

1. Introduction: The Scope

We understand multimedia systems being characterized by the integrated computer-controlled generation, manipulation, presentation, storage, and communication of independent discrete and continuous media. Both synchronization and resource management are considered to be fundamental problems in multimedia systems. In this context, the notion synchronization covers a wide range of aspects, such as the synchronous play-out of multimedia data as well as orchestration the activities of collaborating team members. Resource management forms the basis for communication and computation in multimedia systems, and in particular for synchronization.

Synchronization of multimedia data can be considered at two levels of abstractions, the stream level and object level. Similarly, for the synchronization of user activities in cooperative environments, one can distinguish between the synchronization of user interactions on shared data and the coordination of team members.

Resource management and synchronization are strongly interrelated. Appropriate resource reservation and scheduling techniques are used to ensure the required quality of service (QoS). Those techniques as well as QoS specifications are to be considered at the network and operating system level. In addition, QoS is an issue at the end-user interface.

The seminar brought together researchers from the different areas of synchronization and resource management. The participants investigated what new requirements for synchronization and resource management are emerging from advanced multimedia applications, such as virtual reality, interactive TV and interactive games. The goal of the seminar was beyond a discussion of the impact of those requirements on synchronization and resource management methods at the various levels of abstraction; it was to indicate the interrelationships between these levels and areas and to discuss the issue of an integrated synchronization architecture. Topics treated at the seminar include:

- Requirements emerging from new multimedia environments
- QoS specification for synchronization and resource management at various abstraction levels
- QoS mapping
- Models and algorithms for stream and object synchronization
- Models and algorithms for synchronization in cooperative multimedia environments
- Resource reservation versus no reservation
- Scheduling techniques for continuous media
- Dependencies between resource management and synchronization
- Integrated synchronization architecture
- Demos of applications

We thank all our colleagues for their very active participation and contributions to our research community.

*Nicolas Georganas
Thomas D.C. Little
Kurt Rothermel
Ralf Steinmetz*

2. Program

The seminar was structured into the following sessions:

- Applications, New Environments
- CSCW, Telecooperation
- Synchronization, Internet
- Synchronization
- Adaptive Applications, Media Scaling, Filtering
- System QoS I
- System QoS II

Associated with each session was a one hour panel discussion. The abstracts of all presentations of these sessions are included in this report.

3. Applications, New Environments

3.1 Architecture for Interactive Networked Multimedia Applications Tobias Helbig, Philips Research Labs, Aachen, Germany

The buzzword “Multimedia” is around for some time now. So are the dreams associated with it, the prototypes and first products. Increasing performance of networks and digital equipment step-by-step allows the digital information technology to move into high-end professional domains. There, new challenges pop up: High-end quality still brings our technology to the limits and requires algorithms that are not yet in place. Paying customers in the professional domain require a reliability of products that is very hard to achieve with current hardware and software. And, last not least, there are analog and even “old-fashioned” digital solutions in place that are proven, offer incredibly good features and have to be watched with regard to pricing.

The talk focused on a multimedia middleware architecture. Such an architecture is expected to be the core of an evolution from current application support functions for professional applications to future ones. The evolution will be based on improvements in available technology and the emergence of new protocols. I presented results from our experiences with a prototype implementation and of our discussions with regard to the decisions we made.

3.2 Issues in Interactive Networked Mobile Media Manipulation Ralf Steinmetz, Darmstadt University of Technology, Germany

The aim of this research work at the GMD-IPSI is to “nicely” combine visualization issues and mobile/mobility environments. Today we encounter very different networking infrastructure capabilities (from less than 9.6 kbps up to more than 150 mbps, latencies between a few milliseconds up to more than a second, disconnected operation vs. on-line, etc.) and heterogeneous devices (high resolution SGIs down to e.g. Nokia Communicator with small screen and very limited processing power). Media contents exist for all of these environments as there are many

different documents and document models. The ultimate challenge is to have one single document/media content and adapt it properly to “whatever” environment we encounter. This implies to know about the contents and scale media data to a large degree. We call this approach/challenge “Scaling in the large” which includes deep knowledge and making use of QoS as well as pricing constraints.

3.3 Authoring Support for Synchronized Multimedia Delivery **Dick C.A. Bulterman, CWI, Amsterdam, Netherland**

The basic thrust of my work is that the specification of synchronization constraints is inherently a content-based matter. Nearly all of the decisions that (should) guide the selection of QoS or adaptation control are influenced by these content-based relationships.

