



THE OPEN MEETING

A WEB-BASED SYSTEM FOR CONFERENCING AND COLLABORATION



Roger Hurwitz, John C. Mallery

Abstract

An asynchronous collaboration system was developed for Vice President AL Gore's Open Meeting on the National Performance Review. The system supported a large online meeting with over 4000 participants and successfully achieved all its design goals. A theory for managing wide-area collaboration guided the implementation as it extended an earlier system developed to publish electronic documents. It provided users access over SMTP and HTTP to hypertext synthesized from an object database and structured with knowledge representation techniques, including a light-weight semantics based on argument connectives. The users participated in policy planning as they discussed, evaluated, and critiqued recommendations by linking their comments to points in the evolving policy hypertext. These policy conversations were structured according to a link grammar that constrained the types of comments which could be attached in specific discourse contexts. Persistent actions enforced constraints on man-machine tasks, such as moderation workflow. Timely delivery of newly moderated comments kept the conversation gain at a level comparable to tightly-focused mailing lists threading out from specific points in the hypertext. After reviewing the architecture and performance of the system in this Open Meeting, the paper closes with discussion of lessons learned and suggestions for future research. **Key-words:** Collaboration, form processing, HTTP, information access, link grammar, National Performance Review, organization theory, persistent actions, semantic network, SMTP, surveys, typed links, URN, White House, World Wide Web

Introduction

A recent Open Meeting of several thousand Federal workers under the auspices of Vice President Al Gore demonstrated that the World Wide Web can support productive, wide-area collaboration for policy planning and problem solving. The collaboration system implemented for this meeting enabled the participants to find and discuss proposals for bureaucratic reforms, which had been prepared by the National Performance Review (NPR). The resulting policy conversations, which crossed traditional agency boundaries, mobilized support for the proposals, helped refine them and gave NPR feedback on their recommendations. The collaboration architecture was effective because it built upon a theory that identified the interactions needed for productive discussion and problem solving, and

also provided ways to reduce the obstacles to such interactions that arise with a large, dispersed group of participants.

After first describing the Open Meeting event, we examine the general theory it embodies as a guide for other wide-area collaboration applications. Next, we review the architecture of the system, and then, evaluate its performance during the Open Meeting. The conclusion offers some suggestions for refining the system and managing wide-area collaboration over the Web.

The Open Meeting Concept

The Open Meeting application implemented the idea that messages posted to an online discussion can be linked by a light-weight semantics into structured discourse and that the discourse can be modeled as extensible hypertext. The idea for

the event itself originated with National Performance Review staffers, who sought an online meeting to disseminate and discuss NPR proposals for reinventing government operations. In their view, a successful meeting would involve several thousand workers, from a wide range of government organizations, who easily access texts relevant to their interests and link their comments in coherent, virtual conversations. The meeting would itself demonstrate a key element of these proposals—the use of computer networks to coordinate policy planning and actions across traditional organizational boundaries. NPR recognized that the conventional technologies for online asynchronous discussion—*viz.*, listserve, newsgroups, and electronic bulletin boards—were not well-suited for such a meeting.

Our research group at the M.I.T. Artificial Intelligence Laboratory considered the meeting an opportunity to implement and test our ideas for managing public access and participation via the Internet in government inquiry and regulatory processes. In such processes, a government agency typically has its experts prepare proposals, and then, invites comments from its relevant public. This public, in turn, hires consultants to prepare briefs and speak at hearings. Because Internet use will broaden and cheapen access to these processes, it will dramatically increase the number of responses to the proposals, and consequently, subject agency officials to information overload. This undesirable result can be attenuated by intelligent routing that decomposes the comment stream by policy proposal and directs comments to the responsible officials.

In an initial public comment system, people would attach their views to documents under review according to the type of their comment. Officials would retrieve these comments by target and type from a database of review documents. Alternatively, the submitted comments could be matched against profiles that indicated the relevance of these categories for the individual officials. For officials reviewing public comments on a proposal the system functions as an annotation

server, which enables them to retrieve specified types of comments on individual proposals. When users both receive and reply to one another's comments, it supports discussions that are composed of the typed and threaded comments.

The Open Meeting application was designed to deliver the necessary support for the meeting organized by the NPR staff. It would build on the Communications Linker System (COMLINK), which was developed as a publication system during 1992-1993 for handling subscription and distribution of documents, based on combinations of categories from a domain taxonomy [7]. The Open Meeting would extend COMLINK mainly by adding typed links between documents in the database.

Given the anticipated character of the meeting's textual environment, the web was a self-evident choice for data entry and display in the system, but the distribution of computer resources among the prospective users made email access as equally self-evident. In December 1994, the time of the event, fewer than half the registrants had a web client and fewer still had clients which supported interactive forms through which comments could be sent to the server. Since all registrants had email, we provided both Web and email access, and as a result, we were later able to compare their respective effects on users' experience of and satisfaction with the meeting.

Textbase

Together with NPR, we made several non-technological choices that affected the organization and interactions in the meeting. These choices concerned the proposals and background material to be included in the initial textbase, how their texts would be presented and the types of comments participants could make on these texts. To provide common grounds for discussions across organizational boundaries, we selected reports which NPR had recently completed about reinventing Federal operating systems, like procurement or information manage-




Figure 1: The standard node

ment, that are found in all federal government departments and agencies. Because the reports had the same generic parts, namely an *Executive Summary*, a set of *Recommendations* and attached enabling *Actions*, and *Appendices* on the implementations of the recommendations, the set was easily reconfigured into a hypertext. A standard node architecture maintained structural analogies across the main branches of the hypertext to simplify implementation and provide a consistent user interface. A root document, which presented the plan of the meeting, branched to eleven nodes, one for each operating system. The standard node included hyperlinks to the various parts of the reports and to additional relevant documents: an *Overview* of several paragraphs and reports of *Promising Practices*, that fulfilled recommendations for the system. During the meeting, *Newsletters*, which summarized the ongoing discussions, would be attached to their respective nodes (Figure 1).

