[BeSt96] R. Bertram., R. Steinmetz: Varying Audio Quality for Networked Multimedia Services and Tools; The 1996 European SMPTE Conference on Imaging Media, September 1996, pp. 131 – 142

Varying Audio Quality for Networked Multimedia Services and Tools*

Reinhard Bertram, Ralf Steinmetz
Industrial Process and System Communications
Department of Electrical Engineering and Information Technology
Technical University of Darmstadt
Merckstr. 25, D-64283 Darmstadt, Germany
{Reinhard.Bertram,Ralf.Steinmetz}@KOM.th-darmstadt.de

Abstract

Networked multimedia systems and applications must cope with varying availability of resources such as buffer capacity, CPU time and bandwidth. Resource reservation schemes allow the reservation of resources at the moment of connection setup. The allocated bandwidth may not be sufficient to result in the desired Quality of Service for the user of multimedia applications. If, for instance, more bandwidth is available at a later time or if there is no resource reservation at all, the application may control the bitrate of the associated media streams depending on the available bandwidth.

In this paper we show some impacts of the scaling of audio data on the *user* of multimedia applications. It turns out that six levels of quality are sufficient for an implementation which allows the transition from high quality audio (CD) to very low quality audio (GSM-like speech) and vice versa. Scaling towards lower quality should be performed in various steps, while scaling toward higher quality may be done in one or many steps in most cases. For the user of multimedia applications the best scaling algorithm is a cross-fade between quality levels.

1. Introduction

Information processing in a time-sharing environment is performed without any hard time constraints. The system responds to a user interaction as soon as possible and most often lacks support for real-time data. The use of discrete media (e.g. text and graphics) dominates traditional computing, while the integration of continuous media (e.g. audio and video) into existing computer environments creates the new complexity of time-dependent data processing. 'Correctness' in real-time systems is determined by whether deadlines are met. We define the processing of time-dependent data in multimedia systems as the delivery of data in well defined intervals over a period of time, in our terms this process is called a continuous-media stream. Multimedia communication deals with the transfer of discrete *and* continuous media over digital networks.

In networked multimedia systems various entities typically cooperate in order to provide the required real-time guarantees of data that is to be presented at the user interface. Therefore resource management systems that provide mechanisms for streams with guaranteed or statistical Quality of Service (QoS) have become a key issue (for details see, e.g., [Wol96][NaS95]). Those systems take care of the coordination of media streams and the interfacing between layers of protocol stacks as well as further mechanisms (like process and bandwidth scheduling) in order to enforce the appropriate data handling. Most of the involved mechanisms are developed for a complete error-free presentation of continuous-media data at the user interface. In todays networked environments we still encounter many data paths over networks and via communica-

* This work is sponsored in part by: Volkswagen-Stiftung, D-30519 Hannover, Germany.

tion protocols which are not capable of providing a guaranteed real-time service. In such setups it is a key issue to decide which data item must be presented at the user interface and which data items may be discarded.

QoS is the set of parameters which defines the properties of media streams. In accordance with [StN95] we distinguish between four layers of QoS: User QoS, Application QoS, System QoS and Network QoS. The user QoS parameters describe requirements for the perception of multimedia data at the user interface. The application QoS parameters describe requirements for the application services possibly specified in terms of media quality (like end-to-end delay) and media relations (like inter/intra stream synchronization). The system QoS parameters describe requirements on the communication services resulting from the application QoS. These may be specified in terms of both quantitative (like bits per second or task processing time or ordered delivery of data) and qualitative (like interstream synchronization, error recovery or ordered delivery of data) criteria. The network QoS parameters describe requirements on network services (like network load or network performance).

Multimedia applications negotiate a desired QoS during the connection setup phase either with the system layer or possibly directly with the network layer, if the system is not able to provide QoS for the application. If both of them are not capable of providing the desired QoS, many of todays multimedia applications try to set up an end-to-end connection and to take care of QoS by themselves. Alternatively data is transferred with the best effort approach (see e.g. the various Internet schemes). This may happen even if the network is able to reserve a specific amount of bandwidth (but less than demanded) for multimedia applications. The Resource ReServation Protocol RSVP [ZDE+93] reserves in every intermediate system component from end-system to end-system the respective resources. The QoS that an application can provide to the user reflects the status of the available resources at the time the connection was set up. The negotiated QoS may be "lower" than the QoS the network may provide at a specific moment. This happens for example when a connection is established during peak traffic load hours and lasts until the network is not heavily loaded anymore. In such a situation a dynamic adaptation of the data rate would be useful, depending on the application. Sometimes data must be transferred over "non reserved" paths which means that loss of data is likely occur. In such a situation it would be helpful (if not required) to discard less significant data and to adapt the media stream to the available resources (scaling down).

