Corrie, B., Wong, H., Zimmerman, T., Marsh, S., Patrick, A.S., Singer, J., Emond, B., & Noël, S. (2003). Towards quality of experience in advanced collaborative environments. Paper presented at the *Third Annual Workshop on Advanced Collaborative Environments*, Seattle, Washington, June 22.

Towards Quality of Experience in Advanced Collaborative Environments

Brian Corrie[1]
Hong-Yee Wong
Todd Zimmerman

New Media Innovation Centre
600-515 W. Hastings Street
Vancouver, BC, Canada
V6B 5K3

Stephen Marsh
Andrew S. Patrick
Janice Singer
Bruno Emond

National Research Council of Canada
1200 Montreal Road
Ottawa, ON, Canada
K1A 0R6

Sylvie Noël

Communications Research Centre
3701 Carling Avenue
Ottawa, ON, Canada
K2H 8S2

Abstract

Collaborative environments have the potential of truly supporting distributed teams but there are still a number of barriers preventing seamless collaboration. These barriers are a result of problems in the following four domains: 1) a lack of understanding of the tasks that people perform when they are collaborating; 2) a lack of understanding and fulfillment of users' needs during collaborations; 3) the high complexity of collaboration services; and 4) limited access to a wide variety of technologies for use in complex, heterogeneous, and dynamic environments. The goal of our Advanced Collaborative Environment (ACE) project is to support seamless collaboration by removing these barriers and improving the users' Quality of Experience (QoE). This paper describes our view of a QoE-ACE architecture of the future, an architecture that takes into account not only available and emerging technologies, but also the users, their individual needs, and the uniqueness of the tasks they set out to pursue.

1. Introduction

The goal of an Advanced Collaborative Environment (ACE) is to bring together the right people and the right data at the right time in order to perform a task, solve a problem, or simply discuss something of common interest. We conjecture that it is not enough to be able to provide tools to users and hope that they use them effectively. It is our belief that, by focusing on the individual and social needs of the users in the context of the collaboration task, we will be able to adapt to the requirements of the situation and deliver the best Quality of Experience (QoE) to each user in the environment, regardless of task, technology, or individual [5] [9] [14] [17].

The main objective of this paper is to describe our view of an ACE architecture of the future, an architecture that takes into account not only available and emerging technologies, but also the users, their individual needs, and the uniqueness of the tasks they set out to pursue. As a starting point, we will present some barriers to the use of collaborative environments. We then introduce the concept of QoE in ACE. QoE is central to the current research as it is the end-result of the interaction of people with collaboration technologies. Following this, we outline a vision of an ACE sensitive to QoE issues. From this vision, we introduce the four domains of intervention to assure quality of experience (task, needs, services, and technology). Finally, the paper presents a QoE task/needs matrix, a data model for supporting ACE applications, our current implementation status, and a plan for future work.

2. Barriers to the Use of Collaborative Environments

Collaborative environments today involve several barriers to use and adoption [17]. There is, for example, the complexity of existing tools which discourages users [11] without continuous support from an expert node manager, while poor sound [4], video [10] [13] [21], or both [12] [19] limits usefulness. In addition, tools may

[1] Contact e-mail: Brian.Corrie@newmic.com

arbitrarily limit the number of users who can participate simultaneously.

There is also the fact that technologies in use today do little to adapt to the task at hand [7]. This is problematic since collaboration is never the same from one session to the next as a result of the uniqueness of individuals and the broad spectrum of tasks that can be accomplished through collaboration.

Moreover, from one moment to the next, the collaboration environment itself may change. Noise pollution, network performance, software glitches, changes in lighting, and other characteristics all contribute to the nature of collaboration, especially across a distance.

Additionally, there exist a number of different collaboration tools today, ranging from powerful video and audio conferencing tools to the more common tools of e-mail, pen and paper, and the telephone. At any time, one specific system or mode of interaction may be more appropriate than another. A truly successful collaborative environment should support the most appropriate tool for the tasks at hand, and provide other tools that can enhance collaboration.

Finally, we believe that powerful systems can be developed that are capable of working on whatever platforms are available at any given time. This principle applies if the platform is a traditional telephone, a handheld device, a desktop computer or a room-based collaborative environment. We do not claim that the same experience is achievable with, for example, the telephone when compared to a large wall-based videoconferencing system, but we do claim that we can maximize the QoE for each user taking their available technologies into account.