Interestingly, the textual components of the standard node correspond to the generic parts of a strategic model or plan for reforming the operating system: The *Executive Summary* states the problem, the *Recommendations* propose solutions and means of obtaining them, the *Actions* describe tactics and the *Promising Practices* are

example solutions. On this view, the conversations about these texts during the meeting are part of a problem solving process that generates refinements and evaluations as well as support for the proposed solutions (Figure 2).

A recommendation in the Open Meeting environment is consequently an evolving document—its own hypertext—that can be represented by a page with hyperlinks to pages for its original text, the associated enabling actions and the comments in the discussion about it. (Figure 3). The header for each page includes the title, time of submission, author and a location-independent document identifier. To facilitate navigation, each page showed its context with anchors to the immediate parent and to pages summarizing related material. For email users, a text arrived embedded in a form with which one could order one or more the texts subsumed by the present text. A topic node form, for example, included the *Overview* text and an order form for the various parts of the report including the individual recommendations, listed by their titles.

Link Grammars

Comments in discussions are instances of conversational moves which appropriately reply to preceding comments. In ordinary conversation,




Figure 2: Strategic model




Figure 3: One of ten NPR topic areas

speakers implicitly recognize these moves, their intentions, and their expectations of reply. In more stylized discussions, speakers often announce the type of statements they make, e.g., "I have a question," to clarify their relation to a previous statement and to cue the expected type of reply. When comments are threaded through their targets, the identification of a conversational move indicates the relationships between otherwise opaque texts, and the sequence of typed conversational connectives indicates a flow of intentions and expectations.

What link-type grammar is appropriate for an online meeting? By grammar, we mean a set of rules that specify the admissible ways in which comments can be linked to an evolving hypertext based on their type and the context. These rules formalize the quasi-normative order of a conversation and prevent incoherent or inappropriate sequences. Such rules can be enforced at a dynamically reconfigurable interface which limits the choice of link type to those links that can be legally attached to the target comment.

The selection of link types and a composition grammar govern the character and development of knowledge in an online discussion. Conversations that permit only agreement or disagreement [1, 9] are more conflictual or stunted than those

also permitting alternatives, examples given and questions and answers. Since the Open Meeting was convened to discuss policy and rule making, we wanted a set of link types that were familiar in policy debates, and that could express differences of opinion without polarizing participants. After careful consideration, we excluded simple endorsements of a proposal and motions that would call a vote, and narrowed the choices to *Agreement*, *Disagreement*, *Question*, *Answer*, (propose an) *Alternative*, *Qualification* ("yes, but"), or (report a) *Promising Practice* (Table 1). The *Root* document explained these types and asked Open Meeting participants to use this link semantics to frame their comments.

Certain institutional and logical conditions dictated the attachment rules in this grammar. Some NPR assertions has been vetted and were officially beyond debate; consequently, no comments could be attached to the *Overviews*, *Executive Summaries*, *Appendices* and *Promising Practices*. Second, it did not seem reasonable to comment on the *Newsletter* summaries of discussions. Third, other kinds of attachments, namely an alternative or qualification to a question, and an alternative to an alternative, answer or promising practice, were excluded for illogic.

Table 1: Open Meeting Comment Link Types

Icon	Link Type	Description
	Agree	A reason to support the recommendation or action
	Qualify	A qualification that explains exceptions or extensions for a recommendation or action
	Alternative	An alternative way to implement a recommendation or action
	Disagree	A reason to challenge why or how a recommendation or action can work
	Example	A report of a promising practice that illustrates one good way to realize a recommendation
	Question	A question about a recommendation or action
	Answer	An answer to someone else's question

Conversation Displays

An Open Meeting participant submits a comment on a commentable text (*Recommendations, Actions, other comments*) by editing the form attached to that text. The form captures the target's document identifier, lists the comment types that can be attached to the target, and provides queries for the comment title and text. The database creates a document object for the comment and uses the link information in generating a virtual page that displays the current state of the discussion.

The page includes a hyperlink to the recommendation and hyperlinks to the the comments, each listing the comment title, author, time of submission and link type, with the last indicated by a distinctive icon, as well as type name. These hyperlinks are displayed as a recursively indented outline, so hyperlinks that directly attach to the same target are below it, with the same offset. Hyperlinks to all comments in sequences and subsequences attached to one target are listed before the hyperlink to the next target. The layout (Figure 4) provides a synoptic view of the discussion.

Intergovernmental Tax Filing, Reporting, and Payments

Recommendation for Reengineering Through Information Technology

[Intergovernmental Tax Filing, Reporting, and Payments](#)

- Agreements**
 - [Good](#)
 - Michael O'Rourke [11 Dec 1994 19:38 EST]
 - [Provide Incentives to file electronically](#)
 - JOHN R COBARRUVIAS '78 [15 Dec 1994 18:06 EST]
- Alternatives**
 - [Use Of Commercial Software](#)
 - rfs@NRC.GOV [19 Dec 1994 11:11 EST]
- Promising Practices**
 - [Encouraging electronic filing](#)
 - Scott Morizot [15 Dec 1994 23:53 EST]
- Alternatives**
 - [Stop sending paper tax forms](#)
 - JOHN R COBARRUVIAS '78 [15 Dec 1994 18:08 EST]
- Disagreements**
 - [Stop mailing paper forms](#)
 - Scott Morizot [15 Dec 1994 23:57 EST]
- Actions**
 - [Integrate Financial Filings, Reporting, Payments](#)
 - [Eliminate Filing Routine Income Tax Returns](#)
 - Agreements**
 - [Converting to a flat tax rate could make paperless](#)
 - gail_warner@HUD.GOV [9 Dec 1994 14:22 EST]
 - Agreements**
 - [Greatly Reduce IRS Staff and Facility Requirements](#)
 - Gary Romsaas [10 Dec 1994 16:08 EST]
 - Disagreements**
 - [Not everybody has a PC](#)
 - FRED GROSBY [20 Dec 1994 20:46 EST]

Figure 4: A NPR recommendation, implementing actions, and linked comments

Moderation

To minimize the posting of low quality, redundant and inappropriate comments, the Open Meeting was moderated. Moderators were assisted by administrative tools, which include moderation forms, canned response letters, virtual queues to allocate work, and a constraint-based view system. A moderator can use these tools to overview all submissions to the meeting, access unreviewed and otherwise pending submissions for a topic, rate a submission, accept it to make it visible, reject it, return it for revision, or defer a decision to another moderator (Figure

5). Moderation exploits the database support of views, since accepting a comment merely changes status of its visibility to the public (Figure 8). Views then are displays generated by constraints that determine what gets shown to whom. Although this interface generation idea can be used to apportion the textbase according to arbitrary criteria the Open Meeting employed only user and moderator views. Working with their view, moderators could see all submitted documents with their review status and could retrieve comments based on the quality ratings (Figure 6).