2. Related work

Work on the subject of scaling and filtering has been carried on at various locations, nearly all of this research work focuses on video data [DHH+93],[GoV93],[WHD95]. Using software codecs, multimedia applications on workstations or PCs are often capable of adjusting the media quality and the data rates over a wide range. This capability allows for dynamic scaling mechanisms, which means to retrieve the state of the network continuously and adjust the application QoS accordingly [BDS96] [BTW94]. Scaling is most often performed in the following way: Whenever the network becomes congested, the data rate of the media stream is decreased by a high degree. It is increased softly until the value desired by the application is reached or until congestion occurs again. This well-known approach has been used successfully in packet-switched networks to handle data traffic [Jac88]. In opposite to the above references we study adaptation mechanisms at the user interface.

3. Motivation

Any perception of loss of audio quality depends on the type of audio encoding used and which packets are dropped from the specific media stream. We believe that the user does not perceive all errors as equally annoying and he or she may not even notice some of them. In this paper such issues related to audio are evaluated in more detail. Therefore, here we want to find out what a user perceives as an error-free data presentation or – if errors are noticeable – what errors may be tolerated.

Our work outlines impacts on the scaling of audio data, which needs to be kept under limited constraints, because "the ear is surprisingly sensitive to sound variations lasting only a few milliseconds. The eye, in contrast, does not notice changes in light level that last only a few milliseconds" [Tan96]. The real-time nature of audio (and video) data streams and the many incompatible encoding schemes used for audio to achieve best quality at a minimum bitrate for specific kinds of audio data make it hard to find suitable scaling mechanisms. The video frame rate can be varied by means of many individual parameters like the frame rate, the video spatial resolution, and e.g. the DCT parameter quantization. Audio encoding schemes most often result in constant bitrate and the algorithms are optimized for some degree of quality and some range of data rate. Hence for a change of data rate in the order of magnitude 1:100 (this is the approximately range from GSM to CD audio) the encoding must be changed as well.

Human perception of aural data varies from person to person. It is based on the different notion of good sound quality, musical skills, age, hearing defects and many other issues. Hence any experiment to define user QoS must address various different people. Although the individuals which took part in our series of experiments may not be representative for all users of multimedia applications, we have achieved similarities in the results and present them in this paper. All results are based on a set of preliminary experiments.

The remainder of this text is organized in 5 sections: Section 4 outlines human perception of audio data and the influence of the application environment. In Section 5 we provide our assumptions and prerequisites for all tests. In Section 6 the experiments are presented, the respective results are shown in Section 7. Section 8 outlines the implications of our results on audio encoding schemes and scaling mechanisms.

4. Human Perception & Application Environment

Any measurement of an aural impression must take into account the respective environment as well as the experience and expectation of the listening people – the test candidates; We can't expect to find *the typical user* of multimedia applications, hence the tests need to be performed with an average sample of people of both sexes, various ages, different musical skills, etc. In this context it is also difficult to find out the personal awareness for audio, so we restricted ourself to let the people define their capabilities according to our questionnaire (without checking this in detail).

As a first step we defined the target application domain to be networked multimedia applications running at workstations, PCs and mobile equipment. Table 1 outlines individual environments, the related devices and applications (the related details are shaded in Table 1). This list is certainly far from complete – we use it to discuss with the candidates where such an audio effect would appear.

In summary, we are looking for constraints for the scaling of audio data in a workstation/ PC/MobilePC environment used by users with no specific musical skills. The adaptation will result in changes of data rate, which may not necessarily result in a change of the perceived audio quality of the presentation. If we can find a way to minimize the subjective change in the perception of audio data (while still allowing to change the data rate for the related audio data streams), we will be able to design the respective protocols and architectures for QoS processing. Dynamic bandwidth allocation may be integrated into the operating and communication system and it may be controlled by the application. In contrast to related work we focus on the user QoS of multimedia applications.

