

Preface

From June 16 to 20, 1997, leading researchers from eight countries met at Schloß Dagstuhl to discuss High Performance Networks for Multimedia Applications. They came from telecommunications and computer networking backgrounds. It was the third workshop on this topic at the castle; the earlier two were held in 1993 and 1995.

Most of the contributions centered around three major issues: next generation Internet protocols, ATM networking (in particular traffic models and traffic shaping), and quality-of-service for multimedia applications. Already today much of the bandwidth of high-performance networks is used by multimedia applications, transmitting digital audio and video streams. Traditional networking protocols are not appropriate for these new applications; for example, they do not provide guarantees on bandwidth, end-to-end delay or delay jitter, and they do not have addressing schemes or routing algorithms for multicast connections. Also, the formal traffic models published in the communications literature do not capture the properties of multimedia streams very well. The presentations and discussions at the workshop addressed these issues, and they proposed interesting and innovative solutions.

This report is a collection of abstracts of all the presentations given at the workshop. Its purpose is to provide a quick overview of the topics. In order to keep the sessions informal, the participants were not required to submit a paper in advance. However, the organizers volunteered to edit a volume of proceedings if potential authors were interested. And indeed, a selection of ten high-quality papers has appeared under the title "High-Performance Networks for Multimedia Applications", edited by the organizers of the workshop and published by Kluwer Academic Publishers in November 1998.

The organizers wish to thank all the participants for coming to Dagstuhl, and for the many lively and inspiring discussions. In the long evenings, many of us made new friends, and decided to work together more closely in the future. And again, as in 1993 and 1995, the Dagstuhl team provided us with a very pleasant and hassle-free environment: *this place is optimized for thinking.*

André Danthine, Université de Liège, Belgium
Wolfgang Effelsberg, University of Mannheim, Germany
Domenico Ferrari, Università Cattolica di Piacenza, Italy
Otto Spaniol, RWTH Aachen, Germany

Contents

A New Internet Generation.....	6
<i>Guy Pujolle, Université de Versailles, France</i>	
The Importance of Flexibility in the Design of Network and Transport Layer Protocols	6
<i>Stephen Pink, SICS and Lulea University, Sweden</i>	
Active Networking using Java and IPv6.....	7
<i>Doug Shepherd, Lancaster University, UK</i>	
Supranets I	7
<i>Domenico Ferrari, Università Cattolica di Piacenza, Italy</i>	
Supranets II.....	8
<i>Luca Delgrossi, Università Cattolica di Piacenza, Italy</i>	
Crossbow A Toolkit for Integrated Services over Cell-Switched IPv6.....	8
<i>Bernhard Plattner, ETH Zürich, Switzerland</i>	
Error Control for Real-Time Audiovisual Services	9
<i>Georg Carle, Institut Eurecom, France</i>	
Co-Scheduling of Traffic Shaping and DMA Transfer in High-Performance Network Adapters... 9	
<i>Per Gunningberg; University of Uppsala, Sweden</i>	
Recording and Playback of RTP Data Streams.....	10
<i>Wieland Holfelder, Universität Mannheim, Germany</i>	
Distributed Games on the Internet.....	10
<i>Christophe Diot, INRIA, Sophia Antipolis, France</i>	
From Video Content Analysis to Object-oriented Television	11
<i>Wolfgang Effelsberg, Universität Mannheim, Germany</i>	
Implementation of an Internet Video Conferencing Application over ATM	11
<i>Heiner Stüttgen, IBM ENC, Heidelberg, Germany</i>	
Integration of Wireless LANs in the Internet.....	12
<i>Adam Wolisz, TU Berlin and GMD FOKUS, Germany</i>	
A TINA-Based Framework for Collaboration Services	12
<i>Eckhart Körner, Université de Liège, Belgium</i>	