Moderator Review Form

Date: 16 Dec 1994 23:05 EST
From: Benjamin Renaud (benjamin@WILSON.AI.MIT.EDU)
Document ID:

Source: AGREE (Ecosystem Management via Cross-Agency Management)
Current Status: DEFERRED
Subject:

Text Body:

It would be very useful to build common ecological models to facilitate cross-agency management. System dynamics models would be most appropriate.

Review:
☐ Accept ☐ Defer ☐ Reject

Importance:
☐ None ☐ Low ☐ Average ☐ High ☐ Exceptional

[Search Any Open Meeting Documents](#)

Send Email To: (Separate multiple addresses by commas.)

Email Default Text:

Send Email:
☐ No ☐ Yes

Moderator ID:

Figure 5: Moderator Review Form




Figure 6: Moderator Search Results for exceptional comments concerning the Department of Defense

Ancillary Pages

The Open Meeting environment included friendly interfaces for retrieval of particular text types and online help. A search interface supported retrieval of documents satisfying near boolean combinations of reinvention topic (node), link type and government organizations mentioned in the document text. *Promising Practices* and *News* interfaces enabled retrieval of hyperlinks to all the promising practices or newsletters by their reinvention topic. These were implemented by standing search URLs, which pointed to the search specifications for the required documents rather than the documents themselves and hence avoided the problem of updating hotlinks. A general help page listed hyperlinks to Vice President Gore's welcoming letter to the Open Meeting, to his memo authorizing federal workers to participate during work hours, and to various FAQs.

Wide-Area Collaboration

Wide-area collaboration refers to communication and coordinated action among groups that are large, geographically dispersed, and generally, do not know each other. These kinds of systems are distinguished from *groupware* oriented toward small groups precisely because the system must take over many tasks previously performed by people in small groups.

- *Large-scale communication.* When large numbers of people are involved it is no longer possible for individuals to see all of the communications traffic.
- *Decomposition.* Communications must be broken up into smaller packets that are narrowly focused. The decomposition can occur along several dimensions:
 - Time: *Asynchronous* communications becomes the norm. Relaxation of synchronous constraints on participation is essential because large groups are difficult or impossible to schedule, especially across multiple time zones.

- Space: *Geographic* decomposition provides a way to focus collaboration whenever the domain has spatial extent.
- Content: Specialization by interest, role or function provide the most general way to hierarchically decompose a task domain.

In general, task decomposition allows group size and task elements to be scaled down to a manageable size. The key idea is to reduce the volume of communications and increase the locality of communications in order to match information processing levels with people's ability to cope with complexity and with their commitment to the collaboration.

- *Structuring information fragments.* The decomposition of communications creates a later need to reintegrate information for coherent reconfiguration and presentation to people. The reintegration and delivery options depend critically on the structuring techniques used to organize information fragments.
 - Minimize redundancy by recognizing equalities: In wide-area collaborations with large communication flows, it is essential to reduce information that could otherwise obscure new information. Because people are still needed to recognize similarities and equalities in the various pieces of information, the organizing strategy and the user interface must help them discover whether information they intend to link has already been linked.
 - Danger of self-amplifying redundancy: As the quantity of redundant information increases, it is increasingly difficult to recognize prior similarities, and a wide-area collaboration system risks descent into an unmanageable morass at an accelerating rate.
 - Atomic propositions: It is easier to spot redundancies when comments are short and addressed to one point. Several statements of this type are better than a single

large statement, interweaving multiple themes.

- Knowledge-level annotations: A set of statements that are largely opaque to a computer system can be organized into traversable hypertext or a semantic network by making assertions or annotations about them. The more explicit the semantics of these assertions, the more useful computer manipulations becomes possible.
- *Focus activities and interactions.* Communications decomposition should cluster information and actions into meaningful and coherent chunks that match cognitive capacity and motivational level of participants.
- *Locate interest, expertise, resources, responsibility.* Wide-area collaboration involves the coordination of actions and human resources in addition to information assembly. Coordinating action provides a criterion for decomposing information about agents according to several dimensions:
 - Interest in participating
 - Expertise or special knowledge
 - Ability to provide or deploy resources
 - Responsibility for making decisions

Knowledge-Level Techniques

If effective wide-area collaboration depends on a fine-grained decomposition of information structures and communications processes, it also requires a repertoire of *knowledge-level techniques* for structuring information fragments. Knowledge-level techniques refer to a continuum of approaches for organizing information packets based on their semantic or knowledge content.

Systems of categories organized from general to specific, or taxonomies, provide one of the most powerful ways to organize hypermedia nodes. Taxonomies allow inferences about similarity. Typed links are another extremely powerful way to make statements about how hypermedia nodes relate. These important concepts from the

field of Artificial Intelligence comprise the basic building blocks for knowledge-level techniques. In an application, these ideas need to be combined with a *domain theory*

Various knowledge level techniques were applied in the Open Meeting.

