

Wolfgang Effelsberg, Domenico Ferrari,
Otto Spaniol, André Danthine (editors):

**Architecture and Protocols for High
Performance Networks**

Dagstuhl-Seminar-Report; 118
19.06.-23.06.95 (9525)

ISSN 0940-1121

Copyright © 1995 by IBFI GmbH, Schloss Dagstuhl, D-66687 Wadern, Germany
Tel.: +49 - 6871 - 2458, Fax: +49 - 6871 - 5942

Das Internationale Begegnungs- und Forschungszentrum für Informatik (IBFI) ist eine gemeinnützige GmbH. Sie veranstaltet regelmäßig wissenschaftliche Seminare, welche nach Antrag der Tagungsleiter und Begutachtung durch das wissenschaftliche Direktorium mit persönlich eingeladenen Gästen durchgeführt werden.

Verantwortlich für das Programm ist das Wissenschaftliche Direktorium:

Prof. Dr. Thomas Beth,
Prof. Dr. Oswald Drobnik,
Prof. Dr.-Ing. José Encarnação,
Prof. Dr. Hans Hagen,
Dr. Michael Laska,
Prof. Dr. Thomas Lengauer,
Prof. Dr. Christoph Meinel,
Prof. Dr. Wolfgang Thomas,
Prof. Dr. Reinhard Wilhelm (wissenschaftlicher Direktor)

Gesellschafter: Gesellschaft für Informatik e.V., Bonn,
Universität des Saarlandes,
Universität Frankfurt,
Universität Kaiserslautern,
Universität Karlsruhe,
Universität Stuttgart,
Universität Trier,
TH Darmstadt

Träger: Die Bundesländer Saarland und Rheinland-Pfalz

Bezugsadresse: Geschäftsstelle Schloss Dagstuhl
Universität des Saarlandes
Postfach 15 11 50
D-66041 Saarbrücken, Germany
Tel.: +49 -681 - 302 4396
Fax: +49 -681 - 302 4397
e-mail: office@dag.uni-sb.de
url: <http://www.dag.uni-sb.de>

Preface

From June 19 to 23, 1995, leading international researchers on computer networks met at Schloß Dagstuhl to discuss Architecture and Protocols for High Performance Networks. It was the second workshop on this topic at the castle; the first one having taken place in September 1993.

The purpose of the workshop was to bring together telecommunication engineers, computer communication experts and multimedia application designers, three groups of people who not very often talk to each other. All of the participants were experts in their fields, from many different countries, and the presentations and discussions were at a very high level.

Although a wide range of problems was covered at the workshop, most of the contributions centered around two major issues: ATM networking, and quality-of-service for multimedia applications. It is expected that most of the bandwidth in future high-performance networks will be used up by multimedia applications, transmitting digital audio and video. 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 and many similar issues, and proposed interesting and innovative solutions.

This report is a collection of abstracts of all the presentations 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 high-quality papers will appear under the title "High-Speed Networking and Multimedia Applications", edited by the organizers of the workshop and to be published by Kluwer Academic Publishers in December 1995.

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, the Dagstuhl team provided us with an excellent environment: *this place is optimized for thinking*.

Wolfgang Effelsberg, University of Mannheim, Germany

Otto Spaniol, RWTH Aachen, Germany

André Danthine, Université de Liège, Belgium

Domenico Ferrari, International Computer Science Institute, Berkeley, California, USA

Contents

PER GUNNINGBERG	
<i>ATM as A Memory Interconnect in A Desk Area Network</i>	9
DON TOWSLEY	
<i>Use of Multiprocessors in End System Protocol Processing</i>	9
MARTINA ZITTERBART	
<i>Towards Integrated QoS Management</i>	10
ANDREAS POMBORTSIS	
<i>Some Thoughts about the Techno-Economical Problems of an Integrated Computing and Communication Environment</i>	10
JON CROWCROFT	
<i>Policy and Mechanism to Support Multiple IP or ATM Bearer Service Classes</i>	10
OLIVER BONAVENTURE	
<i>Transport Protocols and Classes of Service in ATM Networks</i>	11
NICOLAS GEORGANAS	
<i>Definitions, Merging, Splitting and Cell-Loss Bounds of Self-Similar Traffic on ATM Queues</i>	11
HARMEN VAN AS	
<i>Photonic Communication Networks: A Universal Infrastructure for the Information Society</i>	11
JEAN-YVES LE BOUDEC	
<i>SCONE" Scalability Enhancements for Connection-Oriented Networks</i>	12
ZIGMANTAS L. BUDRIKIS	
<i>Cell Transfer Protocol (CTP) - Vehicle for Support of Unresourced Bearer Capabilities on ATM</i>	12
ANDRÉ DANTHINE	
<i>Is ATM a Continuity or a Discontinuity for the LAN Environment?</i>	12
MICHAEL PATERAKIS	
<i>Protocols for Voice/Data Integrated Channel Access in Microcellular Wireless Digital Networks</i>	13
AUREL A. LAZAR	
<i>Xbind: A Binding Architecture for Multimedia Networks</i>	13
DOMENICO FERRARI	
<i>Service Models for Integrated Services Networks</i>	14