3. Quality of Experience

Quality of Experience [3] [1] can be defined as the characteristics of the sensations, perceptions, and opinions of people as they interact with their environments. These characteristics can be pleasing and enjoyable, or displeasing and frustrating. In the current context, QoE is the end result of the interaction of people with collaboration technologies and distant partners, and ensuring a good experience is the goal when high user satisfaction is desired. Thus, QoE is how the user feels about how an application or service was delivered, relative to their expectations and requirements. QoE can mean different things for different applications. For example, a high QoE for an audio application might be related to the sound fidelity and ability to smoothly take turns in a conversation. A remote video application might have a high QoE if the video image is large and clear when presented to the user.

QoE can be contrasted to two similar terms (see [3]). Quality of Service (QoS) refers to technical measures that can be taken to improve the flow of data on a network (e.g., DiffServ), while Quality of Transmission (QoT) refers to the actual flow of data on a network, as might be expressed in metrics like “throughput” or “packet loss rate.” Maintaining distinctions between QoE, QoS, and QoT can be valuable for understanding how different components of a system interact, and how one set of characteristics can affect the others (e.g., good QoS measures can lead to good QoT, which will result in a good QoE).

4. The Vision of a QoE-Sensitive ACE

Placer Dome is a large hard-rock gold mining company with mines and offices distributed all over the world. On a regular basis, Placer employees need to meet to plan for the year's exploration and mining activities. These meetings include people from a wide range of locations around the world (head office, mine site, national office) as well as from a wide range of areas of expertise (head office management, mine managers, geophysicists, and mine engineers). Because of the complexity of the data that is being discussed and the wide variety of professional expertise that is involved, 3D stereoscopic visualization is considered essential to providing a consistent view of the mine geology to all participants. The meetings are formal, include a large number of people (over 10), and their purpose is both to generate ideas and present plans. The requirements for such a meeting are immersive 3D visualization, good quality auditory communication, and good quality visual communication between collaborators [16].

To meet these requirements Placer would traditionally fly participants to a single location, usually one with a visualization facility, to hold these meetings. Recently, Placer has started to explore using multiple visualization facilities, connected together over a network with a shared 3D visualization tool. Because of the complexity of the data set and the visualization requirements, standard video conferencing tools do not suffice.

Setting up such a meeting is a truly daunting task that includes the transfer of data sets to the appropriate computers, starting up the 3D application on the correct data set, and making the appropriate audio, video, and 3D application connections between the sites. Expert knowledge is required in a number of areas: the audio and video applications and their controls (to provide the appropriate audio and video quality); the 3D application and how it is configured for collaboration; bandwidth availability and requirements; and IP numbers and ports for connecting applications. Due to the complexity of this environment, any technical problems with hardware, software, or network connectivity, or even something as simple as someone typing in an IP number incorrectly, will result in a very poor collaboration experience for all of the users. Requiring this level of knowledge from such a set of users is unacceptable.

We envision a world where QoE is guaranteed. The users walk into a collaboration room. The system recognizes them, lists the set of tasks that they have carried out in the past, and asks them what they want to do. They say they want a formal meeting, with a large number of people, over multiple sites. They state that it requires a 3D collaboration capability using geophysics data and they tell the system where the data is. The system asks with whom they want to collaborate. Upon learning of the other collaborators, the system determines their locations. The system determines that none of the currently existing session definitions are suitable for this meeting, and therefore the system creates a new collaboration session. Based on the type of task the system determines that audio, 3D visualization, and video (in that order of importance) are the critical components. By understanding both the importance of the collaboration services and the connectivity available between the key sites, the system can set appropriate quality settings for the various applications.

When the collaboration is initiated at each site, the system manages the transfer of the data set, the launching of the applications, the connection of the applications, and the quality levels of the application parameters (audio, video). After the two main sites join the session, the applications are linked and the collaboration can proceed. In this scenario one of the users cannot make it to either visualization room as he is held up at the airport. The user connects to the session via their laptop on the wireless network at the airport. The system recognizes that this user has a scaled down 3D capability on their laptop, but does not have the bandwidth to run audio, 3D, and video services at the same time. Because of the QoE needs of the collaboration, only the audio and the 3D application are started. During the collaboration, the network connectivity between the two main sites suddenly becomes congested. The system analyses the QoE requirements of the session and realizes that video quality is the least important component. It degrades the quality of video to provide more available bandwidth for the audio and 3D components. When the congestion clears, the video quality is raised back to its original level. At the end of the session, the system stores this type of session as a persistent session for later use by these users if this task is performed again.

