Patrick, A.S., Singer, J., Corrie, B., Noël, S., El Khatib, K., Emond, B., Zimmerman, T., & Marsh, S. (2004). A QoE Sensitive Architecture for Advanced Collaborative Environments. Paper presented at the First International Conference on Quality of Service in Heterogeneous Wired/Wireless Networks (QSHINE 2004), Oct. 18-20, Dallas, TX.

A QoE Sensitive Architecture for Advanced Collaborative Environments

Andrew S. Patrick[1, 2], Janice Singer[1], Brian Corrie[3], Sylvie Noël[4], Khalil El Khatib[1], Bruno Emond[1], Todd Zimmerman[3], Stephen Marsh[1]

Abstract

A layered architecture for advanced collaborative environments has been developed to map the definition of a collaboration task from the requirements needed to accomplish the task, to the collaboration services that can be used to satisfy those needs, and to the technologies on which the services can be delivered. The architecture takes into account the users' goals and needs in configuring the services required for a collaboration task. By understanding and designing for users' needs and requirements, we can ensure a high Quality of Experience (QoE). The architecture was used to create a prototype system that is currently being tested through user studies.

1. Introduction

Users of collaborative environments care about seeing, hearing, and working effectively with their collaborators, not about codecs, protocols, and the other technologies involved in setting up a collaboration session. However, the technologies used and their performance over networks will often determine the final quality of the collaborative experience. If a piece of collaborative software does not perform well over the network then the resulting session will be awkward at best and unusable at worst. In this paper, we propose a collaboration architecture based on the concept of "Quality of Experience" (QoE) that addresses both low-level technical issues and high-level user needs. Specifically, our architecture exploits results from the CSCW literature to provide users with the necessary services and network performance for their collaboration session, while automatically making informed trade-off decisions when resources are limited.

Below we define QoE, which is the central conceptual model used in our architecture. This definition is followed by a description of our prototype system.

1.1 Quality of Experience (QOE)

Quality of Experience [2, 3, 9] 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. Many factors contribute to a user's QoE for a particular device or piece of software, including its appropriateness, effectiveness, learnability, and reliability [2]. QoE can mean different things for different applications. For example, QoE for an audio application is related to the sound fidelity and ability to smoothly take turns in a conversation, whereas QoE for a remote video application is related to frame rate and the clarity of the video image. Moreover, within the same media, QoE can be influenced by the collaboration task. For example, high video latency does not affect the QoE for a video-on-demand service, but it can have a large effect for interactive applications such as teleconferencing [4].

2. System Design

Our contribution lies in developing a concrete architecture and system design that is based on results from the CSCW literature. We start with the premise that the CSCW results can be used to define high-QoE collaboration sessions. That is, the CSCW results can form the basis for a translation from the task description to the technical parameters necessary to define a high quality collaboration session (see [8] for a similar approach).

To do this translation, we have defined a layered architecture that allows us to map between the definition of a collaboration task, the requirements needed to accomplish the task, the collaboration services that can be used to satisfy those requirements, and the technologies on which the services can be delivered. Our architecture not only considers the individual components required to successfully accomplish a collaboration task, but also takes into consideration how the interaction between different types of collaboration services and media affects the user experience during collaboration.

1. National Research Council of Canada, 1200 Montreal Rd., Ottawa, ON, Canada K1A 0R6
2. Contact E-mail: Andrew.Patrick@nrc-cnrc.gc.ca
3. Simon Fraser University, 8888 University Drive, Burnaby, B.C., Canada, V5A 1S6
4. Communications Research Centre, 3701 Carling Ave., Ottawa, ON, Canada K2H 8S2

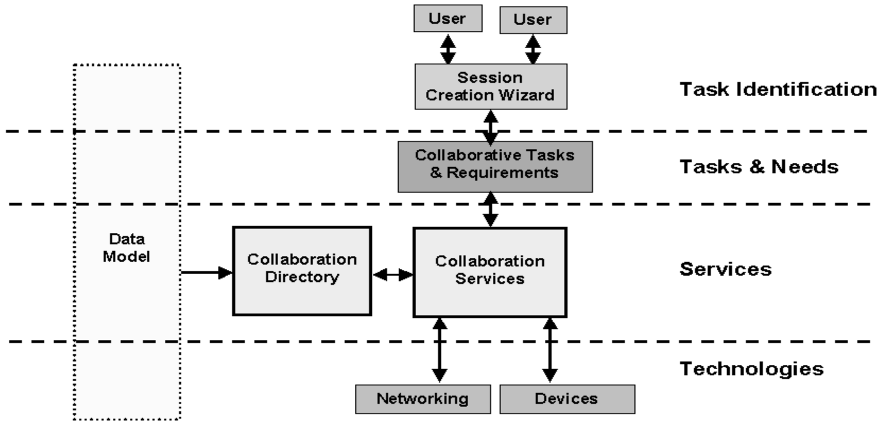

Figure 1. Advanced Collaboration Environment (ACE) Architecture.