- *Information access*
 - Boolean combinations of features: Once features are associated with information fragments, they can be retrieved in sets by combining boolean operators (*e.g.*, AND, OR, and NOT).
 - Taxonomic subsumption: By organizing categories into hierarchies, it becomes possible to make inferences about similarity based on the set of categories spanning a hypermedia node. Additionally, a node inherits certain capabilities based on the set of categories that span it.
 - Typed links: When the links between hypermedia nodes are typed, they can be used to retrieve other nodes with specific relationships to a given node. Additionally, when the links are first-class objects, information about the link instance can be associated with it.
 - Attachments: Nodes can be filtered according to special special-purpose attachments like *Generic Reviews* that provide a characterization along a dimension or *Discourse Contexts* that provide location in organizational processes.
 - Role-based views: Nodes and links may be differentially accessible depending on application-specific roles, for example, moderators vs. users in the Open Meeting.
- *Structure the information base*
 - Fully categorizing the evolving hypertext: Categorization is a key mechanism for hypertext reassembly that allows regions to be found by boolean combinations of categories, and sometimes can uniquely locate hypertext nodes.

- *Category coherence.* If commentary and other hypertext nodes are thematically atomic and adequately covered by their categories, they can be manipulated reasonably by means of those categories. If the content spans additional categories, the value of categorization declines.

- Linking commentary recursively: Linked conversations focus the evolution of debate to the extent that comments remain on topic, *i.e.*, with the range of their categories.

- *Link grammar.* A link semantics adds an important source of coherence when it expresses which conversation moves are possible by particular people in situations. Here, a grammar explicitly the moves (links) and their composition rules. The representation of the discourse context (*e.g.*, Time, Speaker, Affiliation) reflects and guides organizational processes.

Architecture: The COMLINK Substrate

The Communications Linker System (COMLINK) provides a foundation for research into intelligent network services through a general-purpose substrate that is configured by a small amount of application-specific code. The core of COMLINK is a transaction-controlled, persistent-object database. Users interact with the database via email servers and Web servers. These servers present messages or Web pages whose content is generated on the fly from the database. Dynamic Form Processing module [6, 7, 8, 10] manages all interactions with users over *both* email or World Wide Web using a single, unified paradigm that, *inter alia*, validates all user input. Figure 7 summarizes the COMLINK architecture.

The database defines persistent objects related to the domain of network services. These persistent objects are defined with the Common Lisp Object System [4, 11]. They support multiple inheritance, a mix of persistent and dynamic instance vari-




Figure 7: Communications Linker System

ables, as well as multimethods, which allow method invocations to dispatch on possibly multiple arguments.

Basic Database Entities

The database represents the entire range of entities relevant to structuring a hypertext, operating on it, and providing interactive access to it over SMTP and HTTP.

- *Documents.* Document objects can be created from a variety of different mixins depending on the kind of document. The same document can exist in multiple formats, for example, ASCII versus HTML. Although document properties (e.g., categories, dates, authors) are indexed in the database, the body text is stored in a file system but accessed via a transaction on the document. Small documents like comments use all the same machinery as large documents.
- *Persistent Document Identifiers (PDI).* These PDIs are a kind of prototypical URNs which have the form `pdi://logical.authority.dns.name/yearmonth/day/unique-id.document-format`. Because every document stored in the database has a PDI, external references to documents over email or WWW are easy, uniform, and indepen-

dent of physical location. PDIs provide the critical reference resolution capability necessary to link documents with comments and to attach generic reviews.

- *Persistent categories.* All documents have associated categories that characterize their content. Various taxonomic inferences such as subsumption and exclusivity are available. The database actually contains flat features with a one-to-one mapping to categories that are taxonomically structured in dynamic memory. This allows the taxonomy to be reorganized without the need to perform hazardous surgery on a running database.
- *Taxonomic email routing.* Two types of message routing need database support:
 - Static mailing lists: Mailing lists, subscriptions, and subscribers are represented as database objects. Mailing lists are organized in a generalization hierarchy such that messages to a superior are sent to all inferiors. Mailing lists can be active or inactive. User subscriptions connect subscribers to mailing lists and can be active or inactive. Periodically, all active mailing lists are written out to a mailbox table that drives an associated SMTP mailer.

- Virtual mailing lists: Document universes associate collections of documents, categories, and document selectors. A document selector is a pattern of categories that selects documents for transmission to a recipient. When documents are transmitted through a document universe, the categories attached to the document are matched against all active document selectors, and when matches succeed, the document is sent to the subscriber associated with the selector. Currently, document selectors first match against a document intersection of attractor categories, and second, filter documents by a union of repulsor categories. Document distribution occurs within a transaction in order to assure reliable and atomic delivery to all recipients.
- *Ontology of network entities.* Beyond these major database entities, there are comprehensive variety of objects defined for users, contexts, hosts and domains.

Link Representation

The basic ontology provides the database support needed to access or route documents according to taxonomic categories, but it made no provision for representing links between document or making assertions about them. For the Open Meeting, relations were added to the COM-LINK substrate. Borrowing from our research in natural language understanding [5], the approach added bidirectional ternary relations as first-class database objects. This small addition turned the document database into a semantic network with typed nodes.

Ternary relations have three components: a subject, an object, and a relation type. In this case, relations are used to link document objects. The PDIs used a document identifiers make it easy to link documents or comments together, regardless of their physical location. In the Open Meeting application, the relation types were the argument connectives and several internal links. Addition-

ally, relations are explicitly represented as first-class objects so that assertions can be made about the relations as well.

In our natural-language research, we use ternary relation knowledge representations to represent English sentences because they are arbitrarily expressive, they can encode higher order logics, and yet, they support efficient computations. Thus, this approach to light-weight semantics for linking documents together evolves smoothly to heavy-weight semantics as ever more intensive knowledge-level techniques are combined with hypertext.

Generic Reviews

There are many applications that need to attach rankings, reviews, or discrete values to database objects. A generic review system was implemented that uses a single set of entity definitions to implement any range of reviews schemes, provided review values can be encoded in a numeric scale. While database objects in persistent memory are attached appropriately and hold a number representing the application meaning, these numbers are translated for use in dynamic memory as necessary and relevant for the application.

