

Schemes for User-interest Controlled Video Bandwidth Adaptation in a Collaborative Workspace Environment

Stefan Elf¹, Jeremiah Scholl, Peter Parnes

*Division of Media Technology, Department of Computer Science and Electrical Engineering
Luleå University of Technology, SE-971 87 Luleå, Sweden
{stefan.elf, jeremiah.scholl, peter.parnes}@ltu.se*

Abstract

In this paper, bandwidth-sharing schemes controlled by user-interest are presented as an approach to efficient use of resources in a collaborative workspace environment. Two schemes for group video-conferencing are presented and evaluated. Monitoring of user-behavior and message passing to adapt video bandwidth allocation to user needs are key features of both schemes, in which each user's client reports interest in other users to enable them to determine their relative importance. Experimental results using a prototype implementation to sample user-behavior show that one scheme is better suited for high, and the other for low bandwidth-limit scenarios. Measurements also show that the message passing will not add a substantial amount of bandwidth.

1. Introduction

Body language is an integral part of how humans communicate. When people interact, according to some authors [1] as little as 30 % of a message is contained in the spoken words. Thus, voice-only communication cannot entirely replace real-life meetings. The *collaborative workspace* class of applications adds other services to voice communication, such as video, in order to create a *virtual presence*. As video streams are a major component with regards to bandwidth consumption and since bandwidth sometimes is a limited resource, there is a need for applications to use video bandwidth efficiently.

Adapting video to other users' interest is one proposed strategy for efficient use of video bandwidth [2,3]. Schemes of this type are based on a general assumption that a user's focus on an individual video

stream will vary over time. Thus, resources will be used more effectively if a user is provided with higher quality video from those participants that she is currently focused on as opposed to those that she is not. In addition, these schemes also target application usage scenarios where the number of "important" video streams (those that are the object of high user focus), vary over time as opposed to floor control schemes, which grant resources to the single most important user at any given time.

This paper discusses two alternative schemes for user interest controlled bandwidth sharing where session participants who are of high importance are allocated a larger portion of the session bandwidth. An initial empirical study was conducted into how user focus changes over time and simulations of each scheme were conducted based on the data collected. These simulations are used as a basis for comparison and contrast.

1.1. Related work

Resource control by user-interest has been applied to a range of multimedia applications, with each scheme being limited in scope to a specific domain. For example, Kulju et al. [4] investigated user behavior in the context of video streaming, while Ott et al. [5] focused on its use within their own 3D landscape. The SCUBA protocol [2] represents early work in this area, describes the architectural components for the class of schemes discussed in this paper.

Chen [6] designed a multi-party video conferencing system in which low-frame-rate video was sent during idle periods. The frame rate was increased as soon as a user made a gesture that signaled a predefined relevant activity. While Chen focused on low-bandwidth sessions and explicit user interest, our work is not restricted to low bandwidth usage and handles implicit

¹ Stefan Elf is also with Ericsson TEMS AB, SE-931 87 Skellefteå, Sweden. Views expressed in this paper are his own and not necessarily shared by his employer.

user interest.

2. Strategies for bandwidth adaptation

Since users' interaction with the application is a primary component in detection of user interest it is practical to develop schemes related to user interfaces available in a specific application. The work in this paper is based on the Marratech Pro Work Environment [7], a product of earlier research at our university and commercially available for several years, which has been used for prototyping and measurements. The environment comprises several video windows as shown in Figure 1. Each window represents a viewing context and they are (from left to right) private chat window, focus window, and participant window.



Figure 1. Video windows included in the Marratech work environment

The windows that members of the work environment use to interact with each other can be exploited for classification of sender importance. Examples of *user classes* are described in Table 1, and include the use of cross-media events, e.g. where existence of audio transmission is used to allow for increased video bandwidth. A higher class implies a more important sender.

Table 1. User class definitions

Class	User interaction
1	Present in participant window
2	Engaged in private chat
3	Present in at least one other focus window
4	Sending audio

An improved scheme for bandwidth sharing should exhibit a dynamic behavior related to the use of bandwidth in various scenarios. The schemes presented in the following subsections represent two different approaches. The common property is that a small amount

of bandwidth is first reserved for all participants. However, in the first scheme attempts are made to deliver the "minimum" requirements [3] for low priority senders before allocating increased bandwidth to increasingly important senders, as it is assumed that the bandwidth required for low priority senders will not take away a significant amount of bandwidth from high priority senders. The second algorithm takes the opposite approach in that bandwidth is allocated to the most important users first, followed by less and less important users, in order to make sure that important senders can deliver high quality video in more bandwidth constrained sessions.