FOUAD TOBAGI	
<i>Constant Quality Video Encoding and Evaluation of Networks Carrying Multimedia Traffic</i>	14
ERNST BIRSACK	
<i>Scaleable Video Server Architectures</i>	14
DOUG SHEPHERD	
<i>Support of Multimedia Servers in Heterogeneous Environments</i>	15
WOLFGANG EFFELSBERG	
<i>Chaos with QoS</i>	15
RAJ YAVATKAR	
<i>An End-System Architecture for QoS Guarantees over High-Speed Networks</i>	15
GURUDATTA PARULKAR	
<i>QoS Guarantees within Endsystms</i>	16
AHMED TANTAWY	
<i>Building End-to-End Multimedia Systems</i>	16
RALF STEINMETZ,	
<i>Reservation of Resources in Advance</i>	16
HEINRICH STÜTTGEN	
<i>Resource Allocation vs. Traffic Adaptation in Multimedia Communication</i>	17
GEORG CARLE	
<i>Towards High Performance Group Communication Services in ATM Networks</i>	17
SERGE FDIDA	
<i>Multicast Transport Protocol for Multimedia Cooperative Applications</i>	17
ANDRES ALBANESE	
<i>Priority Encoding Transmission (PET)</i>	18
PETER MARTINI	
<i>Real-Time Communication in IEEE 802 High Speed LANs & MANs</i>	18
PHUOC TRAN-GIA	
<i>Performance Aspects of Video Transmission in ATM Systems</i>	18
DEMETRES KOUVATSOS	
<i>Performance Modeling and Analysis of ATM Networks</i>	19
AUGUST CASACA	
<i>An ATM-based Customer Premises Network for the Residential Environment</i>	19
OTTO SPANIOL	
<i>High Speed and Mobility</i>	19

STEPHEN PINK	
<i>A Flexible Network Protocol for High Speed and Wireless Links</i>	20
ADAM WOLISZ	
<i>Wireless Access to High Speed Networks</i>	20
BENOIT MACQ	
<i>Conditional Access and Copyright Protection</i>	20

Abstracts of the Presentations

Per Gunningberg

Dept. of Computer Systems, University of Uppsala, Uppsala, Sweden

ATM as A Memory Interconnect in A Desk Area Network

ATM has been successfully used in Wide Area Networks and Local Area Networks. A possible next step in this evolution is to use ATM for memory interconnect. This presentation discusses how ATM can fulfill the functional and performance requirements and the implications this will have for external ATM communication. We conclude that ATM can meet the throughput requirements, but will have problems with the latency for the transfer of small cache lines. This fixed 48 byte payload in ATM results in inefficient use of bandwidth. Existing ATM adaptation layers do not have sufficient functionality and must be extended. Using ATM internally has a limited synergy effect when connected to ATM based WANs.

Don Towsley

Computer Science Department, University of Massachusetts, Amherst, MA, U.S.A.

Use of Multiprocessors in End System Protocol Processing

This is an overview of lessons learned in the use of multiprocessors for high performance protocol processing. Two approaches are packet level parallelism (PLP) and connection level parallelism (CLP). Under the first, threads shepherd packets up and down the protocol stacks concurrently. The second does the same but restricts at most one packet per connection. Lessons learned include:

- 1) PLP has a difficulty in dealing with protocols such as TCP that provide in-order semantics
- 2) CLP has difficulty providing equal share of bandwidth to connections due to memory reference behavior of connection
- 3) For both PLP and CLP it is important to account for cache affinity when scheduling resources such as processors, trends, buffers.

Martina Zitterbart

*Institute of Operating Systems and Computer Networks
Technical University of Braunschweig, Braunschweig, Germany*

Towards Integrated QoS Management

Decomposition of communication into generic building blocks: protocol subsystem, operating system and network. Each block is comprised of resources and resource managers. A QoS manager takes care about the interaction among the blocks. Specific tasks of the QoS manager are QoS mapping among the service interfaces of the building blocks and QoS control at the provider service interfaces. This way, different types of services, such as guaranteed, statistical and best-effort services can be provided.