- *Appraisal.* These are the generic reviews about an entity that are provided by users or programs. These can be active or inactive.
- *Reviewable Object.* A mixin allows any database entity to be reviewed by attaching an appraisal value.
- *Reviews.* Reviews name a specific scheme for generic reviews and associate a function for asserting, interpreting, and comparing appraisal values. Whenever there are multiple appraisals for objects, reviews maintain appraisal aggregates.

Moderation Workflow

In the Open Meeting, generic reviews implemented the following capabilities:




Figure 8: Moderation Work Flow

- *Quality ratings.* Moderators rated the quality of comments as *low*, *average*, *high*, or *exceptional*.
- *Moderation status.* Comments submitted by users could have any status of the following at a specific time: *unseen*, *pending*, *accepted*, *rejected*, *deferred*, or *removed*.
- *Virtual moderation queues.* Moderation workflow is managed by virtual queues (Figure 8) that allocate moderation tasks. When a moderator pops a review task, the task is locked so other moderators to receive tasks without two receiving the same task. Virtual queues are defined by retrieval criteria:
 - Availability: Document whose moderation status is *unseen* or *deferred* but not *pending* are available for moderation.
 - Ordering: Documents available for moderation are ordered according to the time when they were submitted, thus implementing a FIFO queue.
 - Domain: A boolean combination of categories circumscribe the documents avail-

able for moderation by moderators to a specific region of the hypertext.

This approach allows applications to reconfigure moderator queues in dynamic memory by merely changing the combination of categories that define a virtual task queue for moderation. The flexibility inherent in the approach makes implementation of distributed moderation easy and dynamic load balancing of work over a moderator pool possible.

Persistent Actions

Email servers in COMLINK implement reliable tasking by maintaining a queue of pending requests in a task directory. Although this approach works for tasks invoked by users via email form processing, it does not provide a very general or flexible model that could help with access via the Web. The stateless nature of HTTP means that all information regarding a web transaction exists only within the transaction and disappears afterwards. Persistent actions stored in a transaction-controlled database provide a general, fine-grained, and flexible way to ensure the

reliable execution of tasks in networked environments—which are notoriously prone to availability problems and a range of other exceptional conditions.

Persistent actions represent tasks (computations) as database objects. They transfer the reliability of transaction-controlled database operations to the task domain. Reliable tasking works by posting a persistent action to be executed at a specific time, which may be immediate or in the future. Some actions are cyclic and are repeated at specific intervals. When the execution time is reached, the task runs the operation with all associated parameters in its own thread. If the operation succeeds, the persistent action is removed from the database. If the operation fails, the persistent action is rescheduled for execution after an application-defined delay. Transaction control assures that the task is reliably posted in the first place, and deleted only after successful completion.

In the Open Meeting, persistent actions were used for:

- *Moderation time-out.* A problem with the moderation lock system (discussed above) is that a moderator may lock a document for review, but may fail to complete the review. In this case, nobody else could review the document because it would remain in a *pending* state. This problem is solved in allocation transaction by posting a persistent action to revert the status of a document to *deferred* unless the moderator submits a review within a application-specific interval (1 hour in the Open Meeting). (Figure 8)
- *Document transmission.* When documents are distributed automatically, there are opportunities for failure between the time a system accepts a document from a reliable email server to the time it hands the message off to a reliable SMTP mailer. For example, the system crashes. Since accepting a document involves storing it in the database with a transaction, we reliably accept documents

and assist the reliability regime of the email server. When the document is marked for transmission, a persistent action is posted to transmit the document. The persistent action is deleted from the database only after the document is successfully transmitted.

- *Link transmission.* Transmission of document alone is not enough to reassemble a mirror of the hypertext database. In the Open Meeting, the same mechanism as document transmission was used to transmit a *link View*. This link stream contains the link types and attachment PDIs allows mirroring sites to maintain an exact copy of the text-base.

Persistent actions provide a means to enforce constraints on processes in the face of error and uncertainty. The moderation workflow example illustrates how a human process can be coupled with computer support to reliably achieve a task with a number of unreliable parts.

Context Information

Representing the context of communications is a key element in understanding organizational interactions that may occur in wide-area collaborations. Since one purpose of the Open Meeting was to create a framework for conversations across traditional organizational boundaries, the system needed to track the interactions of participants as representatives of their organizations.

- *Discourse context.* The discourse context, which is known as deixis by linguists and provenance by librarians, is available as an object class that can be mixed into major document classes. The representation builds from a conceptualization of agents, actions, and roles:
 - Communicative act: This is the act of communication by a specific communicator over a specific time interval and originating from a specific location. Possible communicators include: people, organizations, and computational agents.

- Communicative role: Any communicative act can occupy the following roles with regard to a specific document: (1) **source**—the agent who is producing the text; (2) **recipient**: the agent to whom the text is directed; and (3) **audience**: the agent(s) who may also receive the text but who are not the intended direct recipients.
- *Network topology*. Email addresses are associated with representations of human and computer agents. The topology of host addresses is represented for Internet Hosts and X.400 Addresses. Although this representation of hosts and domains was originally intended to support maintenance activities (e.g., failed mail processing), it is helpful for understanding of organizational context to the extent that this is correlated with network topology, which is quite high X.400 addresses.

The discourse context provides a means to ground link grammars organizationally; situations and roles constrain the possible links. (Of course, discourse context also supplies information for natural language systems to resolve intersentential pronouns and indexicals).

Architecture: The Open Meeting Application

Hypertext Synthesis

The primary datastructure of the Open Meeting is the database representation of the hypertext. There were two logical views of the structure:

- *User view*. Users could see only nodes, documents, and links that moderators had accepted. This applied to both browsing the hypertext and searching via categories and link types.
- *Moderator view*. Moderators could see all nodes, documents, and links as well as the

moderation status and any internal quality ratings.

