

Introduction to the Electronic Symposium on Computer-Supported Cooperative Work

KEVIN L. MILLS

National Institute of Standards and Technology

Computer-supported cooperative work (CSCW) holds great importance and promise for modern society. This paper provides an overview of seventeen papers comprising a symposium on CSCW. The overview also discusses some relationships among the contributions made by each paper, and places those contributions into a larger context by identifying some of the key challenges faced by computer science researchers who aim to help us work effectively as teams mediated through networks of computers. The paper also describes why the promise of CSCW holds particular salience for the U.S. military. In the context of a military setting, the paper describes five particular challenges for CSCW researchers. While most of these challenges might seem specific to military environments, many others probably already face similar challenges, or soon will, when attempting to collaborate through networks of computers. To support this claim, the paper includes a military scenario that might hit fairly close to home for many, and certainly for civilian emergency response personnel. After discussing the military needs for collaboration technology, the paper briefly outlines the motivation for a recent DARPA research program along these lines. That program, called Intelligent Collaboration and Visualization, sponsored the work reported in this symposium.

Additional Key Words and Phrases: Asynchronous collaboration, computer-supported collaborative work, human-computer interaction, information management, multimedia collaboration, multimodal collaboration

PROMISE OF COMPUTER-SUPPORTED COOPERATIVE WORK

Computer science researchers have long envisioned using computers to support the work of teams. Research investigating CSCW first appeared coincident with the invention of time-sharing systems. The underlying technology of time-sharing, which enabled many users to simultaneously share a computer and its files, provided a platform on

which researchers could explore schemes for grouping information, sharing files, and jointly authoring documents. In the decades since the invention of time-sharing, information technology has advanced in some surprising directions. Personal computers (PCs) have replaced time-sharing systems, multimedia data has enriched simple text documents, networks have brought users together across vast dis-

The work discussed in this symposium was conducted while the author was on detail to the Defense Advanced Research Projects Agency (DARPA)

Author's address: National Institute of Standards and Technology, 820 West Diamond Avenue, Gaithersburg, MD 20899-8920; email: kmills@nist.gov.

Permission to make digital/hard copy of part or all of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.

© 1999 ACM 0360-0300/99/0600-0105 \$5.00

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE JUN 1999		2. REPORT TYPE		3. DATES COVERED 00-00-1999 to 00-00-1999	
4. TITLE AND SUBTITLE Introduction to the Electronic Symposium on Computer-Supported Cooperative Work				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) National Institute of Standards and Technology, 820 West Diamond Avenue, Gaithersburg, MD, 20899				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 11	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

tances, and a wide range of interaction modes have augmented simple command-line interaction. Even today, we can see changes ahead in the growing population of specialized information appliances with built-in wireless communications. Such appliances will enable individuals to improve their ability to manage personal information. Yet, the main challenge for CSCW researchers remains the same — enabling people to work effectively together in teams mediated through computers.

The CSCW challenge has particular interest for modern society because two trends continue to accelerate. First, with increasing frequency, people work together in cross-organizational teams that form quickly for particular goals and then evolve with changes in the nature of the tasks needed to achieve those goals. Such teams must quickly establish a work plan, divide up tasks, and determine means of coordination and self-regulation. In a growing number of cases, team members work asynchronously, but their work must still be coordinated effectively. The second accelerating trend places extensive computing power into the hands of individual office workers. This means that professional work has become digital. Since individuals work through computers, teams will naturally do so as well. This juxtaposition of personal computing power with the need to work digitally in teams presents the CSCW research community with an urgent challenge: improve the ability of teams to work together through networks of computers.

SPECIFIC CHALLENGES FOR CSCW IN THE MILITARY

This challenge has particular import for the U.S. military because military operations have entered a new era of uncertainty, requiring agility, rapid response, and innovative teamwork. In this heterogeneous and fast-changing environment, information technology in support of adaptive teams can be a critical dis-

criminator for the U.S. military. For this reason, the Defense Advanced Research Projects Agency (DARPA) challenged the computer science research community to develop some new collaborative computing capabilities that address five specific needs faced by military teams.