Devices	Application Examples	Environment	Noise Level
professional audio equipment	production, editing	recording studio	none
PC/workstation with professional audio equipment	multimedia editing, authoring	multimedia studio	none
high end consumer equipment	playback	home	none to low
average consumer equipment	playback	home	low to medium
low end consumer (portable) equipment	playback, background	home, outside	medium to high
dedicated multimedia equipment attached to dedicated computer	virtual reality, multimedia kiosks, video conference	office, lobby, public place	medium to high
multimedia equipment attached to general purpose computer	video conference, CSCW, information retrieval	office, PC lab, private PC	medium to high
multimedia equipment attached to mobile computer device	presentation	office, hotel, outside	medium to high
ISDN Telephone	telephony	office, private	medium
analog telephone	telephony	private, office	medium
GSM telephone	telephony	outside, urban, office, car	high

Table 1: audio devices in their typical environment

5. Assumptions/Preconditions for UIF Test

At the beginning of this study it was difficult to choose which parameters to be adapted and changed and which to be kept constant. Therefore we ran a set of preliminary tests where many parameters were varied. These preliminary experiments showed:

- Scaled music annoys the user more than scaled speech:

  People are accustomed to listen to bad quality speech (like when using mobile phones), yet they are not accustomed to listen to bad quality audio. Hence music is more demanding than speech, so we decided to choose some Jazz/Pop music as a sample for our experiment. Later on we will continue with experiments using speech.
- The choice of the amount and the respective levels of quality:

  Six levels are sufficient in the range from CD to GSM-like quality for the purpose of our tests. We have chosen the lowest level of quality to be GSM-like, because it seems challenging to integrate mobile phones into networked multimedia systems. Therefore audio data originating from mobile telephone equipment is represented by this low quality in the applications.
- How fast should subsequent changes in quality occur?

  It is better to keep the sound quality constant for a minimum time of ten seconds to allow adaptation of the ear to the respective new situation. Shorter times seem to confuse the listener, longer periods unnecessarily lengthen the total test time.
- Test environment:

  The tests should be performed in a low noise environment to keep the tested individual concentrated during the approximately 45minutes required for all experiments.
- Personal awareness of audio quality:

  It is important to distinguish among tested individuals with a musical background, ones working in the field of professional audio and ones having no special training at all. We selected people with different background and asked them to define their aural capabilities.
- Quality of used playback devices:

  One of the major influences on the statements given by the candidates appears to be the device used for the test. All statements need to be weighted by the quality of the audio playback device.

6. Experiments

The tests were performed in three sessions. Each session consisted of the audio files we created in a professional recording studio* using the harddisk recording system ProTools and by applying various digital filters to different sections of a chosen sample CD track. In the beginning of the experiment we introduced the session by explaining the purpose of the tests. We also mentioned that the candidates should concentrate on the *change in quality* in opposite to the quality itself. We did not show the candidates a plot of the different levels of quality prior to the test (as shown in Fig. 1, Fig. 2 and Fig. 3). The listeners had to evaluate the audio sample by specifying a value between 1 and 5 for the subjective perception of the presented material. This scheme is known as the CCIR impairment scale to assess the audio quality and has been used during different international listening tests, for example the ISO-MPEG-1 process, which used in contrast to our series the "triple stimulus, hidden reference" method. The meaning of these values is:

5 = transparent, no difference to original signal noticeable

4 = perceptible, but not annoying

3 = slightly annoying

2 = annoying

1 = very annoying

* SchokoPro, Sonnenberger Str. 82, D-65193 Wiesbaden

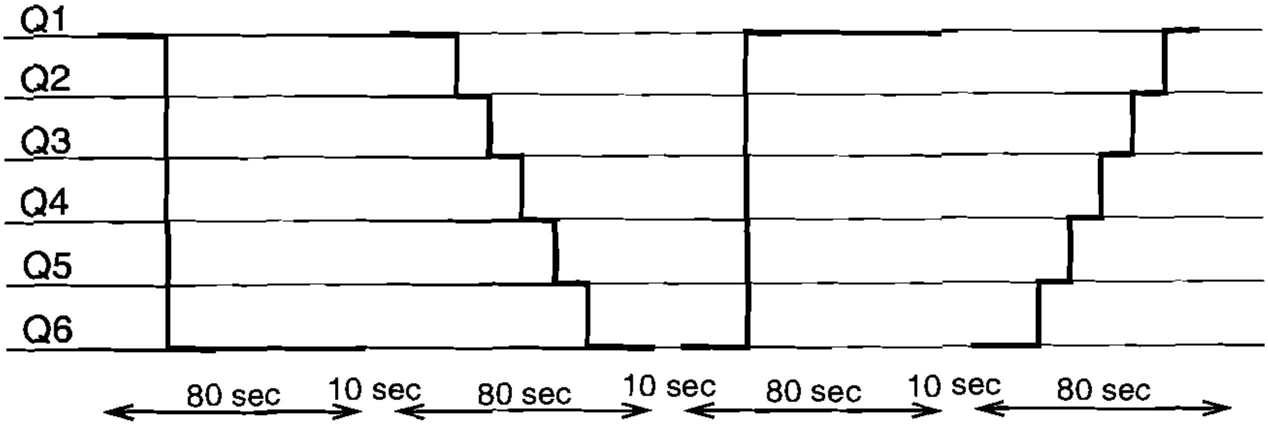

Figure 1: Session 1

At the time being we have performed a set of preliminary experiments with 3 test groups of about 5 candidates each. Based on these results we will carry out the major tests with a set of about 100 candidates later on.

The candidates had to answer the following questions about their musical background first:

- Do you play an instrument, or did you receive any musical education?
- How would you describe your hearing capabilities? (Several answers possible)
    - ❑ I can hear absolute pitches
    - ❑ I can hear changes in tempi
    - ❑ I can hear if a single instrument is pitched different among multiple other instruments
    - ❑ I can hear if a single instrument is pitched different among few other instruments
    - ❑ I can pitch a guitar
    - ❑ I can hear, if the overall sound is good
    - ❑ I am just a listener, without caring too much about sound
    - ❑ No opinion
- How would you describe your musical skills?
    - ❑ very good
    - ❑ medium (like a singer in a typical chorus)
    - ❑ low
    - ❑ none
- How would you describe the quality of the sound system, the test is performed with?
    - ❑ professional studio equipment
    - ❑ high end HiFi system
    - ❑ average HiFi system
    - ❑ high end portable system
    - ❑ average portable system
    - ❑ low end portable system

These questions were typically answered with the aid of one of the testers.

Session 1:

After the introduction we played session 1 according to the plot shown in Fig. 1. The candidates had to evaluate each of the three 80sec sequences with the scale described above. The different levels of sound quality refer to the levels described in Table 2. Each sequence was assembled to