In principle, all views of this structure are synthesized on the fly, whether a user is viewing the structure via email or via the Web. Although the overall views presented over email and the Web are the same, differences in the character of these transport media imposed some asymmetries in the user interface, even though both views accessed the same functionality on the same structure. One invariant across all views and user interfaces was the need to provide context-sensitive navigation. Every presentation to users had a variety of links for stepping around the structure and returning to known reference points.

Email Hypertext

Many Federal workers who participated in the Open Meeting had only email access, and consequently, email hypertext browsing was the key technology that made possible their participation. Email hypertext pages always use ASCII forms that rely on the dynamic form processing facility. Hyperlinks are replaced by analogous queries preceding or following any text body. Because email transport is not realtime, there is no need for special caching to improve performance. Users step through pages at the rate of email roundtrips between themselves and the Open Meeting server. For this reason, it was very important to minimize the number of transactions required to traverse structures or accomplish some task, which is usually the number of form submissions by email. The constraint on minimizing email roundtrips introduced some divergence in the interface models between the Web and email views. For example, a single page might offer more options than the corresponding Web page. Context-sensitive navigation was especially important for email users. Despite these efforts, email access remained substantially more clumsy than Web access due to clients which are limited to linear, text-based interfaces and delay times which are often present in transport and processing.

Despite these drawbacks, the email interface served some very important functions in the Open Meeting:

- *Authentication.* Because wide-spread authentication of users was unavailable in the Web browsers at the time, we used a technique of email authentication pioneered in a precursor Community Forum System that deployed at the MIT Artificial Intelligence Laboratory during the Presidential campaign in October 1992. Namely, if a user can receive and respond to an email form sent to their email address, then there is a high probability that the user actually controls that address and their identity is authentic. This assumption is even stronger at government sites where many of our Federal workers were located because these computers are usually tightly controlled. The trick in the scheme is that the form arrives with query values defaulted to request the desired service, and so, the service is not performed unless the user decides to return the form. This kind of email authentication was applied to:
 - Participation surveys: All participants in the Open Meeting had to complete a participation survey [3] running during several months before the event.
 - Linking comments: Both email and Web users needed to request a comment form while visiting the target node, and then, reply to the email forms they received. This email form contained the document identifier (PDI) for the target node and would accept a range of link types according to the link grammar.
 - Subscription: While visiting a hypertext node, both email and Web users could subscribe or unsubscribe to any comments attached to the node or topically-related nodes. Either choice on both the email and Web interfaces caused the system to send them an email form requesting confirmation. Because the hypertext

was fully categorized, the system knew the exact category combination required to subscribe to any node, and consequently, the users were freed from the need to specify the category combination themselves or for that matter to learn how to specify these in the first place. Similarly, a user could unsubscribe by visiting the hypertext node from which the subscription was originally requested.

- *Notification.* When user subscribed to a node in the hypertext structure, they would receive all comments and newsletters attached within the scope of the categories spanning the node. Of course, new attachments were not transmitted until a moderator accepted the comment. Unlike other comment contexts, here the comment stream arrived in a form that allowed immediate response because the system already had confidence in the subscribed users' identities. Although transaction costs were relatively higher for a Web user to submit their first comment, these costs were neutral for email users, and substantially lower for subscribed users because this notification capability relieved people from a need to constantly check to see if new comments were available. Thus, timely delivery of newly moderated comments kept the conversation gain at a level comparable to custom mailing lists tightly focused on specific regions of the hypertext structure.

Email Caching Strategy

During the Open Meeting, a simple governor limited the rate at which COMLINK accepted messages over SMTP and sufficed to keep computational load within hardware capabilities. The message traffic (e.g., submissions of surveys and comment) was heavily biased towards form processing that invoked relatively expensive database transactions. Fortunately, the SMTP protocol allows an email server to use unaffiliated store-and-forward mailers (Figure 9) out in the net-




Figure 9: HTTP and SMTP Traffic Flow

work to buffer the message traffic. This network buffering allows an overloaded email server to spread out message receipt and processing to periods of lower activity. This email strategy works as long as a server clears the backlog within 24 hours,

Web Caching Strategy

The realtime interactive properties of Web access threatened to put undue load on the main server (Symbolics XL1200 Lisp Machine) that had more than enough work managing the database as it handled all email communications and served web pages to moderators. In anticipation of this bottleneck, we deployed a caching proxy (CERN server) between the main database server and the Web users. (Figure 9) The only traffic at issue was Web-based browsing and searches. Two caching strategies were employed:

- *Forward caching* was combined with incremental page synthesis to maintain updated user and moderator views for browsing. As users submitted comments, the transaction that linked them into the moderator view would also invoke an incremental update of

all pages affected by the change in the moderator view.

- Moderator view: When, and if, a moderator accepted a user's comment, the transaction which changed its status to *accepted* also invoked an incremental update but this time not just to all the effected moderator pages but also to the relevant user pages. The moderator structure involved updating all superior pages to the root (one more page) because information about the review status of comments appeared all the way up.
- User view: For the user structure, moderator acceptance of a new comment required an update to the page summarizing the recommendation to show the new attachment and a incrementing of a number on the main topic node that indicated the number of comments below on a recommendation page.

An important computational property of this update strategy was that these updates only propagated changes upwards in the hypertext structure. Since the HTML structure was a shallow tree with rapid fan-in, this was quite efficient

and imposed no debilitating load on the backend server.

- *On-demand caching* with timeout was used for searches because we did not want to cache all possible searches. By caching searches with a fifteen minute timeout, we were assured of maintaining a relatively fresh cache while removing load from the backend server for high frequency searches. The same strategy was applied for both user and moderator views.

Although the Web caching strategy was designed to allow replication of the caching proxy, loading never became high enough to require additional hardware.

Use Patterns

The Open Meeting achieved its initial goal of attracting the attention of a large number of government workers from a wide range of organizations and geographic locations. Of the 4200 people who returned an online registration survey:

- 85% were government workers, and another 4% were state or local government workers.
- Respondents were drawn from all fifty states, twenty foreign countries, and even US Navy ships at sea. Fewer than 40% were from "inside the Beltway" (DC, Maryland, and Virginia).
- Their profile was similar to government workers in general for age (early 40s) and experience (54% were more than 10 years in government), but they were considerably more senior, technically oriented, educated and male than government workers generally. (60% vs. 25% in supervisory capacities; 47% vs. 15% with MA or more; 66% vs. 11% in information systems or engineering; 30% vs. 42% female.)
- In comparison to the estimates of Internet users at the time (December, 1994), this

population was older, more educated and had a lower percent of males.