Additionally, a so-called Transport Gateway was presented which is in charge of the interconnection among wireless networks and fixed networks. It consists of a Wireless Agent which splits the end-to-end connection into two parts: across the wireless part and the fixed part. This allows to select individually appropriate error and flow control mechanisms and thus, to achieve a good overall performance.

Andreas Pombortsis

Department of Computer Science, Aristotle University, Thessaloniki, Greece

Some Thoughts about the Techno-Economical Problems of an Integrated Computing and Communication Environment

The presentation tries to examine a small part of the spectrum of choices available to network users and providers. In this framework it starts with a typical CAN and examines the problem of forecasting the demand for broadband telecommunication services, the possibility for global networking with ISDN and the migration alternatives to the integrating and interacting of services based on existing technologies. The basic results are as follows: Advances in existing technologies may extend the life cycle of existing networks and services and slow the acceptance of new technologies. Probably ATM networks will coexist with existing STM-based networks for a long time. A gradual transition from N-ISDN to ATM will support the cost reduction of ATM and will further progress towards B-ISDN.

Jon Crowcroft

Computer Science Institute, University College London, London, Great Britain

Policy and Mechanism to Support Multiple IP or ATM Bearer Service Classes

THE IETF Internet Services Group and the ITU and ATM Forum have all devised an (evolving) set of bearer services - in IP6 these are best effort, predictive and guaranteed, while in ATM these are ABR, VBR and CBR (roughly). On a given path through a network created by more than one provider, different business cases will create different policies for the mix, strengths and types of guarantees. This talk is about mechanisms to support hierarchical policies for providers, paths and reservations of different classes based on RSVP and CBQ.

Oliver Bonaventure

Institute Montefiore B28, Université de Liège, Liège, Belgium

Transport Protocols and Classes of Service in ATM Networks

This presentation looks at the currently defined service classes for ATM networks (CBR, VBR, ABR) from a transport protocol's point of view. We show that the transport protocol must be aware of the service class used by the notifying ATM connection. More specifically, some mechanisms must or must not be present in the transport layer depending on the ATM class of service used. Concerning CBR we show that spacing is really necessary in the ATM adapter, but congestion control should not be provided by the transport protocol. We also look at VBR, and particularly at the VBR traffic contract. We then look at ABR and VBR, and we conclude that ATM is not comparable to the classical LAN and that the different classes of service offer a different service to the transport layer.

Nicolas Georganas

Multimedia Communications Research Laboratory, University of Ottawa, Ottawa, Canada

Definitions, Merging, Splitting and Cell-Loss Bounds of Self-Similar Traffic on ATM Queues

Definitions and comparisons of self-similarity; merging and splitting self-similar processes; characterization of the resulting process; model of self-similar process constructed from "on-off" sources; derivation of lower bounds for the buffer overflow probability and cell-loss probability of self-similar traffic in a finite ATM buffer; derivation of cell-delay distribution.

Harmen van As

IBM Research Division, Zurich Research Laboratory, Rüschlikon, Switzerland

Photonic Communication Networks: A Universal Infrastructure for the Information Society

Continuous technological progress, ever-increasing demand, new services and applications as well as severe economic pressure caused the creation of various stages of communications networking. With photonic networks a new communication era appears at the horizon, in which services can be integrated naturally and traditional networks can be included smoothly. New properties such as optical transparency, high-capacity transmission over global distances, optical by-pass switching and dynamic wavelength routing will allow to operate photonic networks in a flexible manner. In addition, the same network infrastructure can simultaneously be used for terrestrial networks for mobile communications, for road and air traffic control as well as for analog and digital television distribution.

Jean-Yves Le Boudec

Laboratoire Réseaux de Communication (LRC), EPFL, Lausanne, Switzerland

SCONE" Scalability Enhancements for Connection-Oriented Networks

We consider the issue of increasing the number of connections that connection-oriented networks, such as ATM, can support. We describe one step that aims at reducing connection awareness inside the network. To that end, connections between the same pair of access nodes are grouped together and made indistinguishable inside the network. The concept of dynamic virtual path trunk is introduced as a support mechanism, and it is shown how virtual path links can be set up and maintained without additional round trip delays.