- (1) *Collaboration across heterogeneous bandwidth and devices.* Military teams must collaborate across multiple media even when some team members participate through handheld devices at the end of low-bandwidth links, while others work with high-speed network connections, large-screen displays, and powerful desktop computers.
- (2) *Collaboration using natural modes of interaction.* Individual soldiers must interact in collaborative sessions using the modes most appropriate for their situation. Speech recognition, gesture, eye-tracking, and verbal and tactile feedback must all work together seamlessly, along with the more conventional modes of interaction.
- (3) *Ready access to information affecting collaboration.* The immense amounts of information produced and consumed by military teams must be made more accessible. Specifically, as documents are generated or discovered they should be routed automatically to team members whose task might require them. In addition, all interactions among team members, whether audio, video, or event data, should be archived and indexed to support rapid search. Beyond mechanisms to manage a team's information, automated mechanisms are needed to help people discover each other when they could potentially profit by collaborating.
- (4) *Collaboration without continuous connectivity.* Military units need to collaborate asynchronously. Given

the nature of the military situation, where links and nodes go up and down and where collaborators come and go, collaborative sessions must persist beyond the lifetime of any current members. Even when no members are present, the session must persist.

- (5) *Evaluation of effectiveness before deployment.* New tools, techniques, and metrics are required to enable effective and efficient evaluation of collaboration technology prior to its deployment for military operations.

The epitome of military collaboration appears during crisis response planning and execution. As a crisis nears, a crisis action team (CAT) in a military headquarters dusts off skeleton plans for similar crises and for operations in a similar geographical area. It is likely the CAT members did not develop the skeleton plans. In fact, this might be the first time the CAT members have worked together on an assignment. The CAT must formulate a set of tasks to quickly complete a specific plan for the current crisis and geography. Often this will involve establishing contact with local operatives at or near the crisis site. Sometimes this contact comes through limited bandwidth or through sporadic connections. As the crisis unfolds, the nature and duration of local contacts can change. Planning deliberations also often include high-level discussions of political and military options. The intent of the commander-in-chief typically shapes the direction of a crisis plan. Bridging the different people, sites, and equipment involved in crisis response presents a major challenge. Over the course of a few hours or days, the current situation must be determined and an adequate plan developed. At that point, a crisis response is "good to go."

When given a "go," the crisis response execution unfolds under the watchful eye of the CAT. As the situation changes, the plan must be adapted. Local reports, while vital, can be sketchy.

Bandwidth can degrade, links can come up and down, and CNN crews can be captured and disabled. Units assigned to tasks can be destroyed or otherwise incapacitated. The CAT must recognize such occurrences and respond quickly and effectively with changes in plan. This might involve drawing on new units to participate. Such units must be brought on board quickly with both the plan and the situation. If a crisis goes on long enough, the CAT itself will undergo shift changes. Newly arriving members must be updated rapidly regarding the situation and changes in plan. Eventually, the crisis is resolved, the troops return, and the CAT for the specific crisis can be disbanded. While responding to a crisis places high demands on the people and systems of the U.S. military, the demands are much higher when multiple crises occur simultaneously. In such cases, an individual can serve at once on multiple crisis action teams, and so must carefully manage time and attention. This scenario illustrates how military crisis management personnel face some serious and difficult collaboration challenges.

RECENT ADVANCES IN CSCW

Addressing such challenges requires contributions from researchers across a range of specialties, such as networking, multimedia, artificial intelligence, visualization, and human-computer interaction. In order to bring such diverse expertise together, in 1997, DARPA initiated a three-year research program, known as the Intelligent Collaboration & Visualization (IC&V) program (for more information, see <http://www.darpa.mil/ito/research/icv>). The seventeen papers collected in this symposium describe some of the advances achieved by researchers working within the IC&V program. The papers are organized in four categories: (1) multimedia and multimodal collaboration; (2) information management for collaboration; (3) asynchronous collaboration;

and (4) test and evaluation for collaborative systems. A brief synopsis of each paper follows, organized by category.

ADVANCES IN MULTIMEDIA AND MULTIMODAL COLLABORATION

In "MASH: Enabling Scalable Multi-point Collaboration," **Steve McCanne et al.** describe a new generation of multimedia conferencing software for the Internet. The MASH software, based on multicast technology, makes possible synchronous collaborative sessions among users connected through a diverse set of bandwidths and access devices, including personal digital assistants, Web-clients, desktop workstations, and large-screen displays. In addition, MASH makes possible archiving, indexing, and replay of multimedia collaborative sessions. In order to achieve these new capabilities, the MASH architecture includes a number of novel technologies: programmable media gateways, scalable reliable multicast protocols, and streaming media transport protocols. Beyond these valuable research results, MASH has had early practical results because the new conferencing tools from the project are already in use at PACOM to support various planning activities among teams in Hawaii, San Diego, and Washington, DC. The paper provides an overview of the new services available in MASH, as well as the underlying technology innovations that make the services possible.