Long-Term Popularity of On-Demand Movies	13
<i>Carsten Griwodz, TU Darmstadt, Germany</i>	
Global ATM Networks Technical Issues for Multimedia Applications	13
<i>Andres Albanese, ICSI, Berkeley, USA</i>	
Protocol Mechanisms for Native ATM Transport Services	14
<i>Raschid Karabek, RWTH Aachen, Germany</i>	
The ATMification of Ethernet.....	14
<i>Peter Martini, Universität Bonn, Germany</i>	
Performance Evaluation of Connectionless Overlay Networks for ATM.....	15
<i>Otto Spaniol, RWTH Aachen, Germany</i>	
Renegotiated ABR Service for VBR Sources.....	16
<i>Serge Fdida, Université de Paris VI, France</i>	
A Comparison of ABR and UBR to Support TCP Traffic	17
<i>Sam Manthorpe, EPFL Lausanne, Switzerland</i>	
A Simulation Study of TCP with the Proposed GFR Service Category.....	17
<i>André Danthine, Olivier Bonaventure, Université de Liège, Belgium</i>	
Communication Services for Multimedia Systems	18
<i>Andreas Mauthe, Lancaster University, UK</i>	
Network Calculus Made Easy.....	18
<i>Jean-Yves le Boudec, EPFL Lausanne, Switzerland</i>	

A New Internet Generation

Guy Pujolle, Université de Versailles, France

The ITU-T and the ATM Forum are defining new ATM bearer capabilities and associate traffic control mechanisms. At the same time Internet is changing to take care of multimedia applications using new protocols with reservation schemes. The most important fact seems to be the IP interface that now was chosen by almost all applications. A good trade-off is IP-switching that permits the network to look like an ATM network for large flows and to look like a TCP/IP network for short connections. The proposal described in the talk is to switch IP packets "à la ATM". The size of IP packets is fixed, and the tag to allow switching is supported by the flow label. As a summary the proposal is to use an ATM network but to replace the ATM cell by an IPv6 packet. IPv6 packets are transported through SONET/SDH links.

A full paper has appeared in the proceedings of this workshop: A. Danthine, O. Spaniol, W. Effelsberg, D. Ferrari (Eds.): High-Performance Networks for Multimedia Applications, Kluwer Academic Publishers, Boston, Dordrecht, London, 1999.

The Importance of Flexibility in the Design of Network and Transport Layer Protocols

Stephen Pink, SICS and Lulea University, Sweden

Network and transport protocols should be designed with flexibility over a number of dimensions. Networks are not only becoming faster, but a wider dynamic range of speed will happen as many new low-speed wireless networks will provide mobile services in the future. There should, as far as possible, be one set of network and transport protocols for all of these networks to satisfy the demands of internetworking. The NP++ protocol has been defined to provide service over a wide range of speeds, hardware architectures, and error conditions. The header of NP++ is adjustable to the speed of the network since its fields can be sent at varying frequencies. NP++ can also fragment and reassemble in the network so that it can be both cell-switched and variable-length-packet-routed, depending on available hardware. This talk shows how the principles of NP++ are applied to IPv6 to provide header compression over low speed links. Also, a more flexible UDP is presented that provides variable length checksumming to cover both real-time multimedia users and traditional data users of UDP.

Active Networking using Java and IPv6

Doug Shepherd, Lancaster University, UK

Current networks can be thought of as simple passive transporters of bits from one end-system to another via intermediate nodes within the network. The computational power is concentrated at the end-systems; within the network computation is restricted to header processing and signaling. The active network (AN) concept perceives the evolution of the entire network – hosts and intermediate nodes – into a more general computation engine. One in which not only data is transmitted but also code which can be executed on the routers and switches inside the network. This code may be used to transform the data as it travels to its destination, for example to filter down a stream for transmission over a lower bandwidth link. Users could quite literally program the network tailoring it to suit their own purposes. There already exist examples of computation inside the network: network firewalls, WWW proxies, multicast routers, etc, showing that there is a USER PULL in the market. Furthermore, there is a TECHNOLOGY PUSH seen on the increasing development of active technologies, such as Java, that could be used to develop an Active Network. This talk outlines one approach to active networking and shows how the facilities provided by the Java language and the new Internet protocol IPv6 can be used in its implementation.