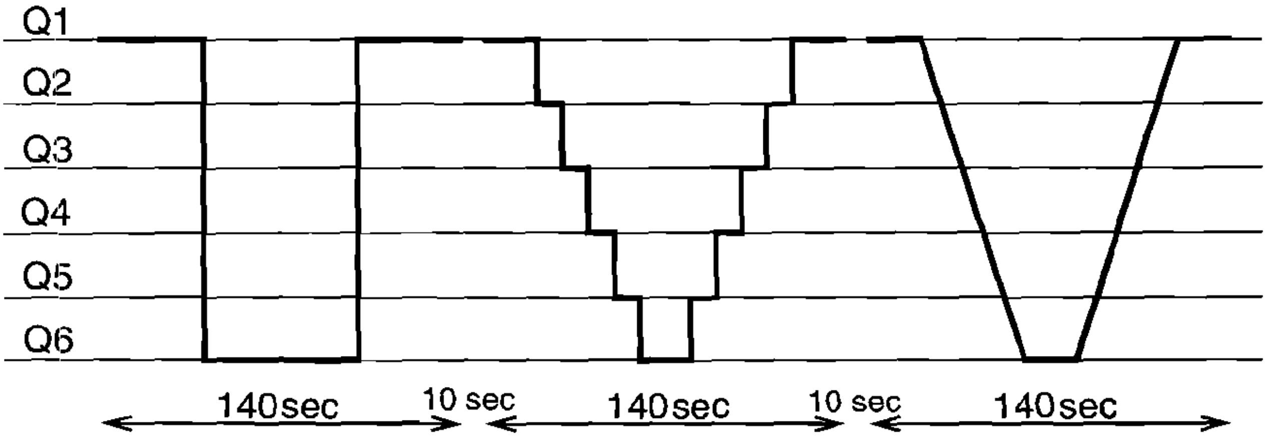

Figure 2: Session 2

have the same length. The candidates were instructed to evaluate each complete sequence as a unit and not each change in quality by itself.
Session 1 was intended to answers the following questions:

- Is it better to scale from a maximum quality level to the lowest possible level with or without intermediate steps?
- Is there any difference in scaling from up to down or from down to up?

Quality	HighPass-Filter	LowPass-Filter	Level	Quantization	Limiter	NoiseReduction
Q1 CD	none	none	0.0dB	none	none	none
Q2 MM-PC	210Hz 18dB	4800Hz 18dB	+3.4dB	none	none	none
Q3 Telephone	300Hz 18dB	3400Hz 18dB	+5.2dB	none	none	none
Q4 Radio	500Hz 18dB	2000Hz 18dB	+8.4dB	none	none	none
Q5 Mini-Radio	560Hz 36dB	1800Hz 36dB	+12.4dB	none	none	none
Q6 approximately GSM-like	830Hz 36dB	1200Hz 36dB	+18.0dB	8bit dithering: type1 shaping: ultra	-0.1dB threshold -10dB	broadband –18dB att: 5ms rel: 2ms smoothing 80% Hi shelving: 1200Hz -6dB

Table 2: Levels of sound quality used during the final tests

Session 2:

The next session is plotted in Fig. 2. Although it looks quite similar to session 1 it is intended to answer different questions leading to new results. Here the candidates had to rank each sequence compared to the others and finally to evaluate them individually using the same scale described above.

Session 2 was intended to answers the following questions:

- Is there any difference in scaling by steps or scaling by crossfade between different levels of quality?
- How does scaling down to the lowest level quality, holding there for a long time and scale up again compares to schemes where the lowest level of quality is only kept for a short time?

Session 3:

The final session (plotted in Fig. 3, where the gap between the individual samples is 5sec) was intended to judge about the chosen number of quality levels and to answer questions about preferences in the direction of jumps in quality. Jumps in quality level occurred in both directions. The start level of a jump was lowered from CD quality to GSM-like quality and back up again. This was intended to show if the candidates will get accustomed to a lower level of quality. The candidates were asked to evaluate every single jump as soon as the sample was played.