In "Semantic Multicast: Intelligently Sharing Collaborative Sessions," **Son Dao et al.** describe an architecture, built above Internet multicast, to distribute multimedia information efficiently on the basis of relationships between user needs and network topology. The architecture has four main components. First, graph discovery algorithms organize users and information sources into distribution trees. Second, various realtime and offline algorithms extract metadata from multimedia content for redistribution on metadata channels.

Third, based on the metadata and the interests of users, multimedia streams are matched and distributed along appropriate network channels. Fourth, a multimedia database provides a repository that can be enriched with offline processing, searched for relevant information, and replayed for users who missed live multimedia sessions. The paper provides an overview of the architecture and its components.

In "DISCIPLE: A Framework for Multimodal Collaboration in Heterogeneous Environments," **Ivan Marsic** discusses a collaboration system, based on a distributed Java-beans event model, which integrates two novel concepts. First, desktop users collaborate through a multimodal interface that integrates gesture recognition, gaze tracking, speech recognition, and speech synthesis with the usual windows, mouse, and keyboard interaction modes. A number of novel technologies implement these modes. For example, a microphone array frees collaborators from a head-worn microphone and a force-feedback tactile glove provides the user with interaction cues. The second novel concept in DISCIPLE is the ability to tradeoff remote method invocation against the use of mobile code. Based on observing interobject message traffic for specific functions, a rule set is learned that enables real-time, knowledge-based decisions about when objects should be moved and when objects should be accessed remotely.

In "Virtue/Orbit: Collaboration and Visualization Toolkits," **Daniel Reed** and **Simon Kaplan** discuss two software toolkits that can be used either independently or together to improve collaboration in data-intensive environments. Orbit provides a software framework to support users simultaneously engaged in multiple collaborative tasks. The Orbit framework visually manages multiple contexts, while enabling users to exercise variable control over their level of awareness of each context. Virtue enables users to collaborate im-

mersed within massive data sets and integrates multiple modalities into a software and hardware system that helps users visualize, manipulate, and annotate complex data. In addition to discussing the main concepts underlying the two toolkits, the authors describe how the toolkits can be used together and sketch a brief scenario to illustrate the collaborative power that can be achieved.

In "Creation and Performance Analysis of User Representations in Collaborative Virtual Environments," **Kevin Martin** reports results from an experiment to determine the frame update rate achieved when using various avatars to represent users in a virtual environment. Martin considers nine representations for head-only avatars and four full-body avatars, and also measures the avatars with and without movement. The results of this experiment should help to inform designers who must choose among particular avatars when designing collaborative virtual environments.

The five papers are on a spectrum from near-term to longer-term. The MASH project has already delivered beta software that permits users to collaborate across a range of bandwidths and devices. The project is now exploring mechanisms for constructing a conference bus to support shared control of devices by the members of a collaborative session. Beyond that, the project expects to deliver improved scalability for reliable delivery of information among groups and to archive multimedia sessions, so that others may search and replay them later. Complementing the archiving capabilities of the MASH toolkit, the Semantic Multicast project will provide technology to aid in the indexing and distribution of multimedia sessions, both in realtime and during replay, while conserving overall network bandwidth. The DISCIPLE project complements the capabilities of MASH in two ways. First, DISCIPLE presents a multimodal interaction model that MASH could incorporate into its confer-

ence bus. Second, DISCIPLE will provide knowledge-based rules that MASH could exploit to select which MASH objects should be replicated for local access. The Virtue project contributes an advanced form of multimodal interaction through which selected collaborators can immerse themselves in the structure and behavior of massive data sets. The Orbit project contributes an information structure and user interface tuned for collaborators who must juggle simultaneously multiple collaborative sessions. One could easily imagine a future software toolkit that merges the best ideas from these projects into a versatile and adaptive collaboration environment. The final paper in this set tackles a specific issue in multimodal collaboration. Low bandwidth often prevents the injection of video representations into a collaborative session. In such cases, avatars can be loaded locally on each machine and then animated based on small messages sent sporadically to represent the state of remote collaborators. In such situations, the specific avatars to choose might depend upon the amount of processing power needed to render and animate them. Martin reports results that can guide designers who must match avatars with processing power.