At CWI, we have looked at authoring support in which a presentation creator can:

- specify which alternatives should be selected if QoS or adaptation takes place
- specify the flexibility that a user can have in altering a presentation
- specify parameters that can guide the “implementation” environment

We use these specifications to drive a particular presentation of hypermedia documents. The hypermedia base provides additional clues that can provide discrete points at which QoS/adaptation decisions can be implemented.

4. CSCW, Telecooperation

4.1 Synchronization and Computer Supported Cooperative Work(CSCW) **Hans Schlichter, Technical University of Munich, Germany**

CSCW requires five different types of synchronization which occur either on the system level or on the human level.

Synchronization of group membership refers to the management of dynamic groups. People may join or depart from an existing group. Despite these dynamic changes all group members need a consistent view of the group. Thus, the distribution of information and messages depends on the group structure and group organization.

The second type is the synchronization of awareness information. Awareness could be interpreted as part of the glue that allows groups to be more effective than individuals. Improved awareness within a collaborating group encourages informal spontaneous communication and keeps the group members up-to-date with important events. Group work incorporates group activities which may be causally and temporally dependent. The actors executing these activities must coordinate their actions in order to manage these dependencies, e.g. resource allocation or sequencing of activities.

Communication during group work encompasses often multiple media streams, e.g. audio, video and multiple data streams (telepointer, shared applications). The communication channels must be synchronized and adapted to communication characteristics (e.g bandwidth, load etc.).

The last type of synchronization refers to the synchronized access to shared resources, e.g. group documents which are jointly manipulated. If the documents are replicated at each site of a group participant, the multiple copy update problem must be handled. Several different approaches of concurrency control are applicable ranging from preventing and reducing divergence to accepting and ignoring divergence.

4.2 The Provision and Management of Multimedia Services: how are they possible with the Telecommunication Architectures? Jean-Pierre Hubaux, EPFL, Lausanne, Switzerland

The Telecommunication Community has been investing a lot in the definition of architectures to create and manage services. The Intelligent Network, aimed initially at the provision of new telephony-oriented services, is now widely deployed and should continue to play an important role at least for mobility services based on voice (GSM, DECT, but also UPT).

However, the success of service deployment over the Internet will oblique the strategy of the Telecoms to be significantly revisited. Indeed, an initiative such as TINA, in spite of its capacity to embrace both service engineering and network management issues, is based on assumptions which are not verified any more. For example, the end-to-end availability of ATM solutions is more and more questionable, due to the expected success of IP-based solutions with a reservation mechanism a la RSVP. On the other hand, there is certainly a need for “hard guarantees” in certain circumstances, for example for the provision of Intranets comprised of interconnected CPNs (called sometimes “virtual intranets”). It is therefore suggested to reconsider the TINA architecture, keeping the concept of information network, and splitting the transport network resources into two parts, the first with soft, and the second with hard guarantees.

4.3 JETS (Java-Enabled Telecollaboration System) and Web Telecollaboration Tools Nicolas Georganas, University of Ottawa, Canada

JETS permits synchronous telecollaboration among users equipped with Java-enabled web browsers (any such browsers). A server (written in Java for portability) sends Java applets (whiteboard, chart, text-editor, VRML, HTML shared space) to clients addressing the server URL. These applets permit real-time telecollaboration among the clients. JETS has a management system interface that permits a “chairperson” to set access and control rights.

The current prototype does not provide voice and video support, but a video H.263 Java applet has been developed. The talk also gives an overview of available web telecollaboration tools, such as MS Net Meeting, Netscape Conference and other. The JETS software can be downloaded from the URL <http://www.mcrlab.uottawa.ca/jets>.

4.4 Multimedia at the User Interface

Max Mühlhäuser, University of Linz, Austria

Net and Operating System support for Multimedia, QoS and Synchronization models are all design-centered around the notion of periodic samples. Recent experiences with User Interfaces challenge this assumption:

1. Multimodal User Interfaces, where compute-intensive, periodic modality-specific processing (e.g. frame grabbing → gesture) is in closed loop with “logical interaction” (e.g. command/position) and application core (e.g. VR game), all distributed.
2. Comic Actors, representing Internet Agents or serving as Avatars. While the final-form Comic Actors behave similar to movies, they have to be assembled from building-blocks on the fly which is compute intensive and act in closed loop with the (Internet) application and the user with whom they interact.
3. The video Web-interface “DBVI”, based on a hierarchy of 3D VR models (“shoebox” filled with key frames of video). On the web, DBVIs must be computed/loaded on the fly, the data-stream is real-time, non-periodic.
4. The “WorldBeat” music station based on Virtual Batons:
 - a. coordinates of batons must be translated into (non-periodic) coordinates of “beats”
 - b. the system works on the non-periodic, highly abstract model of “aesthetic music piece” instead of MIDI (“notes”) or wave forms.
5. The “PinUp” cooperative interactive whiteboard with its object-oriented “movie” model: Users can go back on the timeline, hence non-periodic variable volume data must be processed in real-time, possibly (traditional) user audio/video recordings are attached.

All in one, the examples represent a closed-loop real-time processing of non-periodic time-dependent data with periodic processed and/or closely interacting users.

5. Synchronization, Internet

5.1 Talking Heads in Multimodal Interfaces

Georg Fries, Deutsche Telekom Berkomp, Darmstadt, Germany

Today, most computers use window-based graphical user interfaces. By adding speech recognition and speech synthesis, limited voice capabilities can be included. Advances in micro-electronics are now providing low cost sensors, which can be used to implement further input modalities to the human computer interface. So in the future, we might expect additional input skills of the computer which will lead to real multimodal interfaces.

In such a multimodal environment, an animated synthetic face (Talking Head) can represent the computer part of the interaction, to give user the impression of a face-to-face communication. Several 3D-face models have been implemented. Such models can be seen as a geometric description of the facial surface, which is deformed by varying face parameters. In Darmstadt we use a cartoon face which shows some advantages compared to a complex 3D face:

- The effort of modification and rendering is much smaller.
- We can use exaggeration for modeling articulation and facial expressions.

To give the talking head a voice, we prefer the use of speech synthesis instead of “canned speech”, because of its flexibility. Furthermore, the TTS frontend provides the lip-synch information. In our future work, we want to define a high-level API with sophisticated control parameters, so that we can exchange the face model without changes in control parameters.

5.2 Internet Synchronization Architecture **Henning Schulzrinne, Columbia University, New York, USA**

Multimedia synchronization in the Internet can be divided into three separate functionalities:

- frame and sample-level synchronization, in particular lip-synch;
- synchronizing several streams, possibly originating at different servers, to start playing together;
- describing sessions of multimedia objects synchronized in complex timing relationships.

Frame and sample-level synchronization has to happen at the transport layer and can be provided by RTP. RTP synchronization assumes that data sources (but not necessarily receivers) have access to a synchronized global clock. Modern versions of the Network Time Protocol (NTP) can readily synchronize clocks to within a few milliseconds across the Internet; thus, it is not necessary for multimedia synchronization to reinvent clock synchronization. RTP data packets carry a timestamp that increments at the nominal sample rate, such as 8 kHz for telephony-quality audio or 90 kHz for video, regardless of the actual source sampling clock rate. Every few seconds, RTP sources emit an RTCP packet containing a mapping between synchronized “wallclock” time and a recently set RTP timestamp. The receivers can thus compute and equalize their absolute playout delay for all synchronized streams.

Synchronization of stream start times is provided by RTSP (Real-Time Stream Control Protocol), a “VCR remote” for the Internet. RTSP can initiate streams at distributed servers to start sending at roughly the same time. Fine synchronization is provided by RTP, as described above, so that differences in start times only incur additional buffering at the receiver.

Session description “languages” provide higher-level specifications of the synchronization relationships of media streams. The Session Description Protocol (SDP) currently in use for Mbone multicast sessions can only associate several media streams proceeding in parallel, while efforts within W3C SYMM WG and languages like RTSL are based on a nesting of timing and choice structures.

5.3 Temporal Synchronization of Multimedia Scenes **Chabane Djeraba, University of Nantes, France**

We present an approach for temporal and interactive relation composition of media. The approach is based on time-intervals and dependency temporal relations. We consider the seven relations of Allen (equals, meets, finished, starts, before, overlaps, during) with the following features: Firstly, the temporal relations are designed to specify relations between multimedia

objects of both determined and undetermined duration. Secondly, the temporal relations describe both the existing arrangement of media, and dependency relations between media. Finally, a powerful temporal petri net is generated automatically to model interactive and temporal relations.