Supranets I

Domenico Ferrari, Università Cattolica di Piacenza, Italy

The networks of the future will have to offer the most important features of human society in order become a true basis for such society. Among the features that we feel should be reproduced is the ability to form groups and to support the requirements of those groups. These requirements have been categorized into the following six types: membership, topology, capacity, security, connectivity, multicast. We propose a toolkit available to all users of a network such as the Internet that will allow any user to write and manage a virtual network satisfying the requirements of these six types specified by this user. The virtual network will be created on top of (Latin: "supra") this physical network; for this reason, we called it a *supranet*. The design of a supranet architecture, which is based on its own components, address space, and layer (the SN layer, which we put between the network and the transport layer), plus all the necessary mechanisms for security, connectivity and multicast, shows that the toolkit is feasible. We propose that a virtual network service based on such a toolkit is offered by future networks, as its applications are important and very numerous.

A full paper has appeared in the proceedings of this workshop: A. Danthine, O. Spaniol, W. Effelsberg, D. Ferrari (Eds.): High-Performance Networks for Multimedia Applications, Kluwer Academic Publishers, Boston, Dordrecht, London, 1999.

Supranets II

Luca Delgrosso, Università Cattolica di Piacenza, Italy

In my opinion, the networks of the future – and the Internet is no exception – will need to provide new types of services, and among these, adequate support for group communication. The notion of *supranet*, which is at the basis of our current research efforts at CRATOS, is a first attempt to introduce a higher degree of control and security into the existing network architectures. We feel that, by doing this, we will allow many groups of users to reproduce their own collaboration environment and to tailor it to their needs. In other words, we would like to give the users the possibility to define, create, manage, and delete their own virtual networks built on top of a physical one. Today, supranets are only a new idea, but so far many discussions with colleagues and students demonstrated that there is a large number of potential applications. Also presenting this idea at the Dagstuhl seminar has encouraged us to continue our efforts, and has provided new perspectives to be worked on. In particular, by introducing appropriate security mechanisms, supranets will be able to offer message integrity, sender authentication, confidentiality, and access control both to the network itself and to the resources.

A full paper has appeared in the proceedings of this workshop: A. Danthine, O. Spaniol, W. Effelsberg, D. Ferrari (Eds.): High-Performance Networks for Multimedia Applications, Kluwer Academic Publishers, Boston, Dordrecht, London, 1999.

Crossbow

A Toolkit for Integrated Services over Cell-Switched IPv6

Bernhard Plattner, ETH Zürich, Switzerland

The next generation Internet protocols will play a major part in the deployment of the future global information infrastructure. While some of the building blocks of the future Internet such as the Internet Protocol Version 6 (IPv6) have been defined and tested others are still in the research domain and will need a testbed for experimentation. This especially applies to the protocols enabling the Internet to provide integrated services. The *Crossbow* project is a joint project between ETH Zürich and Washington University in St. Louis (USA). It provides a flexible framework to investigate services and mechanisms, including resource management and packet scheduling for multimedia/multicast applications. It implements the IPv6 protocol suite in a toolkit environment, enabling the user to plug in experimental versions of comparison protocols for evaluation. It uses a Berkeley NETBSD UNIX kernel and an industry standard PC as its platform, and may be configured to serve as an end-system or a router. The *Crossbow* system will be used for research in various areas: specifically we plan to experiment with IPv6 on top of ATM, and we hope to be able to demonstrate synergies between the two technologies with integrated services capabilities.