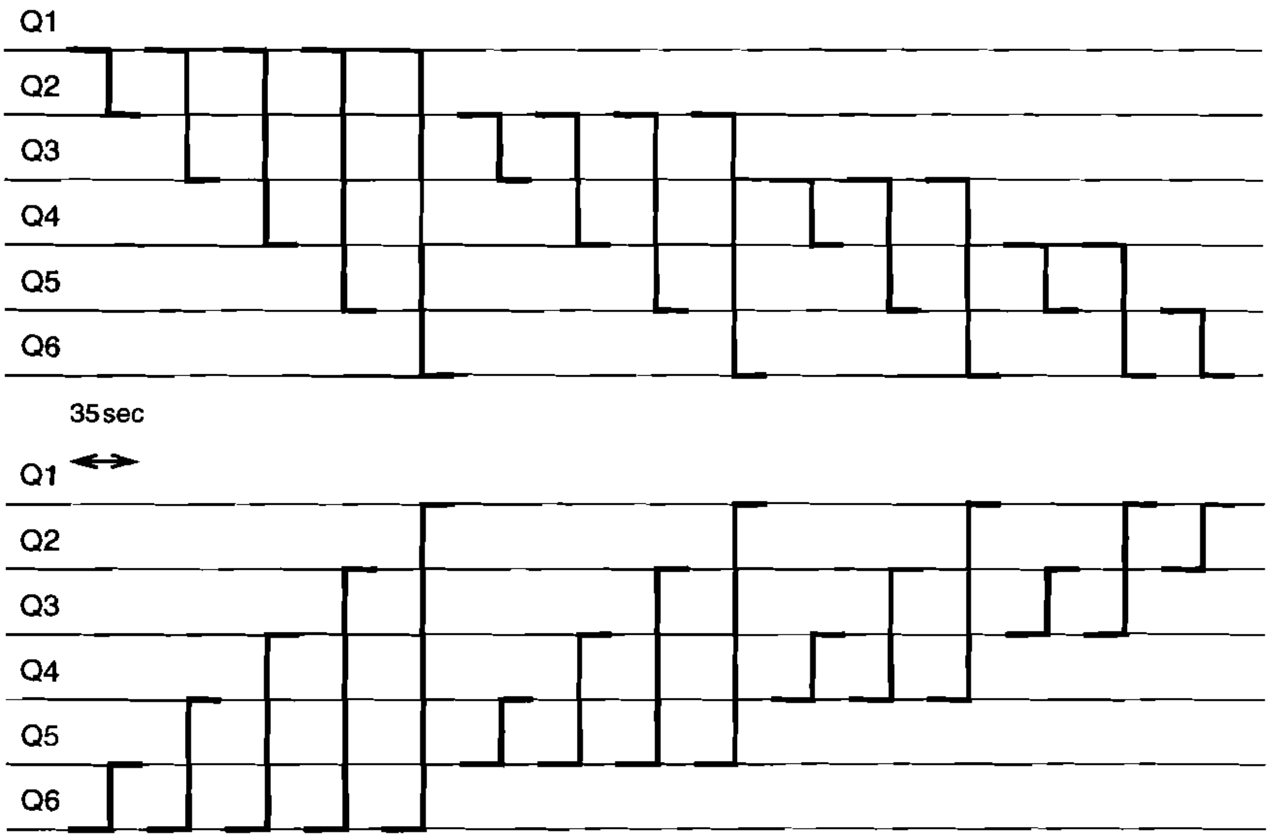

Figure 3: Session 3

7. Results

Our first aim was to find out the appropriate number of steps of differing quality. We also had to select the absolute quality level of each step and to check if we have chosen a suitable alternative. In approximately all cases the candidates still notice the difference of quality of adjunct steps. However the change is sufficient to have such a step defined as another level of quality. The results show that an increase in the number of steps make quality changes during the transition from on step to an adjunct step less noticeable.

Session 1:

In the first session (see Fig. 1) we wanted to check the difference between an increase and a decrease in quality. All candidates mentioned that an increase in quality is better than any decrease – for the average ratings refer to Table 3. This is independent of the kind of increase (stepwise or in one large jump). The interpretation is simple: In our natural behavior as human beings we are more happy to perceive an increase in quality.

Jump down	Step down	Jump up	Step up
1.2	2.1	2.9	3.0

Table 3: Average ratings for Session 1

In the same session we also compared different methods to increase quality where we found out that there is no significant difference whether in increasing the quality in one or in various smaller steps. The situation is different if the quality is decreased: A stepwise approach is significantly preferred over a large decrease in quality.

Session 2:

The main question behind the second session (see Fig. 2) relates to the usefulness of steps and a comparison with today's alternative (one large step) and the smoothest possible approach (fading). Fading turned out to be the best alternative in any case. However it is difficult to fade over a large range (perhaps leading to data rates between 6.2kbps and approximately 700kbps).

Jump	Step	Fade
1.9	1.9	2.7

Table 4: Average ratings for Session 2

Somehow strange were that approximately half of the candidates preferred the change in quality via several steps more (in the scale introduced in section 6. Experiments) than a large jump, while the other half scored vice versa. The average ratings (see Table 4) show an equal value of 1.9 for the jump or stepwise approach. We discussed this with the candidates and got two contradictory informations: The first group told us that they prefer to have small changes, others would mind to have many changes (they prefer one larger step) as they are disturbed by many changes. This is an issue for further research. The constraints must be understood more in detail and more experiments need to be carried out.

Session 3:

In the third session (see Fig. 3) we compared the impact of different sizes in quality changes, first by decreases and later by increases in quality. In general all candidates rated larger changes in quality worse than changes by smaller steps. The average ratings are listed in Table 5 for the decrease of quality and in Table 6 for the increase of quality. Corresponding steps up or down are listed above each other and shaded cells mark single step transitions. Decreases in quality are rated worse than an increase. A significant dislike of lower qualities in general is our interpretation, so it seems obvious that the candidates not only rated the changes in quality, but also the absolute quality level.