ADVANCES IN INFORMATION MANAGEMENT FOR COLLABORATION

In "Rough'N'Ready: A Meeting Recorder and Browser," **Francis Kubala et al.** describe a prototype system that automatically constructs a browsable index from audio streams. The system integrates four individual technologies (speech-to-text transcription, speaker demarcation, named entity extraction, and topic classification) to consume audio streams and to produce a tagged metadata file containing an index into the raw audio. The system also includes a Web-client plug-in that helps the user browse and query the audio stream for particular speakers, topics, and named entities. The authors discuss the perfor-

mance of each technology and outline their future research plans. To date, the system has been applied to broadcast-quality audio from news sources. Future plans include investigating the feasibility of applying the technology to low-quality audio in conversational applications, such as multimedia conferencing.

In "Dynamic Collaborator Discovery in Information-Intensive Environments," **David Payton et al.** present a novel way to discover potential collaborators based on comparison of individual patterns of Web browsing. As individuals browse the Web, their accesses are typically logged at a firewall. By converting these accesses into graphs for each user and then matching the graphs, similarities and differences can be detected and users brought into contact. Payton et al. describe the architecture of a prototype system deployed inside their company. The authors also address the mechanisms to overcome privacy concerns that might otherwise be raised by their approach. They briefly present visual designs that enable users to explore possible matching interests with other users. Finally, the authors discuss experimental results obtained by using the system.

In "Informedia Experience-on-Demand: Capturing, Integrating and Communicating Experiences Across People, Time and Space," **Howard Wactlar et al.** discuss a prototype system that uses speech, image, and natural language processing, combined with GPS information, to capture, integrate, and communicate personal multimedia experiences. The main premise of this research is that individuals, augmented with audio, video, and GPS recording devices, can capture a complete multimedia record of their experiences. Such experiences, uploaded and integrated in a database, can then be indexed, integrated, summarized, and searched. The authors identify a number of potential applications for such a system, including allowing a military commander to better understand an unfolding situa-

tion by combining the individual views of subordinate units.

In "Task-Based Information Management," **Michael Wolverton et al.** discuss a prototype system for automatically routing incoming documents to specific individual team members. The system relies on Bayesian task graphs, where each node in the graph contains a key question, a template of parameters, and a set of relevant keywords. Probabilities are assigned to each item in each node. A Bayesian task graph, when instantiated with specific parameters and probabilities, is used to generate a query weighted for the related task. As documents arrive into a repository, the generated queries are applied and the documents are routed to users who meet the requirements of those who are needed to help complete a task. After describing how the system works, the authors indicate the current status of the prototype and the plan for additional research.

During large-scale undertakings that involve numerous collaborative sessions with many participants per session, massive amounts of multimedia information are generated, shared, and consumed. The four projects outlined above consider means to improve the accessibility of such information. Three of the projects, Rough'N'Ready, Dynamic Collaborator Discovery, and Experience-on-Demand, rely on *post-hoc* analysis of passively recorded information, while the fourth, Task-Based Information Management, actively queries information for matches against task models. Rough'N'Ready aims to produce multi-valued indices for a collaborative session, based solely on analyses of the session audio. This technique should prove very versatile because most collaborative sessions involve some form of voice communication. The second project, Dynamic Collaborator Discovery, performs graph matching against passively collected records of Web-surfing behavior among a group of people. The matching is used to identify and suggest people who might profit from

collaborating with each other. The third project, Experience-on-Demand, collects and integrates multimedia data from the viewpoint of many individuals. Once collected, the multimedia data is fused, indexed, and summarized so that individuals can query the entire situation for patterns and information not known to them when looking strictly from their own viewpoint. The fourth project, Task-Based Information Management, requires preparatory work to design, select, and parameterize task models prior to beginning collaborative sessions. During collaborative sessions, active queries, generated from the instantiated task models, monitor newly discovered or generated documents and route them to particular people that might need to see them in order to complete a task. Combining all of the monitoring and indexing methods discussed here can substantially increase the usefulness of any archive of collaborative sessions.