Error Control for Real-Time Audiovisual Services

Georg Carle, Institut Eurecom, France

One approach for the provision of real-time audio-visual services is the use of cheap network services in combination with strong error control mechanisms in the end systems. Cheap network services show relatively high loss rates as well as wide delay variations, which may lead do additional errors due to late delivery. Existing applications focus on recovery of lost packets by forward error correction (FEC). This presentation shows how retransmission-based error control can be applied advantageously for real-time audio-visual services, and that retransmission of parity allows significant performance improvements in multicast scenarios. The Real-Time Multicast Protocol (RTMCP) is presented. It has a hybrid error control scheme and comes with guidelines for the selection and dimensioning of protocol parameters. For networks which support different priorities, the protocol allows not only the recovery of corrupted or lost packet but also supports the novel approach of recovering from errors of packets delivered past the deadline by transmitting redundancy with higher priority, and also by re-transmitting data and redundancy with higher priority. Applying a charging scheme with higher prices for high-priority packets is proposed in order to limit the amount of high-priority traffic.

o-Scheduling of Traffic Shaping and DMA Transfer in High-Performance Network Adapters

Per Gunningberg, University of Uppsala, Sweden

The presentation discusses how to combine traffic shaping, scheduling of a fixed sized packet and DMA transfer from host memory to a network adapter. The proposed adapter is based on a state-of-the-art CPU. By coordinating these scheduling activities less buffer memory is needed, bandwidth can be used more efficiently, DMA latency is reduced and the overall behavior becomes more predictable. The scheduling algorithm is implemented on a Pentium Pro, Ultra SPARC and DEC Alpha. Each connection is individually shaped according to its traffic contract, and its state is stored in memory. The algorithm will, for each ATM cell sent, calculate the next time a cell of this connection should be send according to the shaping, insert this time value into an ordered ready-list, and possibly initiate a DMA transfer from host memory to the adapter. The ready-list used is a binary heap since it has good scalability. Measurement results show that for 100 connections the algorithm takes about 250 CPU cycles on a Pentium Pro and 600 cycles for 64 connections. The algorithm maintains fairness between connections in overload situations.

A full paper has appeared in the proceedings of this workshop: A. Danthine, O. Spaniol, W. Effelsberg, D. Ferrari (Eds.): High-Performance Networks for Multimedia Applications, Kluwer Academic Publishers, Boston, Dordrecht, London, 1999.

Recording and Playback of RTP Data Streams

Wieland Holfelder, Universität Mannheim, Germany

Videoconferencing and shared applications are becoming more and more integrated into our day to day environment. RTP, the Real Time Transport Protocol is the protocol that is mostly used in other Internet environment to exchange data with real-time properties. In addition to exchanging the data a recording and playback facility for RTP data streams is desirable. Since RTP provides a good level of abstraction, recording and playback can be done at the level of RTP packets, even without knowing much about the payload. All information needed for synchronization and identification is provided on the RTP level. Besides the technical issues such as synchronization, indexing, efficient storage etc. we believe that other issues such as fairness and security have to be solved as well. In particular, a recorder should make itself known so other participants know that they will be recorded, and former participants that are being played back should be marked as such.

Distributed Games on the Internet

Christophe Diot, INRIA, Sophia Antipolis, France

There is not that much to say on distributed games except that it is expected that they might be the most important traffic on the Internet in five to ten years from now. Another point of interest is that behind distributed games, there are a lot of very exciting applications such as distributed simulation, air traffic control, and real-time cooperative applications. This could be enough justification to try to understand how these applications could work on the Internet. But there are more reasons: First, we don't know how these applications will behave on a network. They carry a completely new type of traffic (interactive actions) that has very strong temporal constraints. This traffic is very different from network traffic known today. Second, there are plenty of nice problems to resolve, mostly linked to the size of the groups of participants, and to synchronization between flows and between participants. If you're not convinced, please download MiMaze from our Web site (<http://www.inria.fr/rodeo/MiMaze/>), play with us from where you are, and join this new exciting research area.