2.1. Minimum-requirements-first

As stated above the minimum-requirements-first approach first allocates a small amount of bandwidth to users of all classes. Additional bandwidth is then allocated according to increasing importance until all available resources have been allocated.

Step 1. Bandwidth is divided evenly between all senders of class 1 – 4, until each sender can send at some minimum bandwidth suitable for the Participant window.

Step 2. If there is still bandwidth available after step 1, it is allocated between the senders of class 2 or higher until they are sending at some bandwidth suitable for the private chat window. Otherwise no more bandwidth is allocated.

Step 3. If there is still available bandwidth after step 2, it is divided between senders of class 3 or higher until they can send at some bandwidth suitable for the Focus window. This is done first for class 4 senders, and then for class 3 senders. If, in any of these two sub-steps, bandwidth is exhausted, the allocation process terminates.

Step 4. All possibly remaining bandwidth is divided evenly between each sender in class 3 and 4.

2.2. Important-senders-first

The important-senders-first approach also begins by allocating a small amount of bandwidth to each member but continues by attempting to meet the needs of senders of high importance before allocating additional bandwidth to lower priority users. The definition of "needs" is not given here, but is touched in the reference literature [3,6] and remains an object of further research.

Step 1. A small amount of bandwidth that is suitable for the participant window is allocated to each member.

Step 2. If additional bandwidth is available, senders who are also sending audio, class 4 senders, are allocated an additional amount of bandwidth suitable for the focus window plus some extra amount due to their significance as speakers. Otherwise the process stops here.

Due to the nature of speech, there is most likely only one class 4 sender in a session irrespective of the scenario.

Step 3. If there is still bandwidth available, a suitable amount of bandwidth for the focus window is allocated to the class of senders that are shown in at least one other user's focus window, class 3 senders. Otherwise the process stops here.

According to the results from our study presented in section 3, there may be several class 3 senders in a session, and the bandwidth allocated for this user class is distributed according to the number of other users that are viewing each important sender.

Step 4. If there is still bandwidth available, a suitable amount of bandwidth for the private chat window is allocated to class 2 users. Otherwise the process stops here.




Figure 2. Number of unique class 3 senders

3. User study and Simulations

A prototype with the purpose of logging user behavior relating to class 1, 2, and 3 events was implemented. The event logs were parsed and analyzed by a log analyzer, which also implemented the bandwidth-sharing schemes, allowing for the schemes presented in subsections 2.1 and 2.2 to be simulated and analyzed.

Events were recorded in a group of 9 experienced users during one day, corresponding to a discussion scenario for 7 hours and a lecturing scenario during the

last hour. The prototype recorded 1109 events, where e.g. "is-looking-at" or "video-window-selection" events signified that the object of the event belonged to class 3, while "opened-private-media" events implied a sender of class 2. The events were used to classify (Table 1) senders to determine the bandwidth to which they were entitled.

The total amount of events that were transmitted during this 7-hour period corresponded to 161221 bytes, i.e. a traffic volume of 6.4 bytes/s, which is an insignificant amount of messaging in a completely non-optimized messaging system.




Figure 3. Actual average bandwidth used in relation to the session limit

From Figure 2, which shows the number of unique video senders that were regarded as important by other users, it is evident that the number of such class 3 senders can be significant. This is illustrated by the fact that, during some periods, almost half the participants were actually class 3 senders. This motivates a method for distributing bandwidth unevenly between class 3 senders, and may require a more fine-grained approach.

The performance of the two proposed bandwidth-sharing schemes was simulated using different session bandwidth limits based on the events logged. As shown in Figure 3 there were significant differences between the two schemes. In particular, although the two schemes delivered a comparative amount of bandwidth to the most important senders, the minimum-requirements-first scheme used a larger percentage of the total available bandwidth. This shows that the important-senders-first algorithm may be more appropriate in situations with a lower session bandwidth limit is available and it is deemed more critical to provide the most important senders with a larger

share of the session bandwidth.