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
3.65	2.64	1.82	1.09	1.00	4.09	2.91	1.64	1.27	3.27	1.91	1.27	2.36	1.55	2.82

Table 5: Average ratings for session 3 – decrease of quality

30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
3.64	3.09	4.27	3.09	3.45	3.91	2.64	2.82	2.73	3.18	2.09	2.09	2.18	2.18	3.09

Table 6: Average ratings for session 3 – increase of quality

For the increase of quality it does not make much difference if a transition occurs between 2 or more steps. This result is interesting, as session 1 showed that an increase in quality by the stepwise and the one large jump approach is similar in terms of subjective quality perception. Our interpretation it that smaller steps are preferable if the quality is expected to change often, while a large jump is preferred when the higher level of quality is expected to stay constant for a longer time.

We carried out all experiments on different type of equipment and in diverse environments. We found out, that the expectation and the actual equipment influenced the results: Some playback experiments were (intentionally) performed using a portable cassette recorder, others were played back on high quality HiFi equipment. Most candidates which listened to the portable device gave better grades than those listening at high quality equipment. A careful analysis also showed that humans tend to accept a variation of 3 levels in quality (related to the level of quality of the play back equipment). I.e. so far in any case a jump up to 3 steps (with respect to the quality of the playback equipment) downwards can be tolerated. Larger steps (downwards) should be avoided.

8. Implications

The results above lead to the following implications:

User QoS may be raised if audio coding is scalable to a higher extent than it is today, such that application and system software may control the audio data rate dynamically. It would also be helpful to have smooth transitions defined for changing the data rate. This would allow to activate fading capabilities and hereby enhance the audio user interface.

Session 1:

The comparison of increase and decrease in quality showed that any multimedia scaling system is, in general, free to choose the most suitable way and timing to adapt the QoS of the respective media stream whenever quality is going to be enhanced. With the results from session 3 this is true as long as a decrease in quality is not expected to occur approximately less than 10sec after the last increase, a situation where a stepwise approach is preferable. However, any decrease in quality should be performed in smaller steps (and not in one single jump). This result is of no advantage for the communication and operating systems as the reaction to congestion would only be a smooth decrease of the load.

The transport of audio data should not be performed according to the traditional TCP congestion avoidance algorithm with it's well-known slow start and fast stop behavior. Our results show, that scaling has to be done in the opposite direction, e.g. fading down slowly towards a lower quality, while scaling up to a higher quality can be done much quicker, although this scheme does not fit into the congestion avoidance algorithms used in today networks.

Session 2:

System and network QoS should be restricted to a certain lower limit, i.e. during the connection setup phase a minimum QoS should be reserved and adaptation should only occur above this lower limit. We recommend to have a predefined lowest value according to the level of quality Q5.

Without better audio compression for low bandwidth it will be difficult to have a seamless integration of cost effective mobile equipment into existing networked systems and applications.

The fading approach results in a value of 0.8 higher compared to a stepped change of quality, so this scheme should be used whenever scaling of audio data is needed. The same is expected for video data. Most of todays audio encoding schemes are optimized for specific bitrates, so new encoding schemes, which allow scaling over a wide range need to be implemented. Further research should be performed on this specific field.

Session 3:

A more sophisticated increase in quality (instead of simply going to the best quality level) can be achieved by changing slowly step by step. However it is an open question how to specify the best value for the "slow" adjustment. In case of an increase by higher degrees it is recommended to make a substantial change immediately, as long as the higher level of quality can be hold for a certain amount of time. Any decrease in user QoS should be performed in a slowly manner in all cases.

We understand this as a preliminary set of experiments which will lead to new ideas of what, when and how to scale audio data in networked multimedia systems. In particular multimedia tools and services, including multimedia authoring tools, need to be enhanced by components to handle scaling of different media in different ways to achieve maximum quality for the end user and not only to avoid congestion or to achieve maximum throughput.