From Video Content Analysis to Object-oriented Television

Wolfgang Effelsberg, Universität Mannheim, Germany

The first generation of multimedia systems was only able to route digital audio and video streams over the network and through the end systems; apart from compression and decompression the contents of the streams were not processed. Today our computers are powerful enough to analyze the contents of the streams. At the University of Mannheim we have built a system called MoCA (Movie Content Analysis) that derives 25 basic (syntactic) parameters of video and audio streams automatically. Examples include motion vectors, color histograms, edges of objects in images, and audio frequencies and amplitudes. Based on these syntactic parameters we compute semantic characteristics for a video, such as cut frequencies, object motion, camera motion (panning, zooming), silence vs. speech vs. music etc. These can be used to classify movies into genres, to produce video abstracts automatically, to identify commercials in a video etc. Whereas traditional television systems have transmitted "flat" signals (either analog or digital) from the sender to the receiver, we will soon be able to create models of objects with much more semantics on the sender's side, and to develop compression and transmission applications, such as *really* interactive television and immersive telepresence.

Implementation of an Internet Video Conferencing Application over ATM

Heiner Stüttgen, IBM ENC, Heidelberg, Germany

We describe a complete implementation of a video conferencing application taking advantage of the Quality of Service (QoS) capabilities provided by ATM. RSVP has been used to exchange the QoS parameters between the peers. The RSVP QoS parameters are mapped to ATM parameters and dedicated ATM VCs with the corresponding ATM parameters are established by the underlying classical IP over ATM implementation. The implementation is based on existing components like VIC, RSVP and classical IP over ATM. The presentation explains the required modifications of all these components to support QoS communication, and the adaptation for their integration.

Integration of Wireless LANs in the Internet

Adam Wolisz, TU Berlin and GMD FOKUS, Germany

Looking from the transport protocol point of view, two main problems have been identified as obstacles in connecting wireless LANs as Internet subnetworks:

- the excessive packet loss rate
- service disruptions because of hand-over between the picoells.

Based on measurements in a real environment using standard WLAN (Wavelan) Technology we demonstrated that because of the active immediate retransmission function within the MAC protocol the packet error rate as seen on the IP level is *not* higher than in fixed networks. On the other hand, because of this MAC level retransmission, there is a significant, highly dynamic, variability of packet *delay*. So it is packet delay variability, not packet loss, the transport protocols (say TCP modifications) have to cope with! This observation pertains both to the recently proliferated wireless LANs and emerging standards (IEEE802.11, HIPERLAN). As for hand-over minimization upon changing the cell, this is a problem of increasing significance, as the size of the cells tends to decrease. We advocate an approach based on the concept of distribution system introduced, but not specified with IEEE 802.11. Our main idea is to *multicast* data addressed to the mobile to several base stations in his neighborhood This data, invalidated after a pre-selected delay, waits for the mobile. Minimal hand-over time and no loss (but possibly duplication of data) is provided.

Credits are due to my Ph.D students Berthold Rathke and Jost Weinmüller.

A TINA-Based Framework for Collaboration Services

Eckhart Körner, Université de Liège, Belgium

The TINA Service Architecture supports the provisioning of multi-media and multi-party services. As part of the TANGRAM project at GMD FOKUS (Berlin) a TINA multimedia communication service has been designed and implemented. It offers the means to end-users to exchange audio-visual streams according to various topological communication forms. In this talk the evolution path of this service to a sophisticated collaboration service is described. Three design patterns are presented that support the coordination among a group integrity criteria. Thanks to the frameworks approach taken in TANGRAM the additional design patterns can be introduced into the service very rapidly. The voting and group integrity criteria components are also rarely found in current collaboration services.

Long-Term Popularity of On-Demand Movies

Carsten Griwodz, TU Darmstadt, Germany

In wide-area environments, the distribution of large data objects will, even with increasing bandwidth in the networks, continue to take advantage of caching and pre-distribution along the distribution path. We have started our work on effective mechanisms for this with a division of possible application types into groups. One of them are video-on-demand systems on a nation-wide level. Studies of existing literature showed quickly that no applicable user model has been defined yet. In this talk, we look at existing models and compare them with real-world data from video rental shops. The result of these studies is a set of arguments against the use of these models for wide-area or long-term video-on-demand systems. To overcome this lack of a good model, we create a new one based on the real-world data, which is presented in this talk. This will be the basis for further investigation of distribution algorithms.