It is towards this vision that our ACE framework is striving.

5. Getting There: Tasks, Needs, Services, & Technologies

In order to achieve this vision of an advanced collaborative environment, we have decided to focus on the following domains:

- The task domain – understanding the tasks that people perform when they are collaborating [18];
- The needs domain – understanding the users' needs as they perform the tasks;
- The services domain – understanding the collaboration services that are required to satisfy the needs;
- The technology domain – understanding the available technologies and how they can be used to effectively provide a service.

A simple example is useful for understanding this framework. Suppose that Bob is traveling on business and he has a task of collaborating with his boss on a project proposal. The task is to create the project proposal, and this involves sub-tasks of communicating ideas, preparing a draft document, correcting errors, and submitting the final version. In considering these tasks it becomes clear that there are a number of needs that must be satisfied. For example, Bob and his boss need to communicate rapidly and effectively. They also need some kind of document preparation capability, and they need to be able to check and correct their work. Finally they need to distribute the results of their work to the right people at the right time. There might also be other needs, such as those resulting from the boss's tight calendar, which might mean that they have to work very quickly in a real-time fashion.

Each of these needs can be satisfied by a number of services. For example, the need to communicate can be satisfied by a postal letter, an e-mail, a telephone call, or a videoconference session. By understanding all the needs, such as the need related to time constraints, some of these services can be rejected as unsuitable.

In turn, each of the services can be delivered using a variety of technologies. For example, a telephone call can be provided by a traditional telephone, a VoIP software client, or a cell phone. By first understanding the tasks and needs, better decisions can be made about the services and technologies to ensure that they provide an appropriate environment and a high QoE.

6. Tasks and Needs

Our work is based on the premise that the users' tasks and needs can be operationalized and then used as the basis for providing services via particular technologies.

To do this, we have created a task/needs matrix with tasks represented on the rows and needs represented on the columns, following the leading work done on the ETNA project [2] [15]. We have defined five basic collaboration tasks: meetings, collaborative work, education, presence, and entertainment. The tasks can be further subdivided as needed. For instance, meetings can be large or small, formal or informal. The columns of the matrix are a set of needs, including auditory communication, visual communication, audio/video synchronization, a shared workspace, a presentation space, decision support, turn-taking ability, privacy controls, and meta-communication.

Within the task/needs matrix, each cell is assigned an integer value to represent the importance of the particular need for that particular task. Where it exists, we have used research findings to assign the cells' value. Where no research exists, guesses as to the cells' value have been made. These will need to be confirmed via further research.

To illustrate the matrix, consider the row created for large informal meetings with colleagues. Each need for this task is described with an importance rating (in parentheses) for that need. Auditory (10) and visual communications (9) are extremely important for this task. Audio/video synchronization is moderately important (5). Shared textual and visual workspaces are moderately important (5). Presentation spaces (7) are more important than workspaces, but still less important than audio and visual communications. Turn-taking support is very important (8), more important than a presentation space. Decision support (7) is at the same level of importance as a workspace. Compared to the other needs, privacy controls (4) are not very important while a meta-communication channel is moderately important (6).

In contrast to large informal meetings with colleagues, consider as another example the row for remote monitoring presence applications. Both an audio (9) and visual channel (10) are extremely important in remote presence applications. Neither shared workspace nor a presentation space is important (1). Turn taking is not important (1), and neither is decision support (1). On the other hand, privacy controls (8) are moderately important, as are meta-communications (7).

In the current version of our user interface, the user defines the collaboration task by answering a series of questions about their goals. This identifies a row in the tasks/needs matrix. Once the task has been specified, the matrix row is used to determine which needs are important for task success. The set of needs that emerge from using the matrix are used to select and control the collaboration services and technologies that are used to accomplish the task.

7. Services and Technologies

Understanding the task, and therefore the needs of that task, is important for delivering a high quality collaboration experience. It is equally important to understand the collaboration services that can satisfy those needs and, in turn, understand how to deliver those services on a given set of technologies. To that end, we have defined a data model that helps us bridge the gaps between needs, services, and technologies.

The data model encapsulates information about both the tangible aspects of collaborations (nodes, users, etc.) and the more intangible aspects such as sessions and services. Furthermore, the relationships that exist between these aspects are also collected. The main entities of the data model are described in detail below.