The petri net is stored in an object called “scenario object”. The user may request the interpretation and the simulation of the scenario object which leads to scenario presentations with domain expert interactions. The petri net permits a formal specification and a proof of scenarios.

5.4 What is Really Hard About Media Synchronization, Even in Low-Latency Networks?

Daniel C. Swinehart, Xerox - Palo Alto Research Center, USA

In media capture, transmission, and rendering environments where a very high quality of service can be guaranteed, in particular, data loss and transmission jitter from end-to-end is negligible, playout synchronization of related and/or contemporaneously produced but geographically separated media is straightforward but nevertheless necessary. Under the assumption that all hosts in modern internetworks can obtain accurate knowledge of absolute global time, the only adaptation required for tight synchronization is adaptation to skew in frequency between sending and receiving devices. This requires increasing or decreasing, at the destination, the number of frames or samples presented to the rendering engine. In high-quality environments it is important to accomplish this in a manner that is imperceptible to the user.

This presentation focused on audio adaptation, suggesting a number of effective mechanisms for sustaining a continuous presentation while inserting or removing samples. Simple drop-a-duplicate methods, sample rate conversion, single-bit add-a-drop, or more advanced techniques which are under development can be used. If this approach is taken, synchronization between related audio and video takes care of itself.

6. Synchronization

6.1 Temporal Representation for Constructed vs. Construed Multimedia Content

Thomas D.C. Little, Boston University, USA

Scripting is a typical production activity during multimedia content creation. However, script information becomes implied in the final, rendered production, and is no longer accessible as information on which to infer relationships among structural or conceptual entities; nor is it suitable for database access in the context of a universe of multimedia presentation objects.

In contrast to this “construction” process, composed multimedia content can be “construed” to yield similar script information, although of a more conceptual, or subjective, nature derived from the final, rendered juxtaposition of multimedia components.

In this talk I consider a common metadata representation to characterize both structural and conceptual entities derived from construction (scripting) or from analysis. Uses of the resultant metadata representation include the support of content- and structure-based access (query) on a collection of multimedia entities, temporal analysis of conceptual entities (e.g., persons in video scenes), and automatic assembly of components in object repurposing. As proof-of-concept, we focus on applications in personalized news video delivery and automatic news item composition from a large collection of annotated news clips.

6.2 Synchronization Models in Multimedia Standards

Wolfgang R. Herzner, Forschungszentrum Seibersdorf, Austria

Multimedia Standards provide several approaches or models for the specification of synchronization of multimedia presentations, where “synchronization” is understood as presentation control or, in some respect, as “layout in time”.

Basically, two approaches can be identified: HyTime (IS 10744), which is an extension to SGML for hyperlinking and temporal aspects, adapts the concept of “layout in time” by allowing to place presentable objects in coordinate spaces build up of different axes including time. Besides a rather high richness provided by HyTime with respect to specify positions and extents in such spaces, this model allows to express all possible mutual dependencies of two temporal durations (of unknown length) including overlap and end-synchronization. However, this occurs on the cost of a) needing a scheduling step before presentation, and b) making it difficult to consider external events like user inputs.

The other approach, which is taken by most other standards, is event-driven, which makes it easy to consider user-input or system events and does not need any prescheduling, while giving up the complete expressiveness with respect to temporal relationships (i.e., overlap, inclusion and end-synchronization cannot be expressed). For instance, MHEG (IS 13522) supports presentation control by means of its “link” object class, which “condition” describes when its effect – an “action”-object – shall be executed. In addition, MHEG knows the concept of a timeline attached to each scene, where timer can be set to raise corresponding events at selected points in time.

VRML knows a similar notion of temporal axis; in addition, it allows routing of events between the nodes of a VRML-scene.

Finally, PREMO (DIS 14478), which provides a development environment for graphics and multimedia applications, specifies several object types for synchronization like the “controller” or the “event handler”, and also has a concept of time-line which can be compared to that of MHEG.

6.3 Modelling and Scheduling of Adaptable Interactive Multimedia Documents

Stefan Wirag, University of Stuttgart, Germany

Multimedia presentations are applicable in various domains such as advertising, commercial presentations or education. If multimedia documents can be accessed on-line via different net-

work types and be presented on various types of terminals, different amounts of resources may be available at presentation time. Hence, it can happen that there are not enough resources to render a multimedia document according to the specification. For usual multimedia documents resource scarcity implies an arbitrarily reduced presentation quality. To handle resource scarcity in a better way, multimedia documents should be adaptable to different resource situations.