A full paper has appeared in the proceedings of this workshop: A. Danthine, O. Spaniol, W. Effelsberg, D. Ferrari (Eds.): High-Performance Networks for Multimedia Applications, Kluwer Academic Publishers, Boston, Dordrecht, London, 1999.

Global ATM Networks Technical Issues for Multimedia Applications

Andres Albanese, ICSI, Berkeley, USA

The Internet has evolved into a Global Area Network driven by multiple applications that have enabled users to access information worldwide, and to reach higher levels of productivity than those achieved over local and metropolitan area networks. This global connectivity is also addressed by the deployment of ATM (Asynchronous Transfer Mode) services, widely supported by the telecommunication and computer industries to provide QoS to the Internet. Experimental results have demonstrated the feasibility of collaborative applications on a global scale, and have shown the need for solutions to the following technical problems:

1. round trip delays of up to 200 milliseconds,
2. multiple management domains,
3. different time zones, and
4. heterogeneity of end-stations, access speeds, and access services.

Proposed solutions are being implemented to address these technical issues. Examples are diagnostic tools that collaborate at finding the fault in the network, a video conference recorder to record and delay the play of the session, distributed security encryption protocols, selective forward error correction using Priority Encoding Transmission to enable video and audio transmission over lossy networks. And an *intermedia service* to translate messages and comply with specific service profiles of senders and receivers.

Protocol Mechanisms for Native ATM Transport Services

Raschid Karabek, RWTH Aachen, Germany

In the previous presentations of this workshop IP seemed to be much more of an issue than ATM. But in this presentation we take a look at purely homogenous ATM networks. Examples for such networks range from small special-purpose networks to factory environments or even a Global ATM Network. The main objective of this presentation is to identify and implement suitable mechanisms on the transport protocol level for reliable message/data-transfer in ATM networks. These mechanisms include Flow Control, Error Control and Congestion Control (for ATM-UBR channels). The objective of Flow Control in native ATM systems is not only to protect the destination station from being overloaded by the sender, but also to protect the sending station's ATM driver from being flooded by applications sending at a faster rate than the corresponding ATM channel's peak cell rate. A suitable back-pressure scheme is introduced, and related implementation issues are discussed. Error Control in Native ATM networks should take advantage of ATM's property of "In-sequence-delivery". This enables the receiver to detect packet loss as soon as a subsequent packet is received. A suitable Error Control is presented. An end-to-end Congestion Control mechanism is necessary to support data transfer via ATM's "best effort" service, UBR (Unspecified Bit Rate). Our current work on such a mechanism based on *Persistent Rate Adaptation* is introduced.

A full paper has appeared in the proceedings of this workshop: A. Danthine, O. Spaniol, W. Effelsberg, D. Ferrari (Eds.): High-Performance Networks for Multimedia Applications, Kluwer Academic Publishers, Boston, Dordrecht, London, 1999.

The ATMification of Ethernet

Peter Martini, Universität Bonn, Germany

The presentation started with a discussion of commercial issues of ATM networks, addressing both the residential and the business area. It presented results from field trials in the AMUSE project which is an ACTS project providing and evaluating multimedia services to residential users. From measurement results and from questionnaires it came clear that the users were much more interested in Internet access and in browsing the WWW than they were in Video-on-Demand or News-on-Demand. Even the most interested users are not willing to pay more for these multimedia services than they do for comparable services such as newspapers or video rental today. A comparison with today's cost of ATM service clearly shows that it will (at least) take a long time to make B-ISDN universal or near-universal. Lacking the short-term perspective of universal end-to-end ATM service it seems almost impossible for ATM to replace Ethernet in the local area. To make it even harder new activities of the "Ethernet camp" now address all areas where ATM was supposed to offer additional benefits: higher speeds (100 Mbit/s, 1 Gbit/s ...), multimedia support

by priority-based traffic expediting, and full-duplex transmission with flow control are important milestones discussed in our presentation.