3. The ACE Architecture

The ACE architecture is comprised of four distinct layers, as is shown in Figure 1 (see [5] for more details). The Task Identification Layer helps the user identify the type of collaboration task that is being undertaken. This results in a high-level task description. The Tasks/Needs Layer determines the collaboration requirements to accomplish the task as described. The Service Layer, in turn, uses these requirements to determine the services necessary for successful collaboration. Finally, the Technologies Layer instantiates the collaboration services using a set of appropriate parameters that can achieve the desired QoE.

3.1 Task Identification

Based on the work of Sasse et al. [6] we have developed a technique to guide users through the identification of a collaboration task. Our current implementation is a web interface that guides the user through a set of task-focused questions. The questions are context-specific and will vary based on answers to earlier questions in the task definition phase.

The questions are carefully chosen to be non-technical in nature. First, the user is asked to identify the type of collaboration session in which he/she will be engaged. Based on a review of the CSCW literature, we have initially defined five basic collaboration tasks: meetings, collaborative work, education, entertainment, and presence. Depending on the type of session, the user is then asked to specify more details about the session. For example, if the session is a meeting, the user needs to specify the number of parties that will be involved, the formality of the task, and the participants' familiarity with each other. If it is a collaborative work session, the user must specify the type of work that will be accomplished (e.g., generating ideas and plans, choosing, executing or performing, or negotiations). If it is for educational purposes, the user must specify the interaction style (lecture or open discussion) and the media. If it is an entertainment session, the user needs to specify the media that will be used. Finally, if the session is for presence, the user must specify whether it is for active monitoring purposes or if it will be in the background.

Once the type of task has been described, the user must then select the participants. In our system, a session can be public or private and groups or individual participants can then be added to the session from a list of available users.

3.2 Task/Needs Matrix

The task identification phase is used as the basis for defining the services and requirements for a collaboration session. That is, once a particular task has been identified, the system consults a task/needs matrix, which is an elaboration of a taxonomy developed by Patrick [7].

As discussed in the previous section, we have defined five basic collaboration tasks that are represented on the rows of the matrix: 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 an extensible 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.

Each need can have a set of characteristics that further describe its basic attributes. For example, visual communication is described by frame rate, resolution, and other video parameters. Auditory communication is described by fidelity, latency, and reliability. A shared workspace can be textual, visual, or both.

Within the task/needs matrix, cells are assigned integer values if that need is appropriate for that task (and left blank if the need is not appropriate for the task). Higher numbers indicate greater importance. Where they exist, we have used research findings to assign the cells' value.

Where no research exists, interpolation and experience are used to determine the cells' values. These values will need to be confirmed by further research.

It is important to note that these values are relative to each other across columns. That is, we are attempting to consider the impact of the collaboration needs on each other as well as the requirements to accomplish a task. In situations where there are limited resources, such as restricted bandwidth, we are able to make trade-off decisions based on the relative importance of the needs.

The questions presented to the user identify a row in the task/needs matrix. Once the task has been specified, the matrix row is used to determine which needs are important for task success. The needs that emerge from the matrix are used to select and control the collaboration services and technologies that are used to accomplish the task.

3.3 Services and Technologies

To tie the architecture together, we have defined a data model that helps us to 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 identified. The main entities of the data model are described in detail below.

3.3.1 Nodes. A node is a particular entrance point into a collaborative session. Nodes are any device that has the capability of connecting to sessions, such as a single computer or a group of machines. Information acquired about nodes includes IP address, physical location, and relevant technical information about the device such as network bandwidth, memory capacity, processor type/capability, and supported audio and video codecs. Also captured is information about specific collaboration services available on the node.

3.3.2 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 in ensuring the best QoE for a planned session is also maintained, such as links to the 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 is maintained.

3.3.3 Service Classes and Services. Service classes are abstract representations of a particular aspect of a collaborative session. The service class captures the required quality levels for a specific task, as defined by the task/needs matrix. Services describe specific resources that provide a collaboration capability to the end user. Services either support collaboration directly (such as a video or audio communication tool, a shared application, or a shared visualization tool) or they provide an enabling service to a collaboration session (such as an audio or video bridge service). Each service belongs to one or more service classes, which helps to define its capabilities. For example, video communication is a service class and specific video applications are services within that class.

3.3.4 Sessions. A collaboration session is a set of services and other information that has been assembled based on the needs and requirements of a collaboration task. Thus, a session is an instantiation of one of the tasks from the task/needs matrix described above. All of the information gathered during the task identification stage is associated with the session.

4. The Implementation

Our QoE sensitive collaboration environment is currently being developed and tested. The heart of the system is the data model that contains the information about tasks, needs, sessions, services, users, and nodes. Commonly used tasks are predefined and stored in a database along with the requirements that must be satisfied to accomplish those tasks effectively (an instantiation of the task/needs matrix).