The presented document model provides abstractions to specify multimedia documents with alternative presentation parts. Further on, the presentation behaviour of media objects can vary within specified limits. Hence, the presented document model allows to compose presentations which have a defined behaviour when resource restrictions occur and resource scarcity need not result in an arbitrarily reduced presentation quality. The presented scheduling algorithm adapts the presentation regularly with regard to the resource situation.

7. Adaptive Applications, Media Scaling, Filtering

7.1 Resource Reservation for Distributed Multimedia Applications **David Hutchison, Lancaster University, Great Britain**

The aim of this talk is to present and discuss our approach to resource management for distributed multimedia applications. We assert that resource reservation management is important to meet diverse Quality of Service (QoS) needs of multimedia telecommunication networks and applications, for both individual communication services and for the network as a whole. Management, in our work, involves monitoring and control of resources in the entire end-to-end progression of setting up, maintaining and closing down a session. It encompasses short timescale control and signalling as well as longer timescale issues of development of resources. We are developing a QoS-based architecture consisting of a set of hierarchical regions from application through transport and network layers to the physical communication paths and the endsystem resources. We expect the innovations of the work to be the following: the definition of QoS-based Service Interfaces to each region in the architecture; a QoS manager that maps application requests into managed end-to-end resources and a signalling system that provides the engineering basis for resource management.

7.2 QoS Filtering in an Internet Environment **Andreas Mauthe, Lancaster University, Great Britain**

The provision of QoS in a heterogeneous communication environment is still one of the most pressing problems in distributed multimedia systems. The Internet resource reservation protocol RSVP is receiver oriented and allows different levels of QoS for different receivers. However, it does not specify how the QoS requirements of a data stream can be reduced to accommodate the specific QoS requirements of a receiver.

The QoS filter development at Lancaster provides mechanisms to change the structure of a media stream and hence to adapt its QoS requirements. In order to provide optimal QoS it is necessary to tie resource reservation and QoS to integrate QoS filter control in the functional model of RSVP. A special QoS service class called QoS adaptation service was defined for the speci-

fication of resource and QoS filtering requirements. Additionally, a new RSVP object (QoS ad-spec) is used to convey QoS filter control information. Our enhanced version of RSVP is capable of interoperating with standard RSVP such that system openness is ensured. A prototype is currently being implemented within the Lancaster experimental environment.

7.3 Optimal Adaptive QoS Provisioning in Multimedia Database Systems **Erich Neuhold, GMD-IPSI, TU Darmstadt, Germany**

Due to the prevailing limitations in today's computing infrastructure, e.g. limited bandwidth, on request multimedia services should incorporate demand-driven Quality of Service adjustments that will provide „the best“ service under the changing resource availability that currently is caused by network instability, changing demands from single users or interaction/interference from multiple users.

A quality adjustment protocol was presented that distributes available resources among competing presentations, but also optimizes resource utilization between the different multimedia streams of a single presentation. The problem represents itself in a multi-dimensional frame as identical savings of resources may still result in a different total Quality of Service or conversely the same Quality of Service may be achieved by different amounts of utilized resources. Linear programming models have been developed by Heiko Thimm, also at GMD-IPSI, both for global (multi-presentation) and local (single presentation - multiple stream) situations. Preliminary experimental results have shown the viability of the approach in a number of different application settings, e.g. tourism, sports, teaching, where multi-stream presentations over Intranet or Internet present the most important rapidly growing markets.

7.4 Database Support for Interactive Multimedia Presentations **Wolfgang Klas, University of Ulm, Germany**

In the framework of a cooperation project with several medical clinics at the University Hospital of Ulm we aim at the development of a database system-based multimedia repository. The repository should serve as a central information server for several applications:

- doctors who specialize in their particular field would like to retrieve high-quality material for the purpose of training and qualification as well as for getting background material to be used in the context of diagnoses.
- lectures should be able to generate presentation material in a highly flexible way for courses and seminars.
- students should be able to retrieve material from the repository in the context of their particular studies.
- patients should be able to retrieve high level explanations related to their diagnosis and should be able to get help with respect to their information need.