7.1. Nodes

Nodes are a particular entrance point into a collaborative session. Nodes can be a single computer, a single computer representing a group of machines, or any device that has the capability of connecting to sessions and is available for collaboration.

Information acquired about nodes includes IP address, (current) physical location, and relevant technical information about the device - such as system bandwidth, memory capacity, processor type/capability, and screen size. Also captured is information about specific services available on the node.

7.2. Users

Users represent the actual people available to set up or participate in a collaborative session. User identification information such as name, e-mail address, and current physical location is stored. Other information useful for ensuring the best QoE for a planned session is also maintained, such as other physical locations to which the user has access, as well as other nodes that the user may have the right to use. Furthermore, the concept of groups of users is introduced and differing abilities and rights of users are maintained.

7.3. Sessions

A collaboration session is a set of services and other information that has been created based on the needs of the collaboration task that is being undertaken. The needs are met through the provisioning of an appropriate set of services that deliver a high quality collaboration experience. A session is an instantiation of one of the tasks from the QoE matrix described above. Users are invited to sessions and can join sessions by connecting from a node. Sessions are persistent in that they exist until deleted, but a session will only be accessible to those users who have been invited to attend that session. All of the information specific to a unique session is stored, such as the time and date of the session, the particular users or groups involved, and whether it is open to the public. Furthermore, information about specific services that are available for the collaboration are maintained, complete with any unique parameters required to facilitate the use of these services during the collaboration.

7.4. Service Classes

This is a generalized categorization and abstract representation of a particular aspect of a collaborative session. Examples of a service class would include a visual channel, an auditory channel, and a shared workspace. These service classes serve a dual purpose. The first purpose is to describe and rank the service classes associated with a particular task in the task/needs matrix. For example, a particular collaborative session type may place more importance on visual aspects of the collaboration over auditory aspects. These importance factors are maintained within this portion of the data model.

The second purpose is to help describe the capabilities of particular services. Each particular collaboration service is associated with one or more service classes that define deliverable aspects of that service. For example, a service such as H.323 is associated with the auditory channel, visual channel, and shared workspace service classes, whereas the RTCP/RTP/H.261 service is only associated with the visual channel service class.

This dual purpose facilitates the mapping of importance factors from the task/needs matrix to the actual abilities of services. To define the elements that make up a particular service class, they are further dissected into associated characteristics. These characteristics help to define certain defining aspects of the class itself. For example, the visual service class has characteristics such as latency, reliability and fidelity - each with a definable level of requirement for a specific session type - and a definable deliverable level for a specific service.

7.5. Collaboration Services

Collaboration services are those that provide a collaboration capability to the end user. They are either end-user services that provide a collaboration service to the user directly (such as a video or audio communication tool, a shared application, or a shared visualization) or ancillary collaboration services that provide a service to a collaboration session (such as an audio or video bridge service). Each collaboration service is classified as belonging in one or more high-level "service classes" such as audio, video, application sharing, chat, whiteboard, or visualization. As noted above, some services, such as H.323-based video conferencing, provide more than one class of service (audio, video, and application sharing).

We categorize our services as having at least (but not necessarily all of) a control protocol (H.323, RTCP, RTSP), a transport protocol (RTP), and a codec (H.261, MPEG2). This allows us to fully describe how each service communicates information and allows us to determine the compatibility of applications on client machines. To facilitate interoperability between applications, it is important to move from applications-based collaboration towards standards-based collaboration. Through the use of our categorization scheme, we are able to describe collaboration services based on standards and allow any client application that is compatible to be integrated into the environment. Collaboration services are not tied to applications until the service is deployed on the client node.

Although this may seem limited to media services like audio and video, in general this is not the case. Any collaborative application (ranging from a chat program to a 3D immersive collaborative application) is defined in the same way: it has a control protocol (how it connects, initiates, and controls communication), a transport protocol (how data is transported between collaborative applications), and a codec (how data is encoded to be transported). Many such applications do not differentiate these levels of communication, but they do exist.

Most of the information captured about services is kept at an abstract level. By maintaining only high-level descriptions, actual application-level decisions can be delayed and ultimately determined by the client when a decision can be made based on platform-specific issues and current operating conditions (e.g., current network load). However, there is some less abstract and more session-specific data that is useful and/or required when using a service during a collaborative session. This information is kept in the form of definable parameters associated with services. For example, when setting up a RTCP/RTP/H.261 service for a particular session, the IP and port information for that particular session must be maintained.

