of the discussion domain and the opportunity to investigate a new information technology. "Reinventing government" was frequently mentioned by the media at the time in the context of downsizing and hence a possible threat to their employment, but fewer than half the registered reported knowing much about it or even having seen an NPR report. The meeting also offered concerned workers possibly their only opportunity to influence the implementation of the recommendations. This incentive was strengthened for some workers by their belief the Vice-President would personally review the discussions[8]. More group oriented motivations, like helping to form a group or speaking with one's peers, may have motivated some users to revisit or track discussions. In informal, post-meeting evaluations, these users reported having a "we" feeling when they discovered people from other organizations who were contending with the same problems.

[8] In addition to contributing a letter of welcome to the meeting, the Vice-President Gore did post two questions.

The focus on individual recommendations for common operating systems placed the discussions at a level close to that of concrete operations, without requiring users to know terminology or conditions of a particular organization. It was consequently easy for users to draw on their own experiences in reading and responding to the recommendations and comments. Because the focus also limited the attraction and participation for any particular discussion, potential contributors could have realistic expectations that their opinions would be noticed and could influence the discussion. This sense of being critical encourage some users to expend the intellectual energy and time to comment, rather than just lurking. Note the meeting was arranged to reduce other costs of participation by enabling attendance from their desktops and at any time convenient for them during work hours – with the blessings of the Vice-President. Indeed the potential for the Internet to decrease the transaction costs of participation is a major reason to imagine it can help revive a public sphere.

The presentation format for the recommendations and the structure provided by the discourse grammar were sufficient cues for participants to recreate the collective policy or program reviews, that many were practiced in their work. This result could be expected, given their background organizational cultures and bureaucratic roles, but was not inevitable. Since roles were under-specified, and the discourse moves were general, the situation was open to interpretations. One participant, a line worker and union steward, submitted scores of questions, agreements and disagreements, which emphasized the need for the government as employer to improve working conditions. He seemed to be following a script or personal role of questioning authority and only stopped, when a moderator suggested by private email communication that he could get more attention by posting fewer comments. The incident highlights that discourse grammars do not constitute a universal pragmatics, that is sets of moves (speech acts) and rules which are understood similarly by all and which construct processes that are transparent to everyone.

As implied above, the second source of success was the people who attended the meeting. Most had training that made them ideal participants: They were information workers, socialized to a culture of meetings, with a professional or managerial interest in what was being said and pre-existing familiarity with how to say it. Such attributes can also be found among the information workers, who have some experience of collaboration over networks. But while their rapidly increasing numbers make them politically significant in the United States[12] these workers hardly constitute the universe, and newsgroups demonstrate that discussions of political and social issues among networked information workers are no more rational and orderly than those of the general population[13]. Online public spheres open to politically and culturally heterogeneous participants are consequently likely to have problems of governability, beyond the capacities of current types of controls.

3.3 The Open Meeting as a Public Sphere

How well did the Open Meeting realize the ideal of a public sphere? How good a model is it for an on-line democratic institution? One answer is that the participatory process through which a community becomes an active agent on an issue requires members' access to relevant information, their deliberation, decision, privacy, commitment to the decision and time. Privacy here refers to protection of participants from intimidation, insult, defamation or similar means intended to chill participation in public discussions. Time refers to the examination of a significant issue and the expression of all interested parties' views on it taking a considerable amount of time – weeks or months. By these standards, the Open Meeting does half the job; it successfully provides a platform for accessing information and participating in orderly discussions.

As implemented for the meeting, the major shortcomings of the system appear in regard to decision and developing a community spirit that would facilitate commitment by participants to a decision. As noted, there were reasons for not implementing a decision mechanism, and such a mechanism, arguably, might not be needed, since argument grammars have an implicit decision mechanism. They terminate when challenges to a claim are exhausted, a challenge cannot be countered or the opposing sides agree to disagree. These endings are more in keeping with a consensus seeking process based on force of reason alone. Nevertheless, few communities really have the time for exhaustive arguments and few scenarios allow decisions to wait for them. Realistic, effective support of on-line processes in which communities clarify issues and take stands has to include event-creating moves like cloture and voting procedures. These moves are event creating; they change the phase of a deliberative process by ruling out some previously permitted moves and activating other ones. More generally, an effective Open Meeting system will need a library of tested procedures, represented by discourse grammars, with projected outcomes and identified risks, against which conveners of community meetings or collaborations can measure their needs.