The technical challenges for the multimedia repository can be summarized as follows: To serve various output channels like online services on as well as off the university campus, CD-ROM-based presentations, and printed material. Presentation material has to be stored presentation-neutral, e.g., text documents by using SGML, image, audio and video material has to be available in different formats with respect to quality (quantization, resolution, color, etc.). Material

is to be stored by means of presentation fragments which can be composed to larger presentations. The composition has to be supported in a highly flexible way which calls for modularity in presentation representation and quality description. This also requires for a model of quality at the level of the application which can be mapped to the various notions of quality at the system levels. Solutions currently under investigation include the usage of object-oriented database technology as well as extendable relational database technology. In this context first implementations of DataBlades for DICOM (Digital Imaging and Communication in Medicine) data and presentation descriptions have been realized.

8. System QoS I

8.1 Application QoS Specification for Resource Management

Richard Staehli, Informix Software, Portland OR, USA

This talk defines application-level QoS specification formally as a constraint on the distance of a presentation from the ideal. Our authoring and playback tools describe the ideal representation. The distance measure is based on an error model that satisfies formal criteria for soundness and completeness. The constraint is expressed as a range of acceptable values for a quality function that models user perception as a function of the error interpretation.

The QoS specification technique does not refer to the properties of the physical devices or data representation and is therefore device and data independent.

The implementation of the SQUINT multimedia player demonstrates that heuristics can be found to perform resource planning and the QoS guarantee algorithm efficiently. High level QoS requirements are guaranteed by distributing responsibility for the error components to service components which understand how the implementation may impact quality.

8.2 Network Resource Management for Distributed Multimedia Presentations

B. Prabhakaran, University of Maryland, USA

Multimedia Applications need to have an initial estimate of the Quality of Service (QoS) requirements before negotiating with the network service provider. We suggest methodologies for estimating the initial requirements and for negotiating with the network service provider. This can be done in a generic manner based on the synchronization characteristics of a multimedia presentation.

8.3 Audio Scaling

Reinhard Bertram, Darmstadt University of Technology, Germany

As a result of listening tests performed with users of multimedia applications, constraints for a dynamic adaptation (scaling) of audio data streams to the available system and network resources were presented. For the speech medium, any severe deformation of frequency characteristics should be avoided, whereas shortening and expanding of speech pauses is suitable to cope with

varying delays. For the music medium the best scaling algorithm is a smooth fade between quality levels, whereas it is often desirable to use discrete steps, because it is easier to switch the audio encoding than to scale within a single encoding. Scaling resulting in lower quality should be performed in multiple steps, whereas any approach can be chosen when scaling towards higher quality. It appears that six levels of quality are sufficient for scaling of music from high quality (compact disk) to very low quality (cellular phone) and vice versa.

Enhancements to networked environments based on a Dynamic QoS Centered Architecture, where the bitrate and scheduling of associated media streams are adapted dynamically, were presented. Scaling information may be embedded into media streams by using digital watermarks, which will impose an impact on performance, but will lead to new mechanisms to transfer control data without breaking existing applications.

8.4 Dynamic Bandwidth Reservation for Video on Demand over ATM **Torsten Braun, IBM ENC Heidelberg, Germany**

Dynamic bandwidth reservation (DBR) allows to adapt the reservation of network bandwidth resources to the application needs. Simulations showed that DBR can result in significant benefits for the ATM network provider and the service user simultaneously. The provider can support more connections to its customers while the user can decrease the amount of reserved bandwidth in order to save costs.

An algorithm has been developed which calculates the optimal negotiating times and the corresponding ATM traffic parameters to transport compressed (MPEG, M-JPEG) video streams. This algorithm shows better results (more admissible connections, more bandwidth savings) than other known algorithms.

To prove the concept, a prototype for M-JPEG video on demand transmission has been developed based on an AIX client/server environment. The protocol stack consists of RTP/AAL5/ATM and allows the video server to directly select the ATM CBR/VBR signalling parameters.

Unfortunately, the available signalling protocols such as UNI 3.1 do not support renegotiating very well. Nevertheless, we were able to get promising bandwidth savings up to 25% even for short M-JPEG movies, which are much less bursty and less suited than longer MPEG movies.

9. System QoS II

9.1 Qualitybased Heterogeneous Multicasting **Martina Zitterbart, TU Braunschweig, Germany**

The emerging communication infrastructure will be inherently heterogeneous with respect to the networking infrastructure (wireless vs. wired networks) as well as regarding the end user equipment (e.g. high end workstation vs. PDA). Therefore, a communication subsystem is needed that supports quality based heterogeneous group communications.