ADVANCES IN ASYNCHRONOUS COLLABORATION

In "Interface Issues in Computer Support for Asynchronous Communication," **Jim Morris et al.** argue that generic computer tools, such as email readers, provide insufficient support for collaborators working together intermittently on a specific task. Instead of generic tools, they suggest that next-generation tools for asynchronous collaboration need to support emergent, task-specific representations that can be shared in whole or part under control of individual collaborators. In addition, they posit that effective tools for asynchronous collaborators must provide a means to grasp the context of a collaborative session quickly, whenever collaborators return to a joint task. The authors support their ideas by reporting results from a case study. While not covered specifically in this paper, PACOM has expressed interest in using a prototype tool designed by the authors to enable

flexible querying and presentation of incoming electronic mail.

In "Integration of Synchronous and Asynchronous Collaboration Activities," **Larry Jackson** and **Ed Grossman** describe extensions to Habanero, a Java-based system for synchronously sharing applications. These extensions enable Habanero sessions to persist even when all participants are absent, and also permit collaborators to participate in Habanero sessions without using Habanero. In addition, Jackson and Grossman describe how Habanero can be used to collaborate effectively when a session can obtain only limited bandwidth. Habanero, including its extensions for persistent sessions, is freely available for download and use.

In "Consistency Management for Distributed Collaboration," **Sankar Virdhagriswaran et al.** propose a new mechanism for configuration management in distributed systems where teams asynchronously update jointly authored documents. The main contribution of their approach is to separate proposed changes to a document space from the orthogonal issues of concurrency control and repository management. Specifically, as a collaborator updates a copy of a shared document, the updates are recorded in change proposals that track information the collaborator expects to revise and that record consistency relationships which must be maintained. Once recorded, change proposals can themselves be treated as shared documents. The change proposal mechanism is implemented in a prototype system for asynchronous collaboration. The proposal outlined in this paper is included in Concorde, a new collaboration system offered by Crystaliz.

In "Intermediary Architecture: Interposing Middleware Object Services Between Web Client and Server," **Craig Thompson et al.** present an architecture that places distributed object services between a Web client and server. By adding a distributed object services plug-in, Thompson hopes to improve the ability of Web-based technology to im-

plement distributed applications. The paper discusses two sample applications: a Web annotation service and a personal Web performance monitor. The paper also discusses some lessons learned during the design and implementation and identifies some open issues that remain to be addressed.

Perhaps the most difficult issues for collaborative systems revolve around providing support for asynchronous collaboration, where progress occurs based on contributions from independent people who might not meet at the same time, even in electronic form. The four papers outlined above make modest contributions toward solving some of these difficult issues. The first paper investigates how to empower users to express task-specific information in a succinct and meaningful form. Should its ideas prove feasible, then users no longer need to resort to heroic means to track the state of task-specific collaborative activities. The second paper describes extensions made to Habanero, a successful system for synchronously sharing applications. The extensions provide for persistence of collaborative sessions and also provide a mechanism through which nonHabanero users can participate in Habanero collaborative sessions. This platform could provide an interesting base on which researchers could experiment with user interface improvements for asynchronous collaborators. The third paper contributes a mechanism to propose changes to systems of replicated objects independently. Such a mechanism provides a fundamental service for asynchronously collaborating around copies of documents. The fourth paper discusses an approach to inserting distributed object systems between a Web client and server. Integrating distributed object services into the Web client-server model improves the flexibility available to build asynchronous collaborative applications. While making only modest contributions, the four papers discussed here attempt to improve asynchronous collaboration processes. Ultimately, improvements in

asynchronous collaboration will likely provide the largest payback for the research invested.

ADVANCES IN TEST AND EVALUATION FOR COLLABORATIVE SYSTEMS

In "Evaluation for Collaborative Systems," **Laurie Damianos et al.** discuss the challenges inherent in evaluating collaboration environments and present a four-level framework to evaluate collaborative systems using a scenario-driven approach. Each level of the framework includes components and related metrics. At the top level, the framework considers the characteristics of a collaborating team, as well as requirements for work tasks, for transition between tasks, and for social protocols. At the second level, the framework encompasses specific capabilities to support the requirements of a team. For example, such capabilities might include shared workspace, shared view, and group communication. Pushing lower, the third level documents the menu of services that could be used to provide the capabilities of level two. For example, services could include whiteboards, text chat, 3D visualization, application sharing, and so on. At its foundation, the framework considers specific products and implementations that can provide needed services. Using this framework, evaluators can identify the specific level and component they wish to evaluate and the metrics to use. In addition, given a set of requirements, users can navigate to the capabilities, services, and implementations that can support their requirements. Conversely, given a specific implementation, developers can identify the services, capabilities, and requirements that can be provided.