Performance Evaluation of Connectionless Overlay Networks for ATM

Otto Spaniol, RWTH Aachen, Germany

ATM networks (which have been the dream of everybody but which appear to have an unclear future seeing all the competitors such as Gigabit Ethernet etc.) have been proposed for the interconnection of different local area networks (LANs) over a wide-area range. Several possibilities and problems for that came from the fact that ATM is connection-oriented whereas the vast majority of LANs is connectionless. In our work (jointly done with Marko Schuba) we have studied and evaluated the *direct connectionless server approach* which is one out of at least five different alternatives. The principle works as follows:

- a. Place "some" connectionless servers (CLS) into the ATM network. A CLS is a "normal" node plus some additional functionality.
- b. Assign a LAN interworking unit to a "home CLS"
- c. Construct an overlay network which interconnects the CLS's to each other.

There are plenty of different possibilities for CLS placement and their interconnection. Different alternatives can be composed by objective functions such as "minimum average delay" or "load balance" or or or ...Evaluation was done by Monte Carlo simulation of randomly constructed ATM graphs. The overlay network was assumed to be

- a full mesh
- a ring
- a hierarchy of nodes where fully meshed subgroups are interconnected.

The best topology was found with three different techniques:

- full enumeration (only possible for small networks),
- optimal extension heuristic,
- optimal extension heuristic plus simulated annealing.

The hierarchical approach turned out to be best since it lead to almost optimal results, and since it could be evaluated in a rather short time.

A full paper has appeared in the proceedings of this workshop: A. Danthine, O. Spaniol, W. Effelsberg, D. Ferrari (Eds.): High-Performance Networks for Multimedia Applications, Kluwer Academic Publishers, Boston, Dordrecht, London, 1999.

Renegotiated ABR Service for VBR Sources

Serge Fdida, Université de Paris VI, France

We discuss two problems related to network services and traffic management for variable bit rate sources. We argue that most applications will require a service that delivers a minimum cell rate (MCR) and a fair-share (FS) of the bandwidth left unused in the network. The FS part can be considered as providing a best-effort service. These characteristics fit very well with the Available Bit Rate (ABR) service definition of the ATM Forum (or even GFR of the IETF). Therefore, we first design an efficient ABR algorithm named ERAQLES. It is original because it uses a novel distribution to control the fitting level of the ABR queue. The inclusion of the ABR queue into the algorithm provides a significant statistical gain. Major properties such as fairness and convergence have been demonstrated mathematically. Simulation results show that the algorithm outperforms the other existing solutions, can handle both ER and Relative Marking switches as well as VBR traffic. The second problem deals with the efficient way to handle VBR sources. VBR service is not satisfactory because of the complexity of the traffic descriptions. Therefore, we consider solutions based on renegotiations when the source traffic is smoothed and considered as a step function. Solutions based on renegotiated CBR (R.CBR) have been studied and showed to be either not efficient or to require too many renegotiations (thus increasing signaling). We explore a novel solution based on the concept of renegotiated ABR (R.ABR) where a renegotiation triggers either an increase or a decrease of the MCR. We infer that the probability that a renegotiation gets rejected is lower than in the R.CBR case. This is because of the fair-share bandwidth that provides a safety margin for renegotiated higher MCR. This R.ABR looks very promising, and will be assessed in more detail in the future.

A Comparison of ABR and UBR to Support TCP Traffic

Sam Manthorpe, EPFL Lausanne, Switzerland

We compare the performance of ABR and UBR for providing high-speed network interconnection with TCP traffic. We test the hypothesis that UBR with adequate buffering in the ATM switches results in better overall "goodput" of TCP traffic than explicit rate ABR for LAN interconnection. We find that this is so in a wide range of scenarios. We identify four phenomena that may deteriorate ABR performance, and test whether each of these has an impact on TCP "goodput". This reveals that the extra delay incurred in the ABR end-system and the overhead of RM cells account for the difference in ABR and UBR performance. We test whether it is better to use ABR to push congestion to the end-systems in a parking-lot scenario, or whether we can allow congestion to occur in the network. Finally, we test whether the presence of a "multiplexing loop" causes performance degradation for ABR and UBR. We find our original hypothesis to be true in all cases. We observe, however, that ABR is able to improve performance when the buffering inside the ABR part of the ATM network is small compared to that available at the ABR end-systems. We also observe how ABR can enforce fairness in the network.