We follow an approach that is based on the invention of additional service components, such as adaptive server or synchronization modules which are part of the interworking unit. Therefore, an adequate signalling is required. The current implementation is based on an extension of

RSVP. It interfaces to a filter agent that is capable of merging, activating and re-parametrizing filters in the interworking unit.

The general paradigm behind the presented approach is that certain services (e.g. QoS filtering) are located inside the network. These modules might be loadable by the applications running on the end systems.

9.2 Time-Variant QoS Management

Klara Nahrstedt, University of Illinois and Urbana-Champaign, USA

The existence of heterogeneous computing and communication system platforms (e.g. ATM, Internet, UNIX, WindowsNT) and both old and new multimedia applications have given rise to the following questions: Can end-to-end QoS guarantees be achieved? What should an end-to-end QoS management architecture look like? We propose to have a flexible QoS management architecture, which relies on a new resource model. The resource model includes shared resources while is controlled by the resource scheduler with resource controller and resource broker. The resource broker performs the admission for resource availability as well as negotiation/renegotiation with the client. The resource controller performs adaptation as reaction to user or system resource changes. QoS in this resource model is aimed to be achieved for minimal requirements (QoS_min) as a soft guarantee. The adaptation is aimed for QoS_max as optimal possibility if client requests change.

We implemented the resource model on CPU and memory, where the CPU server provides low jitters for MPEG player in case of sharing the CPU with three other best-effort NRT applications. The memory broker and service control achieves further stability of the jitter requirements for the MPEG player.

Future work will integrate communication, CPU and memory resource control for provision of QoS at the endpoints. Problems are: arbitrary memory resource allocations and translations between applications and resources.

9.3 QoS Negotiation in Distributed Multimedia Systems

Kurt Rothermel, University of Stuttgart, Germany

Multimedia applications can be classified in rigid and adaptive ones. Rigid applications are based on the assumption of a guaranteed service. To ensure guaranteed service resource reservation is required not only in the network, but also on the operating system and application level.

We propose an application-level resource reservation protocol, which sits on top of a real-time network service. This protocol supports arbitrarily structured flowgraphs connecting media processing elements such as sources, sinks, filters and mixers. During the reservation process the filters are adjusted in an appropriate way. Furthermore, the protocol takes into account format constraints of processing elements as well as relationships between streams.

9.4 Resource Management and Adaptation

Lars C. Wolf, Darmstadt University of Technology, Germany

Over the last years, work has been performed in two directions to provide Quality of Service (QoS):

- Resource management systems – using reservation and enforcement, e.g., scheduling methods which require mechanisms in various systems and components,
- Adaptive applications and systems – using scaling and filtering mechanisms,
- (As further approach the overprovisioning of resources could be considered.)

Scaling mechanisms which use feedback information from receivers and networks to the sender have been applied successfully for some years already. Yet, since the data stream is changed at the sender with this method, it is not well suited to support heterogeneous systems.

Filtering tries to support such heterogeneous scenarios by changing the data streams inside the network. “General filters” have been proposed a while ago which perform operations such as translations, etc. on the data. There are, however, several open questions with this approach such as performance, security etc.

A simpler approach for filtering has been the “packet dropping” filter where some data is removed inside of the network. For this, the data must be partitioned into parts, e.g., using hierarchical coding schemes. The resulting parts can then be send independently or dependently over the network. Based on the experience gained in the design and development of these approaches, the various adaptation methods are compared.

9.5 Distributed Virtual Studio Services

Christian F. Breiteneder, University of Vienna, Austria

Virtual studios are a new video and television production technique attempting to overcome the constraints of traditional chroma-keying. In traditional chroma-keying the camera motion in the foreground layer is uncorrelated with the background layer. The result is that spatial relationships between the two layers are altered or lost. In traditional chroma-keying situations, such as a weather announcer and map, the problem is overcome by “locking off” the camera and using a static background. With a virtual studio, the static background is replaced by a dynamic, computer-generated, three-dimensional background. The foreground camera is then free to move but must be “tracked” so that the background can be generated with the proper perspective.

This talk serves two purposes: First, it gives a brief overview of current virtual studio technology. Second, a particular distributed production is discussed, in which 150 people attended the award ceremony of the “International Videokunstpreis” in Baden-Baden. For this production issues of delay and other QoS aspects are considered. The work was performed at the MediaLab of GMD in St. Augustin.