8. Current Status

8.1. QoE-ACE v0.1

Our QoE environment is currently in its second phase of development. The initial phase of development (v0.1) was completed in January 2003. It consists of a data model that contains a subset of the entities described above (services, sessions, tasks, and task needs). Commonly used tasks are predefined and stored in the database along with the task needs (an instantiation of the tasks/needs matrix).

Users create sessions by either choosing a predefined task to perform or creating an ad-hoc session and choosing the services themselves. When predefined tasks are chosen, the system determines the services required to accomplish that task on behalf of the user, and creates the appropriate session. The data model is implemented in a database with a web portal front end. Through the web portal, users can perform administrative tasks such as adding new entities to the QoE environment (e.g., create new services, tasks, and sessions). A typical user would only create new sessions or join sessions that have already been created.

When connecting to a collaboration session, the user navigates to the web portal and is presented with the currently active set of sessions. The user then chooses the session of interest. Once the session is chosen, the collaboration server sends a session description to the user machine. The session is encoded in XML and describes the relevant information for the session, including the session name and the set of services required. Each service is described by its class (audio, video, application sharing etc.), the protocol it uses for communication, and a set of parameters. The parameters contain connection information for the collaboration service as well as high-level service attribute descriptions (quality=high, latency=low).

When the client node receives the session description, the node refers to its platform specification to determine which services in the session it can support. Each platform has a specification file that describes how it maps collaboration services to specific applications on the local machine. The user is presented with a list of the services that are required by the session, with the services that are not available on the local machine listed as unavailable. The user starts the session, and the applications are automatically run with the correct parameters to manage connectivity to the other collaborators as well as the service quality. This framework for collaboration allows different applications to be used to provide a single service (possibly on a variety of platforms) as long as they use the same protocol. For example, a video service might be described as (class=video, protocol=RTCP/RTP/H.261). Any application that uses these protocols and standards (RTCP, RTP and H.261) could provide such a service, and the decision on which application to use is made by the local machine based on the service description.

The v0.1 implementation leveraged much of the AccessGrid 1.x [8] software tools and architecture. We view our system as a layer that sits on top of AG 1.x and other similar tools, providing a QoE capability to these environments.

8.2. QoE-ACE v0.2

Our experiences with v0.1 taught us a number of things about the requirements for a QoE-ACE environment. We found that the web interface was a good approach, as it is a ubiquitous interface and removed the need to operate the tools separately. It was also important that we did not overload the user with detail, and keeping the interface simple was required. We also confirmed our suspicions that to effectively deliver QoE we need to know much more information about the environment in which the user works.

We are currently in the implementation phase of the second revision of our QoE environment. Where v0.1 was focused on creating an infrastructure where some QoE-based collaboration could take place, v0.2 attempts to build upon that framework to create a richer environment capable of more advanced QoE decisions and recommendations. To this end, the data model was extended to include information about users, locations, and nodes. Furthermore, the concept of Service Classes was broadened to better map the relationship between the requirements of particular collaboration tasks and the abilities of the services available.

In v0.2, the user visits the web portal and logs in. This helps to identify to the system user-specific information and abilities that will aid in delivering QoE. From there, the user can choose to go through the session setup procedure as in v0.1; however, to fully utilize the tasks/needs matrix grid, the user can also choose to follow a "session creation wizard" which will aid in session setup and creation. This wizard is the basic user interface to the task/needs matrix. After answering a series of context-sensitive questions, the user's needs for the collaboration are captured. The user then indicates any specific users to invite to the collaboration session. Once this critical information is entered, the system can use this data to begin the decision making process to deliver the best QoE to the user based on the task at hand and the users involved.

As in v0.1, the system initially decides on the required services to accomplish the task described. Furthermore, abstract characteristics for these services (reliability=high, quality=medium) are determined and session-specific parameters (IP numbers and ports) are noted. In v0.2, the system then flags and notifies the session organizer of any possible problems in delivering the QoE for the desired task, such as a participant’s lack of access to specific services or requirements, and makes recommendations to overcome these inadequacies based on the availability of other nodes/locations and additional services such as bridges and transcoders.