A full paper has appeared in the proceedings of this workshop: A. Danthine, O. Spaniol, W. Effelsberg, D. Ferrari (Eds.): High-Performance Networks for Multimedia Applications, Kluwer Academic Publishers, Boston, Dordrecht, London, 1999.

A Simulation Study of TCP with the Proposed GFR Service Category

André Danthine, Olivier Bonaventure, Université de Liège, Belgium

Recently, the GFR (Guaranteed Frame Rate) service category has been proposed to provide bandwidth guarantees in a simpler way than with the ABR service category. Besides the MCR (Minimum Cell Rate) GFR recognizes the limits of the AAL 5 PDV as a unit for tagging and discarding. In this work we first present a study of a FIFO-based implementation, and then the way this service category interacts with TCP/IP in a LAN environment. Our simulations show that with the FIFO-based implementation of the GFR service category TCP is unable to benefit from the bandwidth guarantees. An exploration of the parameter space associated with TCP behavior indicates that no acceptable mode of operation can be found. With the proposed WFQ-based implementation (Weighted Fair Queuing), the performance of TCP is excellent when no losses occur but quickly degrades if tagging is used inside the network. Our work is likely to influence a modification of the specification of the GFR service category.

A full paper has appeared in the proceedings of this workshop: A. Danthine, O. Spaniol, W. Effelsberg, D. Ferrari (Eds.): High-Performance Networks for Multimedia Applications, Kluwer Academic Publishers, Boston, Dordrecht, London, 1999.

Communication Services for Multimedia Systems

Andreas Mauthe, Lancaster University, UK

In the past it was not always necessary to clearly distinguish between service and protocol since usually there was a straight one-to-one mapping. However, with emerging multimedia communications, this is now no longer true. In this context the intrinsic value of a service is its ability to provide a stable and meaningful interface to the service user while allowing the optimization of the underlying communication system. At Lancaster University a service architecture called *GCommS* and end-to-end communication services have been developed considering the special requirements of interactive multimedia systems. Two transport services, the multicast *M-Connection Service* and the n-plex *N2N Connection Service*, have been specified to provide support for interactive multimedia systems. The services are fully specified with a restricted set of service elements that allow to state QoS requirements, integrity and reliability conditions. A modular, configurable protocol architecture has been developed using existing protocols whenever possible. The design of the system is object-oriented with different objects for certain media types and context conditions. The current implementation concentrates on the Internet protocol suite; however, the service can also be provided over other protocols and communication architectures such as XTP 4.0 or ATM.

A full paper has appeared in the proceedings of this workshop: A. Danthine, O. Spaniol, W. Effelsberg, D. Ferrari (Eds.): High-Performance Networks for Multimedia Applications, Kluwer Academic Publishers, Boston, Dordrecht, London, 1999.

Network Calculus Made Easy

Jean-Yves le Boudec, EPFL Lausanne, Switzerland

Network Calculus is a set of computational rules and methods that can be used to analyze deterministic queuing systems. It is based as a non-standard linear systems theory, called min-plus algebra. We give a simple, though in-depth treatment of the fundamental concepts. Then we present applications to the derivation of complex bounds used in the Internet integrated service model. We give examples for shapers, deterministic effective bandwidth, equivalent capacity, convexity of call acceptance regions, optimal parameters of a variable bit rate trunk, and window flow control.

A full paper has appeared in the proceedings of this workshop: A. Danthine, O. Spaniol, W. Effelsberg, D. Ferrari (Eds.): High-Performance Networks for Multimedia Applications, Kluwer Academic Publishers, Boston, Dordrecht, London, 1999.