The differences between the registrants and all government workers highlight the importance of commitment and access for wide-area collaboration. Supervisors and other managers have a greater interest in proposals for bureaucratic reform than those they supervise, and engineering and information system workers have greater access to email and World Wide Web. Access also explains some anomalies in the distribution of the registrants. Generally, the larger and more technology-involved organizations, like the Defense, NASA and Interior, had the largest numbers of registrants, but the under-representation of similar organizations, notably Treasury and NSA was due to their massive use of firewalls to restrict network access and penetration.

Hypertext Access

During the two week meeting itself, there were 35,000 Web accesses from nearly 1500 different hosts, exclusive of those by moderators and maintainers. While low by current standards for prominent government sites, e.g., the White House, this volume compares favorably with traffic at specialized and professional online forums. It is also large enough for the distribution of accesses over the pages to suggest how people navigate a complex information and conversation environment.

Table 2 shows a nearly consistent pattern of attrition of hosts with number of transactions from the root page, with the one exception being the smaller number who accessed the *Newsletters* than the *Comments*, although access to *Newsletters* was closer to the root. However, data for the second week alone would show a consistent pattern of attrition, since the *Newsletters* were not posted during the first week of the event. The newsletters in fact received considerable attention, either as a means by which users caught up on discussions or a substitute for reading the discussions themselves. Relatively few users moved

Table 2: Open Meeting WWW Accesses

Content Area	Hosts	Accesses	Root %
Root	1467	4000	100
Node Structure	1072	10600	73
Overviews	745	2900	51
Newsletters	318	2182	22
Recommendations, Actions	620	3700	42
Comments	514	8000	35
Search Interface	283	605	19
Full Taxonomy	37	78	3

outside the fixed-topology hypertext environment to accessing documents directly through the search interface.

These usage patterns suggest that most Web users explored several topics broadly but shallowly by looking at Overviews, some Recommendations and available Newsletters. They "felt" their way through the information, not sure of what they were seeking and more inclined to quit the search than go beneath the top level information. About one third of total users also explored one or two topics deeply by traversing the hypertext at the comments level and using the search interface. Because the surveys show the registrants as a whole were highly motivated, the main difference between the two groups of seekers was likely information literacy, with the in-depth seekers the more literate. The distribution of the nearly 1000 subscriptions to conversations, for which the in-depth seekers were necessarily responsible tells us members of this group remained tightly focused. The subscriptions clustered around those conversations which attracted the most comments, rather than being used as a means to branch out. The pattern is consistent with our earlier remarks that localizing communication is the means to handle large and complex information flows. It also underscores that both the computer literate and the less experienced need low transaction-cost navigation tools that sketch the whole domain and lead to specific topics of interest.

Linking Commentary

Out of 1300 comments submitted, moderators accepted 1013, which were contributed by 290 different individuals. Some conversations included ten or more speakers and had several branches. Although the moderators did not correct identifications of link types, we observed no mistaken identifications among submitted comments, except for one or two cases where a contributor may have deliberately mislabelled the type. The comments were generally positive and serious, with few flames. Half the comments were *Agreement* and 15% were *Disagreements*. *Questions* (167), *Alternatives* (106), and *Promising Practices* (72) accounted for nearly all of the rest, suggesting the contributors' willingness to use the meeting as a sounding board for ideas. Relatively few of the questions were answered (37), and almost no one used the more cognitively complex *Qualification* link (3).

While the ratio of contributors to accessors compares to the low ratios for newsgroups, the participation rate at the Open Meeting was higher for those reaching the comment level in this more complex environment. That result agrees with our theory regarding the effects of localizing communication. Low contribution levels are predictable in wide-area collaboration or problem solving, especially when participants have only general rather than specific functional roles. But, participants are likely to participate when conversations and work are localized and closer to their

experiences and knowledge. Because the Open Meeting structure localizes communication, we might expect typically higher participation in collaborations that run several months and are free of external distractions, like approaching holidays.

Participant Satisfaction

Web users were generally satisfied with the meeting and they complained only about an overwhelming amount of text to read and being forced to submit their comments and survey responses via email rather than the Web. The email users complained about the clumsiness of email to traverse hypertext, poor instructions and network delays. These differences reflect, on one hand, the greater interactivity of Web GUIs for hypertext, and, on the other, the fact that the least technically experienced people, who needed the most instruction, had the less sophisticated equipment. They also indicate the need for caution in planning to use email as transport in advanced information environments. Contemporary email does not easily support complex processes, like concurrent multilateral discussions of issues or wide-area collaboration. Instead of simplifying to accommodate the email limitations, the demand for broad participation which email satisfies should motivate upgrading the resources of the less technologically experienced and more basically equipped. But the SMTP transport need not be discarded; as we have seen, use of email subscriptions to track conversations spares the user the transaction costs involved in periodically revisiting the hypertext. Indeed recent trends toward closer integration of clients for reading email and browsing the Web may make SMTP a more useful transport media for wide-area collaboration systems.

Conclusion

The success of the Open Meeting demonstrates the importance of taxonomic decomposition and meaningful link types in the organization of wide-area collaboration. The meeting showed

that people can use typed links which they understand to create argument-structured discourse in a policy planning situation.

A desirable next step is to develop links grammar for decision processes. To generate the kind of knowledge process they seek, convenors of wide-area collaborations may select an appropriate set of link types and composition grammar. In the Open Meeting, the link grammar used did not provide for termination of a conversation. Other planning or action grammars can provide termination—like cloture in parliamentary debate. Interestingly, if we vary decision grammars according to different agent capabilities and functional roles, we start modeling reconfigurable organizations. The views into these processes can similarly be generated for different agents according to their capabilities and roles. Thus, power and social relations within the organization come to be defined by what an agent can do based on information accessible to the agent. This functional division of labor and knowledge, in turn, defines the organization as a process. Thus, experiments in wide-area collaboration promises contributions to new organization theories.