4. Conclusions and future work

Two different schemes for bandwidth sharing are presented in this paper, the minimum-requirements-first (MRF) scheme and the important-senders-first (ISF) scheme. The MRF scheme delivers minimal bandwidth for each increasing level of presence before dividing all remaining bandwidth among the most important senders. The ISF scheme takes the opposite approach in that important senders are given the bandwidth required for a high level of presence before the remaining bandwidth is divided among lower priority senders.

We conducted a pilot user study, which logged user interactions that would result in changes in sender importance and used the data collected in order to simulate each algorithm. The simulations were used to show the behavior of each algorithm when operating under increasing session bandwidth limits.

The difference in bandwidth requirements in a conferencing system when using the ISF scheme in relation to the MRF scheme depends on the group size, the spread in group focus and the session bandwidth limit. A typical group at our university has about 15 users of which some 10 users actually send video. With 10 video senders and no bandwidth restrictions the consumed bandwidth can easily reach more than 1 Mb/s with an average of three important users. However, in our simulations, ISF was able to deliver high quality video to each of the important users while consuming less bandwidth on average (Figure 3) than the MRF scheme. Of course, the bandwidth actually used in a specific case, depends on the "suitable" bandwidths that are assigned to each user class, and of requests made by participants in the session.

Thus, although both algorithms make it easier to serve important senders ISF seems to be more appropriate for particularly bandwidth constrained sessions.

4.1. Future work

Future work includes handling of multi-session senders, where user classification may or may not be relevant across sessions. Our current experimental platform does not support sending different-rate video streams to separate sessions and will use the lowest rate to all sessions. Hence, users will be over-provisioned in some sessions and under-provisioned in other, in relation to their actual importance.

When the most important senders change, so should the bandwidth allocation. Our experiences from physical meetings and discussions show that it is more com-

mon that the floor is yielded back to a previous speaker than to a new person. The bandwidth allocated to a sender should therefore not decline too fast.

We will also examine an entirely different scheme involving a procedure with receiver votes, and sender importance that declines with time unless the receivers cast new votes. This will increase messaging due to the voting frequency but there would be no need for sending downgrading events. As audio would be an important criterion for determining the most important sender, voting frequency may be tuned after the characteristics of human speech, e.g. pauses between words and sentences.

In order to evaluate different sharing schemes, there will be a need for implementing a full prototype, and evaluation will include subjective user surveys in addition to objective measurements. One explicit goal for such measurements and surveys is to establish what bandwidth that is "suitable" in various scenarios and for the different user classes.

References

- [1] J. Wang and D.G. Frank, Cross-cultural communication: Implications for effective information services in academic libraries *portal: Libraries and the Academy*, vol. 2, no. 2, pp. 207--216 Apr, 2002. (Abstract)
- [2] E. Amir, S. McCanne, and R. Katz, "Receiver-driven bandwidth adaptation for light-weight sessions," *Proceedings of the 1997 5th ACM International Multimedia Conference, Nov 9-13 1997*, pp. 415-426, 1997.
- [3] J. Scholl, S. Elf, and P. Parnes, "User-interest Driven Video Adaptation for Collaborative Workspace Applications," *Fifth International Workshop on Networked Group Communications (NGC'03)*, 2003.
- [4] W. Kulju and H. Lutfiyya, "Design and implementation of an application layer protocol for reducing UDP traffic based on user hints and policies," *MMNS 2002 - IFIP/IEEE International Conference on Management of Multimedia Networks and Services, 6-9 Oct. 2002*, pp. 263-75, 2002.
- [5] M. Ott, G. Michelitsch, D. Reininger, and G. Weiling, An architecture for adaptive QoS and its application to multimedia systems design *Computer Communications*, vol. 21, no. 4, pp. 334-49, Apr 10, 1998. 0140-3664. C&C Res. Lab., NEC USA Inc., Princeton, NJ, USA.
- [6] M. Chen, "Achieving effective floor control with a low-bandwidth gesture-sensitive videoconferencing system," *10th International Conference of Multimedia, Dec 1-6 2002*, pp. 476-483, 2002.
- [7] Marratech. - The e-meeting company, <http://www.marratech.com>, URL visited Dec 4, 2003.