4.1 Creating Sessions

To use our collaboration system the user visits a web portal and logs in. This identification step enables the definition of user-specific information and abilities that will aid in delivering an appropriate QoE. From there, the user can follow a "session creation wizard" that aids in session setup and creation. This wizard is the user's interface to the task/needs matrix. By answering a series of context-sensitive questions, the user's task is defined and the needs for the collaboration are captured. The user then indicates any specific colleagues to invite to the collaboration session.

Once this critical information is entered, the system can use this data to recommend appropriate services, and adjust the suitable quality settings based on the task at hand and the users involved. In general, these service recommendations are made by comparing the service requirements in the task/needs matrix to the capabilities of the available services.

4.2 Joining Sessions

When users want to join a collaboration session, they navigate to the web portal, log in, and are presented with the currently active set of sessions to which they have access. The user then chooses the session of interest. Once the session is chosen, the collaboration server sends an XML session description to the user's machine, complete with all of the information necessary to join the session.

This session description includes the session name and the set of services to be used as part of the session. 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 characteristic descriptions (e.g., for the video service the characteristics would be latency, fidelity, and reliability).

Client software on the user's node receives the session information and determines the session requirements. Information such as the applications available on the node, the current network connection, the availability of resources, and the use of the resources by specific applications are considered. The system then starts a set of applications that provide the required QoE for the chosen collaborative session on the available technology.

Our 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, service = 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 node based on the service description. During development and testing we use the application tools that are part of the Access Grid project [1]. For example, VIC is used as the video application, while RAT is used for audio communications. Our goal is to create an interface and decision-making layer that sits on top of the Access Grid applications and other similar tools, providing a QoE capability to these environments.

5. Discussion

We have defined a system architecture that draws on results from the CSCW literature to make informed decisions about the services and quality settings that are necessary to collaborate successfully given different task descriptions. We call this a QoE sensitive architecture because our goal is to deliver the highest possible QoE to end-users based on the collaboration task and their system characteristics.

There are some shortcomings to our work. Chief among these is the reliability of the CSCW results in defining the needs for different collaboration contexts. As we use our system more, it will be necessary to refine the task/needs matrix to reflect the end users' perceived QoE in different collaboration contexts. Furthermore, as technologies develop, the mapping between the needs and the technologies will have to be updated continuously. As the CSCW literature matures, we will learn more about tasks and needs, and we will refine and expand the matrix. Since we have defined a layered architecture, however, this should be possible without revising the basic architecture.

We will begin collecting data on the use of the system shortly. We are currently carrying out a review of the web portal front-end in order to identify its usability for non-technical users. In particular, we are exploring whether or not the series of questions that we ask the user are appropriate for identifying the user's desired collaboration task. In addition, we will be performing studies of participants in distributed collaboration sessions to measure their QoE during the sessions. We will also collect network statistics during these sessions. These measures will allow us to explore the mapping between network characteristics and user satisfaction.

6. References

1. Access Grid Project. http://www.accessgrid.org/.
2. Alben, L., (1996). Defining the criteria for effective interaction design: Quality of experience. *Interactions*, *3*, 3, 11-15.
3. Bauer, B., & Patrick, A.S. (2004). *A Human Factors Extension to the Seven-Layer OSI Reference Model*. Retrieved January 6, 2004, from http://www.andrewpatrick.ca/OSI/10layer.html
4. Bhatti, N., Bouch, A., & Kuchinsky, A.J. (2000). Integrating user-perceived quality into Web server design. Proceedings of WWW'00, May 15-19th, Amsterdam, The Netherlands, pp 1-16.
5. Corrie, B., Wong, H., Zimmerman, T., Marsh, S., Patrick, A.S., Singer, J., Emond, B., & Noel, S. (2003). Towards quality of experience in advanced collaborative environments. Paper presented at the Third Annual Workshop on Advanced Collaborative Environments, June 22, Seattle. (NRC 46521)
6. Mullin, J., Jackson, M., Anderson, A.H., Smallwood, L., Sasse, M.A., Watson, A. & Wilson, G. (2002). The ETNA Taxonomy. Available: http://www-mice.cs.ucl.ac.uk/multimedia/projects/etna/taxonomy.pdf
7. Patrick, A.S. (1999). The human factors of MBone videoconferences: Recommendations for improving sessions and software. *Journal of Computer-Mediated Communication, 4*, 3 http://www.ascusc.org/jcmc/vol4/issue3/patrick.html
8. Richards, A., Rogers, G., Witana, V., & Antoniades, M. (1998). Mapping user level QoS from a single parameter. *Second IFIP/IEEE International Conference on Management of Multimedia Networks and Services*.
9. van Moorsel, A. (2001). Metrics for the internet age: Quality of experience and quality of business. Technical report no: HPL-2001-179, http://www.hpl.hp.com/techreports