Another step devises link grammars for knowledge formation in scientific communities, and building research in scientific paradigms. Churchman [2] outlines methods of inquiring that can be constructed on the basis of several famous epistemologies. We should try to correlate each of these with a link semantics and explore their productivity in wide-area scientific collaboration. Finally, since wide-area collaborations will include the coordination of work as well as integration of information and opinions, we need to develop systems that can recognize collaborative situations, infer possible options, and recommend strategies or identify resources. These kinds of wide-area collaboration systems promise to help scientists conduct research more effectively as disciplines grow in complexity and knowledge advances more rapidly.

The World Wide Web offers *unprecedented* opportunities for wide-area collaboration at a

time when nothing less seems likely to cope with endemic and emergent global problems. We have argued that collaboration systems can begin to manage the complexity by supporting the specialization and localization of knowledge, planning and evaluation. Successful systems will then face the challenge of reintegrating all their partial results. ■

Acknowledgments

Benjamin Renaud contributed significantly to operation of the Open Meeting as well as the implementation and design of a number of the application components, including the moderator interface, generation and caching of virtual pages, and some email interfaces. Mark Nahabedian helped us recover from some disk drive failures. The Vice President's 1994 Open Meeting on the National Performance Review was a collaborative effort between The M.I.T. Artificial Intelligence Laboratory, The White House, National Performance Review, Lawrence Livermore National Laboratory, and Mitre Corporation. Randy Katz that made this project happen by bringing together the players, who included Larry Koskinen and Andy Campbell from NPR. Jonathan P. Gill and Thomas Kalil provided inspiration and critical support for the effort. Howard E. Shrobe helped with earlier versions of the Communication Linker System and provided endless moral support. This paper describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the M.I.T. Artificial Intelligence Laboratory's artificial intelligence research is provided in part by the Advanced Research Projects Agency of the Department of Defense under contract number MDA972-93-1-003N7.

References

1. T. Berners-Lee and A. Luotonen, "Web Interactive Talk," Dialectical collaboration system deployed on the net in June, 1994. Geneva: World Wide Web Group, CERN, 1994.
2. C. W. Churchman, *The Design of Inquiring Systems: Basic Concepts of Systems and Organization*, New York, Basic Books, 1971.
3. R. Hurwitz and J. C. Mallery, "Survey of Participants in the Vice President's Open Meeting on the National Performance Review," Cambridge: M.I.T. Artificial Intelligence Laboratory, December, 1994. <http://www.ai.mit.edu/projects/iitp/doc/open-meeting/participation-survey.html>
4. S. E. Keene, *Object-Oriented Programming in Common LISP: A Programmer's Guide to CLOS*, Reading: Addison-Wesley, 1989.
5. J. C. Mallery, "Semantic Content Analysis: A New Methodology for the RELATUS Natural Language System," in V. M. Hudson, *Artificial Intelligence and International Politics* Boulder: Westview, 1991. <http://www.ai.mit.edu/people/jcma/papers/1991-aiip/raiip3.html>
6. J. C. Mallery, "A Common Lisp Hypermedia Server," *Proceedings of the First International Conference on the World Wide Web*, Geneva, May 1994. <http://www.ai.mit.edu/projects/iitp/doc/cl-http/server-abstract.html>
7. J. C. Mallery, "The Communications Linker System: An Overview," paper presented at *The 1994 Meeting of the American Political Science Association*, New York City, September 1994. <http://www.ai.mit.edu/projects/iitp/doc/comlink/overview.html>
8. J. C. Mallery, "Dynamic Form Processing: A General Framework for Interactivity," paper presented at *The Second International Conference on The World-Wide Web*, Chicago, October 1994. <http://www.ai.mit.edu/projects/iitp/talks/www94b-dynamic-forms-abstract.html>
9. N. Rescher, *Dialectics: A Controversy-Oriented Theory of Truth*, Albany: SUNY, 1977.
10. R. Rao, W. M. York, and D. Doughty, "A Guided Tour of the Common LISP Interface Manager," *LISP Pointers*, 4 (1991).
11. G. L. Steele, *Common LISP: The Language*, Bedford: Digital Press, 1990. <http://www.cs.cmu.edu:8001/afs/cs.cmu.edu/project/ai-repository/ai/html/cltl/clm/clm.html>

About the Authors

Roger Hurwitz is a research scientist at the M.I.T. Artificial Intelligence Laboratory, an architect of the Open Meeting System, and a contributor to COMLINK Development. His work focuses on the modeling, measurement, and management of collective actions, public communications flows, and organizational intelligence. It draws on knowledge representation, discourse analysis, organization theory, and social science models and survey techniques. Hurwitz holds a Ph.D. from M.I.T., has taught at M.I.T. and the Hebrew University (Jerusalem), and has consulted for major communication companies and UN agencies. A member of several professional committees on the globaliztion of scientific communications, his publications concern paradigm development in the social sciences, patterns of media diffusion in late industrial societies, and process models of collective actions and knowledge formation. Current projects include the use of the Open Meeting System for interorganizational planning, collaboration, and knowledge

formation, and the creation of an analytic environment for tracking aggregate national and international media flows.

John C. Mallery is technical director of the Intelligent Information Infrastructure Project at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. His research interests center around new ways to model international interactions and new ways to incorporate advanced computational methods into interactive political communication. He has developed computer systems that construct natural language models from narrative text, learn if-then rules from complexly structured event data, and conduct automatic opinion surveys over global computer networks. An electronic publications system that he developed for use during the 1992 presidential campaign currently serves as the primary distribution hub for press releases by the U. S. White House. His current research explores intelligent information access, wide-area collaboration, knowledge-based organizations, and global knowledge webs.