In "Reconfigurable Distributed Scripting," **M. Ranganathan et al.** describe a distributed scripting system, based on mobile code, which can be used to configure and test collaborative applications from recorded log data; thus providing repeatable testing without the

need to repeatedly assemble teams of people to conduct the tests. The authors explain an innovative technology, called *Mobile Streams*, which underlies this distributed scripting system, and then describe the use of Mobile Streams to build a self-reconfiguring chat application. The authors also discuss the mechanism to control configuration policy on three levels: system-wide, node-wide, and stream-specific. A system based on mobile code must support configuration control policies in order to protect the integrity of system components.

In "The MITRE Multimodal Logger: Its Use in Evaluation of Collaborative Systems," **Samuel Bayer et al.** briefly present the capabilities of a multimodal logging tool and then illustrate its application to help evaluate two collaborative applications. A particularly interesting aspect of the logger is its ability to visualize many multimodal data streams in parallel along a time scale. In effect, the visual presentation becomes a user-friendly index into the log. Logged data can also be exported in tab-delimited format to support spreadsheet analysis.

In "The MITRE Map Navigation Experiment," **Laurie Damianos et al.** report results from an experiment designed to assess the effectiveness of various collaboration services, such as audio, text chat, and whiteboard, when used for tasks that require information sharing and collaborative planning. To design the experiment, the authors used the four-level evaluation framework described above; to collect and analyze data, they used the multimodal logger discussed in the preceding paper. For a collaborative route-planning task, the authors found that collaborators produced successful routes significantly faster when audio was used. In addition, collaborators engaged in more frequent interaction and were more satisfied with the collaborative experience when audio was used. The quality of the task results did not vary significantly, whether or not audio was used.

As shown in these four papers, testing

and evaluating collaborative systems presents some unusual challenges. Evaluating collaborative systems requires enlisting, or simulating, the presence of a group of people to participate in test scenarios during which interactions are logged in order to make possible measurements of interest. The first paper above describes the framework for accomplishing all of this. Evaluators can use the framework as a guide and can contribute substantively to the framework by submitting scenarios, metrics, and logging and test tools. The second paper describes a tool that enables multiuser interactions to be logged and replayed later in a range of controllable configurations. Tools of this sort enable collaborative test scenarios, involving real users, to be reused in various ways. When conducting a test, logging and analyzing the multimodal interactions of a collaborative group can be challenging. The third paper shows one means of performing such logging and analysis. When all of these elements come together, as illustrated in the fourth paper, a productive experiment can be conducted and valid results reported.

CONCLUSIONS

Computers and networks hold great potential to improve the ability of people to work together effectively in teams. Yet despite advances in processing power and bandwidth, the exponential growth in the size of the World-Wide Web, and the nearly ubiquitous deployment of a *de facto* standard commercial operating system, effective computer-supported collaboration eludes us. The DARPA IC&V program explored the premise that significant advances in computer-supported cooperative work can best be achieved by mixing together ideas from the most talented researchers across a range of disciplines, including networking, multimedia, artificial intelligence, visualization, and human-computer interaction. As shown in the papers in this symposium, the results to

date have proven significant; however, the best is yet to come, as research connections made during the program foster innovative ideas for future research.

ACKNOWLEDGMENTS

DARPA provides a great environment to identify the best ideas from the research community and then to push those ideas as far as possible, given the resources available. Without DARPA, the work in this symposium would not have occurred. I extend my personal thanks to all the previous and current program managers who contributed two to four years of their own careers to make DARPA possible. In addition, I appreciate the efforts of the National Science Foundation to foster the evolutionary research base from which DARPA programs can select and amplify the most promising results. Most of all, I am grateful to the scores of researchers, including many bright, imaginative, and energetic graduate students, who contributed proposals, software, hardware, papers, and ideas to DARPA. These institutions and individuals have raised to new heights my optimism for the future of information technology. Tomorrow is in good hands.