Acknowledgments

We want to acknowledge the patience and the various excellent comments of all our test candidates. Lars Wolf provided many detailed comments. Thanks to SchokoPro, Wiesbaden, Germany we were able to record and edit all the music samples on professional audio equipment.

Bibliography

[BDS96] I. Busse, B. Deffner and H. Schulzrinne: Dynamic QoS control of multimedia applications based on RTP; Computer Communications, Vol. 19, No. 1, 1996, pp. 49 – 58

[BTW94] J.-C. Bolot, T. Turletti and I. Wakemann: Scalable feedback control for multicast video distribution in the Internet; SIGCOMM, London, UK, 1994, pp. 58 – 67

[DHH+93] L. Delgrossi, C. Halstrick, D. Hehmann, R.G. Herrtwich, O. Krone, J. Sandvoss, C. Vogt: Media Scaling with HeiTS; ACM Multimedia Systems, Vol. 2, No. 4, October 1994, pp. 172 – 180

[GoV93] C. Gonzales, E. Viscito: Flexible Scalable Digital Video Coding; Signal Processing: Image Communication, Vol. 5, No. 1-2, February 1993, pp. 5 – 20

[Jac88] V. Jacobsen: Congestion avoidance and control; ACM Computer Communication Review, Volume 18, August 1988, pp. 314 – 329

[NaS95] K. Nahrstedt and R. Steinmetz: Resource Management in Networked Multimedia Systems; IEEE Computer, Vol. 28, No. 5, May 1995, pp. 52 – 63

[StN95] R. Steinmetz and K. Nahrstedt: Multimedia: Computing, Communications & Applications; Prentice Hall, 1995

[Tan96] A. S. Tanenbaum: Computer Networks; Prentice Hall, 1996

[Wol96] L. C. Wolf: Resource Management for Distributed Multimedia Systems; Kluwer, 1996

[WHD95] L. C. Wolf, R. G. Herrtwich and L. Delgrossi: Filtering Multimedia Data in Reservation-based Networks; Kommunikation in Verteilten Systemen, Springer Verlag, Feb. 1995, pp. 101 – 112

[ZDE+93] L. Zhang, S. Deering, D. Estrin, S. Shenker, and D. Zappala: RSVP: A new Resource ReSerVation Protocol; IEEE Network, Vol. 7, No. 5, September 1995

Biography

Reinhard Bertram studied Mechanical Engineering with a focus on computer simulations at the Technical University of Darmstadt, Germany, where he received the M.Sc. (Dipl.-Ing.) in 1996. He worked at the University's Computing Center as a system administrator from 1994 until 1996. Since April 1996 he is a research assistant at the Electrical Engineering and Information Technology department of the Technical University of Darmstadt. His current research interests include audio scaling, distributed systems and systems support for multimedia communications.

Ralf Steinmetz received the M.Sc. (Dipl.-Ing.) degree and the Ph.D. (Dr.-Ing.) degree, working in the area of Petri nets and concurrent programming languages with focus on communications, from the Technical University of Darmstadt, Germany, both in electrical engineering, in 1982 and 1986, respectively.

He is Professor of Electrical Engineering and Information Technology, Technical University of Darmstadt. He led the successful IEEE Multimedia Task Force Working Group for magazine publication and is now the associate Editor-in-Chief of IEEE Multimedia. He has served as editor and member of the editorial advisory board of several journals and as chair, vice chair, and member of numerous program and steering committees for workshops and conferences. He managed the multimedia department at the IBM European Networking Center, Heidelberg, Germany. There, in 1988, he initiated the first activities on networked multimedia issues. Since then he has been in charge of several multimedia projects. He has worked in academia, supervising numerous theses and lecturing. He wrote the first in-depth technical book on multimedia systems published in German in 1993; a co-authored and restructured edition in English published in 1995. Befor joining IBM, he worked in ISDN development activities for Philips.

SOCIETY OF MOTION PICTURE AND TELEVISION ENGINEERS

The 1996 European SMPTE Conference on Imaging Media

September 19–21, 1996

organized by the European Sections of the SMPTE in association with the Montreux International Television Symposium

during photokina PROFESSIONAL MEDIA

photokina
Köln 1996
Sept. 18 – 23

Montreux 1997
June 12 – 17

Köln Messe