When connecting to a collaboration session, the user navigates to the web portal, logs in, and is presented with the currently active set of sessions to which they have access. The user then chooses the session of interest. As in v0.1, once the session is chosen, the collaboration server sends an XML session description to the user machine, complete with all of the information necessary to join the session.

As was done in v0.1, the client node receives the session information and refers to its platform specification to determine supported services; however, in v0.2 the client node suggests service options based on more advanced information than just availability of applications. Information such as the current capability of the node, the current network connection, the availability of resources, and the use of the resources by specific applications are considered. The user is then presented with a list of optimal applications to obtain the highest QoE for the chosen collaborative session. The user then starts the session, with any user-specific adjustments, and all of the applications are launched with the correct parameters.

The current implementation of this framework includes the data model and a web portal front end. A prototype version of the QoE wizard exists, but the ability to make the QoE decisions that we desire based on users, nodes, and their capabilities is not yet implemented. The client software currently knows simple information about some of the physical attributes of the node on which it is running (multicast enabled, services available, etc.) but the client is not yet able to detect these attributes automatically. We plan to add both of these capabilities in the near future. As part of our v0.2 implementation, we have created a number of advanced services such as RTP-based video bridges that can transcode between video codecs and/or select a subset of the RTP video streams for transmission over a lower bandwidth connection (this service is similar in nature to the VVP software discussed in [22] [23]).

9. Future Work

The immediate plans for the QoE environment are to complete the development of v0.2. This will include completion of the QoE web portal front end, a more elaborate QoE decision process on session creation, and a more intelligent QoE-based instantiation of the session on the client node. We are currently exploring how the AccessGrid 2.x architecture can be leveraged to support our QoE framework. Like v0.1, our intent is to have our QoE framework exist as a layer that interoperates with AG 2.x and other similar collaboration environments.

Once we gain more experience with the QoE environment, our intention is to move the data model from a centralized server model to a distributed architecture where independent software agents will manage state information on behalf of their relevant physical artifacts. Components of the system that will require agents are users, nodes, services, and sessions. In this model, the centralized server will simply become a rendezvous point for the appropriate software agents to find each other and negotiate QoE solutions.

Longer-term work includes the development of advanced collaboration services such as adaptive audio and video services that adjust quality levels dynamically based on network performance while the collaboration is taking place (e.g., [6]). We are also exploring the integration of advanced interaction capabilities (e.g., table-based and wall display devices, touch sensitive input devices, and immersive visualization environments) into the set of services that can be supported by the QoE framework. This will provide us with a rich, dynamic, and heterogeneous collaboration environment in which to explore many of the issues in delivering QoE to the end user.

In addition, over the next year we plan on deploying our QoE platform for WestGrid, a large grid-computing project in Western Canada. The research program based around WestGrid has a significant collaboration and visualization research group. Our intention is to provide our QoE platform as a tool to deliver a high quality of collaboration experience to the WestGrid computational community. The QoE platform will integrate AccessGrid technologies (which have been or are being deployed at all WestGrid sites), other collaboration tools, and advanced collaborative visualization technologies into a seamless collaboration and visualization framework. Such a test bed will allow us to validate our research in a very challenging and diverse environment.

10. References

[1] Alben, L. "Quality of experience: Defining the criteria for effective interaction design." *Interactions*, 3(3), 1996, pp. 11-15.

[2] Anderson, A.H., Mullin, J., Jackson, M., Smallwood, L., Sasse, M.A., Wilson, G.M. & Watson, A. "Audio & video guidelines for networked multimedia applications: Applying the ETNA taxonomy." 2002. Available from the WWW: http://www-mice.cs.ucl.ac.uk/multimedia/projects/etna/avguidelines.pdf

[3] Bauer, B., & Patrick, A.S. "A Human Factors Extension to the Seven-Layer OSI Reference Model." Unpublished manuscript, April, 2003. http://www.andrewpatrick.ca/OSI/10layer.pdf

[4] Bertram, R. & Steinmetz, R. "Scalability of audio quality for networked multimedia environments." *Proceedings of ICMCS'97*, Ottawa, Canada, 1997, pp. 294-301.

[5] Bouch, A., Sasse, M.A., & DeMeer, H. "Of packets and people: A user-centered approach to quality of service." *Proceedings of IWQoS 2000*, June 5-8 2000, Pittsburgh, PA, pp. 189-197.

[6] Busse, I., Deffner, B., & Schulzrinne, H. "Dynamic QoS control of multimedia applications based on RTP." *Computer Communications*, 19, 1996, pp. 49-58.