BIBLIOGRAPHY OF CSCW ON THE WEB

- (1) *An annotated bibliography of computer-supported cooperative work*. This bibliography is a part of the Computer Science Bibliography Collection, 7/3/98. <http://www.cs.monash.edu.au/mirrors/bibliography/Distributed/CSCWBiblio.html>
- (2) *DCR: Computer-supported cooperative work sites*. Internet resources on computer-supported cooperative work, compiled by Lee Honeycutt (honeyl@rpi.edu). CSCW Web resources; P. S. Malm's unOfficial Yellow Pages of CSCW; Collabra's Groupware; Central WWW Virtual Library's CSCW Entry; EuroPARC's CSCW, 2/6/97. <http://www.dcr.rpi.edu/cscw.html>
- (3) *CSCW Bibliography*. CSCW pages @ Technische Univ. Muenchen, Germany. CSCW bibliography, CSCW home - CSCW bibliography - CSCW links. This is a gateway to the bibliographic database on CSCW and related topics maintained at Applied Informatics, 1/10/99. <http://www.telekooperation.de/cscw/cscw-biblio.html>
- (4) *An annotated bibliography of computer-supported cooperative work*. This bibliography is a part of the Computer Science Bibliography Collection, 7/3/98. <http://sgi.felk.cvut.cz/~biblio/Distributed/CSCWBiblio.html>
- (5) *CSCW*. CSCW research. My work on CSCW information. Bibliography of my collections; USENET news group; comp.groupware; FAQ of bibliography (local copy); FAQ of products (local copy); groupware, 9/30/98. <http://salmosa.kaist.ac.kr/~wwyi/cscw.html>
- (6) *Abstract: An annotated bibliography of computer-supported cooperative work*. S. Greenberg, Dept. of Computer Science, Univ. of Calgary, Calgary, Alberta, Canada T2N 1N4, 12/26/98. <http://www.cpsc.ucalgary.ca/grouplab/people/car1/g1Papers/1991/91-Biblio.SIGCHIBull/abstract.html>
- (7) *ACM*. ACM Conference on Computer-Supported Cooperative Work and Human-Computer Interaction. Bibliography (HCIBIB Project); CSCW'88 (Portland, OR); CSCW'90 (Los Angeles, CA); CSCW'92 (Toronto, Canada), 11/3/95. <http://fas.sfu.ca/0/projects/ElectronicLibrary/Collections/CMPT/cs-conferences/CSCW/>
- (8) *Bibliography*. Bibliography groupware technology and applications. D. Coleman (collaborative strategies); R. Khanna, Stanford Univ. (collaborative

- strategies). www.collaborate.com/resources/resources. PC Magazine (ago95), Byte, PC Week, PC World, www.witi.cs, 4/17/98. <http://selva.dit.upm.es/~proy/doct/dlarra/groupware/tsld002.htm>
- (9) **CSCW Resources.** The following CSCW and groupware resources are available: Distributed systems and industrial automation; group abstracts for the Journal of Virtual Environments (JOVE); Arnaud Dufour's home page; Agent Info AAI/AI-ED Group, 11/29/95. <http://www.scis.nova.edu/~sharma/resour.html>
- (10) **Team IT-Collections.** The Forum for Computer-Supported Collaborative Working. Collections of information bibliographies, directories, and collections of technical reports and papers, information on the Internet. The excellent CSCW pages at the Technische Univ. Muenchen, 9/29/96. <http://www.csc.liv.ac.uk/~team-it/collect.html>
- (11) **The Groupware Home Page.** A source for groupware information, where we could write a definition of groupware. Groupware Links; CSCW Yellow Pages; Groupware Yellow Pages; CSCW Directory DT - 5 Research Project; CSCW Bibliography WWW Collaboration Projects. 11/15/95 <http://www.isu.edu/~aytekgreg/groupware.html>
- (12) **Towards a CSCW framework for scientific cooperation in Europe.** Computer science, computer-supported cooperative work (CSCW), scientific cooperation for researchers, professionals, and scientists. Level: Monograph, H.P. Lubich, Stiftung SWITCH, Zurich, Switzerland. Towards a CSCW framework for scientific, 7/15/96. <http://www.springer-ny.com/catalog/np/mar95np/DATA/3-540-58844-2.html>