On the Development of a Model for Distributed Collaboration

Sami Rollins and Kevin C. Almeroth
Department of Computer Science
University of California
Santa Barbara, CA 93106-5110
{srollins,almeroth}@cs.ucsb.edu

In the wake of networking advances such as higher network bandwidth and greater end user computing power, collaboration between parties at geographically distributed sites is becoming a reality. User can communicate using not only text-based messaging, but a variety of media including audio and video. While many of the primitives exist to generate and distribute content, as the application space grows, the number of media types grows as well. Applications must support dissemination, storage, and retrieval of a growing, heterogeneous information base.

Many current solutions have focused on using straightforward multicast to distribute information among users. While multicast is an extremely efficient solution, many collaboration applications present scalability bottlenecks at the application layer. For example, in a distance learning application video streams of the instructor, audience, and remote participants may be produced. Simple multicast of all streams to all participants is not a sufficient solution. The challenge is present the user with the most useful subset of the entire information base. In this work, we focus on developing a model for collaborative applications that focuses on semantically enhancing media streams. We describe a four phased data dissemination architecture that looks at content generation, dissemination, storage, and manipulation. To illustrate the concepts we describe, we look at the use of the model in developing a distance learning application.

Phase 0 - The Basic Infrastructure: Any group collaboration system must have a basic hardware infrastructure that is capable of producing content. The infrastructure support defines the input and output sources of the application. In the case of distance learning, the infrastructure includes video feeds of instructors and students, content from electronic whiteboards, video output from laptop computer presentations, and annotation tools such as digital pens.

Phase 1 - Raw Information Dissemination: Data collected from the basic infrastructure can be disseminated, in its raw form, to the users of the application. This means that streams coming directly from a device such as a video camera may be sent to participants of the collaboration. To achieve this level of information dissemination, we must consider media encoding types, network bandwidth, and end user restrictions. Most distance learning applications perform this type of raw information dissemination. Video of the instructor on classroom audience is streamed over the Internet to individual students in remote locations.

Phase 2 - Semantic Data Representation: As the number of input sources becomes larger and more diverse, the content produced becomes too overwhelming to be useful in its raw form. For example, a user cannot simultaneously view multiple video streams. Therefore, users need a way to filter information and selectively choose which information is most relevant at a given time. We can accomplish this task by adding semantic, machine understandable information to each media stream. For example, if a professor indicates that she is making an important point in the lecture, that information is added to the video stream such that students watching from home will see video of the professor at that point rather than video of other students. The challenge is to look at the input sources of a given application and determine what kind of information can be added to the media streams, and how that information should be added. While some research has focused on post-processing of media streams to add such information, automatic, real-time addition of semantic information has not been fully realized.

Phase 3 - Data Services: Once the raw data have been decomposed into semantically enhanced media streams, the final goal is to define and develop a set of services that allow users to filter, aggregate, and manipulate the real-time as well as stored media streams. One such service might allow a user to browse a lecture video by searching for keywords. A more advanced service might moderate user interaction by defining profiles for each user. When users send questions or comments about the given topic, their content is broken down, semantically enhanced, and redistributed to users who have expressed an interest in the topic via their profile.