[7] Campbell, J.A. "User acceptance of videoconferencing: Perceptions of task characteristics and media traits." *Proceedings of the 33rd Hawaii International Conference on System Sciences*, 2000.

[8] Childers, L., Disz, T., Olson, R., Papka, M.E., Stevens, R., and Udeshi, T. "Access Grid: Immersive Group-to-Group Collaborative Visualization," *Fourth International Immersive Projection Technology Workshop*, Ames, Iowa, June 19-20, 2000.

[9] Elstrom, L., Ousland, A., Poulain, G., Grein, K., Ranta-aho, M., Leppinen, A. & Roella, A. "Recommendations on user interface designing of video conference based CSCW systems." *Proceedings of the International Symposium on Human Factors in Telecommunications*, 12 - 16 May, 1997, Oslo, Norvege, pp. 255-257.

[10] Ghinea, G. and Thomas, J.P. "QoS impact on user perception and understanding of multimedia video clips." *Proceedings of ACM Multimedia'98*, Bristol, UK, 1998, pp. 49-54.

[11] Ghinea, G., Thomas, J.P. & Fish, R.S. "Multimedia, network protocols and users - Bridging the gap." *Proceedings of ACM Multimedia'99*, Orlando, Florida, 1999.

[12] Hollier, M.P. & Voelcker, R. "Objective performance assessment: Video quality as an influence on audio perception." Presented at the *AES Convention*, New York, September 26-29 1997. Audio Engineering Society Preprint 4590.

[13] Kies, J.K. "A psychophysical evaluation of frame rate in desktop video conferencing." *Proceedings of the Human Factors and Ergonomics Society 41st* Annual Meeting, vol. 1, 1997, pp. 310-314.

[14] Krawchuk, C. "Usability requirements for collaborative environments." Hewlett-Packard Labs, Bristol, Internet Research Institute. Technical Report HPL-IRI-2000-2, 2000.

[15] Mullin, J., Jackson, M., Anderson, A.H., Smallwood, L., Sasse, M.A., Watson, A. & Wilson, G. "The ETNA Taxonomy." 2002. Available from the WWW: http://www-mice.cs.ucl.ac.uk/multimedia/projects/etna/taxonomy.pdf

[16] Nardi, B., Schwarz, H., Kuchinsky, A, Leichner, R., Whittaker, S., & Sclabassi, R., “Turning away from talking heads: the use of video-as-data in neurosurgery.” *Proceedings of INTERCHI ’93*, pp. 327-334, New York: ACM Press.

[17] Patrick, E.J. "Barriers to collaboration: User-centered research and the Access Grid." *AccessGrid Technical Retreat Proceedings*, January 30-31 2001, Argonne National Laboratory.

[18] Ranta-aho, M., Poulain, G., Roella, A., Mirabelli, M., Ousland, A., & Norgaard, J. "Task-dependent user requirements for quality of service of videoconferencing-CSCW services." *Proceedings of the International Symposium on Human Factors in Telecommunications*, 12-16 May 1997, Oslo, Norvege, pp. 251-254.

[19] Rimell, A.N., Hollier, M.P. & Voelcker, R.M. "The influence of cross-modal interaction on audio-visual speech quality perception." Presented at the *AES Convention*, San Francisco, September 26-29 1998. Audio Engineering Society Preprint 4791.

[20] Sonnenwald, D.H., Bolliger, R., Solomon, P., Hara, N., & Cox, T. "Collaboration in the large: Using video conferencing to facilitate large group interaction." University of North Carolina at Chapel Hill, School of Information and Library Science, Technical Report TR-2001-02, 2001.

[21] Thakur, A., Gao, C., Larsson, A., & Parnes, P. "The effects of frame-rate and image quality on perceived video quality in videoconferencing." Lulea University of Technology, Department of Computer Science and Electrical Engineering, Technical Report LTU-TR-01/07-SE, 2001.

[22] Thorson, M., Leigh, J., Maajid, G., Park, K., Nayak, A., Salva, P., & Berry, S. "AccessGrid-to-Go: Providing AccessGrid access on Personal Digital Assistants", In *Proceedings of the Access Grid Retreat, 2002.*

[23] VanderBaan, K., Eschbach, J. "Handheld devices with AG nodes." Demonstration poster, *Access Grid Retreat,* March 2002.