As noted earlier, the strategy of decomposing large meetings into small discussions has the potential to isolate users and block their sight of the large community. A low cost means of meeting this problem is an interface display of the distribution of interests over the various topics and proposals, as measured by user visits or comments. Users would be able to gauge where their own discussions fit in the distribution of the community's attention and to directly access popular or other spots, if they so wanted. Slightly more expensive is the automatic generation of cross-links, based on statistical information retrieval methods and patterns of semantic links, that point users to other discussions with themes or information similar to their own. As well as discovering cross cutting themes or issues to integrate the knowledge produced in the discussions, this technique can also be used to discover construct wider circles of participants with similar problems.[9]

[9] The second idea follows from [15], an implemented application of agent technology, using cluster analysis of email to find individuals with similar interests. Inversely,

By showing under certain condition, suitable mechanisms can produce a successful, widely-attended discussion of issues, the Open Meeting was a useful step toward building democratic practices in networked communities. However, because of the importance of shared organizational/ cultural backgrounds for the success of these discussions, it is more realistic to think about multiple public spheres, each serving a different subcommunity. Each group could develop its own discourse grammar, and agents in each sphere would seek cross-links to other groups by similarities in issues discussed. While that procedure does not create consensus, it might promote mutual awareness.

References

1. Habermas, J.: The Structural Transformation of the Public Sphere. M.I.T., Cambridge, MA (1989), first published as Strukturwandel der Offentlichkeit. Luchterhand, Berlin (1962).
2. Knapp, J.: Essayistic Messages: Internet Newsgroups as an Electronic Public Sphere. In: Porter, D. (ed.): Internet Culture. Routledge. New York London (1997).
3. Hurwitz, R., Mallery, J.: The Open Meeting: A Web-Based System for Conferencing and Collaboration. Proceedings of The Fourth International Conference on The World Wide Web. Boston (1996) 19-36 http://www.w3j.com/1/ hurwitz.349/paper/349.html
4. http://www.npr.gov
5. http://www.whitehouse.gov/WH/publications/html/Publications.html
6. Searle, P.G., Searle, J.: Speech Acts. Cambridge U. Press, Cambridge (1970).
7. Toulmin, S.: Uses of Argument. Cambridge U. Press, Cambridge (1958).
8. Rescher, N.: Dialectics: A Controversy-oriented Approach to the Theory of Knowledge. State U. of New York Press, Albany (1978)
9. Eemeren, F. v., Grootendorst, R.: Speech Acts in Argumentative Discussions. Foris Publishers, Dordrecht, Holland (1984).
10. Brown, P., Levinson, S.: Politeness : Some Universals in Language Usage. Casmbridge University Press, Cambridge (1987).
11. Grosz, B., Sidner, C.: Attention, Intentions and the Structure of Discourse. Computational Linguistics 12 (1986) 175-204.
12. Winograd, M., Buffa, D.: Taking Control: Politics in the Information Age. Holt, New York (1996).
13. Schneider, S.: Expanding the Public Sphere through Computer-mediated Communication: Political Discussion about Abortion in a Usenet Newsgroup. Ph.D. dissertation, Political Science Dept., M.I.T. (1997)
14. Docter, S., Dutton, W.: The First Amendment Online: Santa Monica's Public Electronic Network. In Tsagarousianou, R., Tambini, D., Bryan, C. (eds.): Cyberdeomocracy: Technology, Cities and Civic Networks (1998).
15. Foner, L.: Yenta: A Multi-Agent, Referral Based Matchmaking System. The First International Conference on Autonomous Agents Marina del Rey, California (1997).

cluster analysis of the distribution of participants over the discussions may suggest cross links that are not apparent at the text level.