Zigmantas L. Budrikis

*Australian Telecommunications Research Institute (ATRI), Curtin University of Technology,
Perth WA, Australia*

Cell Transfer Protocol (CTP) - Vehicle for Support of Unresourced Bearer Capabilities on ATM

Proposes creation of ATM layer bearer capability based on generic flow controlled transfer of cells. Requires Generic Flow Control (GFC) at UNI (already in place) and GFC at NNI (still to be defined). Proposes that classical sliding credit window, applicable globally to controlled traffic class, be adopted as GFC at NNI. Introduces a quota based cell transfer protocol (CTP) applicable to ATM switches and suitable for provision of the CCT bearer capability efficiently.

André Danthine

Institute Montefiore B28, Université de Liège, Liège, Belgium

Is ATM a Continuity or a Discontinuity for the LAN Environment?

The answer is that ATM is a discontinuity. A LLAN (Legacy LAN) is a shared medium with an access control but an ATM network is a shared throughput with an access control which protects the traffic of the other connections. In an ATM network a connection traffic descriptor is necessary to allocate the bandwidth. In LLAN there is no connection traffic descriptor. In an ATM network each source must comply with the traffic defined in the contract. In a LLAN there is no obligation but no guarantee of access. An ATM network may drop cells not conforming to the contract while in a LLAN the access method avoids packet loss at the access. ATM is a discontinuity even if IETF and ATM Forum are proposing solutions to allow existing applications to be ported on an ATM LAN. But multimedia requires a new approach due to the QoS and it is a pity that SAR and convergence sublayers are not separated.

Michael Paterakis

*Department of Electronic & Computer Engineering, Technical University of Crete, 731-00
Chania, Greece*

Protocols for Voice/Data Integrated Channel Access in Microcellular Wireless Digital Networks

We present the results of a study that explores the performance of two promising reservation random access (RRA) protocols for transmitting voice packets over a common radio broadcast channel in a microcellular radio environment. We examine the protocols with respect to voice transmission under ideal and adverse channel conditions. In addition, we investigate the ability of the protocols to support efficient voice/data channel access integration within the system. The examined RRA protocols clearly mark the end of the voice contention period (VCP), thereby enabling all of the terminals within the microcell to differentiate between available voice and available data slots. Resolving the contending voice packets first enforces the priority of the voice traffic. Each protocol can be combined with efficient collision resolution random access protocols for transmitting data packets. Such a voice/data integration mechanism eliminates the potential voice degradation caused by competition between voice and data terminals for available slots. The results show that the combined protocols provide stable and robust performance under adverse channel conditions and that they can be employed to sustain voice/data channel access integration under heavy system loading.

Aurel A. Lazar

Department of Electrical Engineering, Columbia University, New York, NY, U.S.A.

Xbind: A Binding Architecture for Multimedia Networks

Service creation, deployment and management is one of the key features of multimedia networking. We noted that such networks are based on open transport and management architectures. However, network control is typically closed. The xbind Project proposes opening up network control. The architecture proposed is based on an object-oriented model of the network, where both logical entities and QoS abstract are used to realize (via binding) network services.

Domenico Ferrari

Computer Science Division, University of California at Berkeley, Berkeley, CA, U.S.A.

Service Models for Integrated Services Networks

The talk reviewed and compared some service models that have been proposed for service integration within the Internet community, the ATM community, and the Tenet Group. The services were classified, to facilitate the discussion, into essential and additional, the essential being the "best-effort" services for non-real-time traffic and the "deterministic guaranteed" ones for real-time traffic. The emphasis was on what QoS parameters and what additional services (e.g., multicast, setup type, media scaling, and others, besides any "intermediate" services, i.e., those falling into the gap between the two essential extremes) have been proposed in the approaches that were reviewed. The conclusion was that the differences in the QoS parameters are very small, and those among the additional services somewhat larger but not very large; the conjecture was put forward that the latter are most explainable with the different stages of development reached by the various proposals. Thus, there is good potential for a conveyance of the Internet and ATM technologies in the area of integrated services offerings, but a large amount of experimental work remaining to be done with real applications to see whether we are on the right track.

Fouad Tobagi

Dept. of Electrical Engineering, Stanford University, Palo Alto, CA, U.S.A.

Constant Quality Video Encoding and Evaluation of Networks Carrying Multimedia Traffic

In order to achieved a given video quality at all times, the encoder parameters must be appropriately adjusted according to the scene content. We propose a video encoding scheme which maintains the quality of the encoded video at a constant level. This scheme is based on a quantitative video quality measure, and it uses a feedback control mechanism to control the parameters of the encoding. We also evaluate various networks carrying video traffic. We take into account the end-to-end delay and consider a wide range including low delay for interactive communication. We measure the effect of packet delay and thus packet loss in terms of rate and duration of glitches. We also evaluate the effect of data traffic (bursty and non-bursty) on the quality of video.

Ernst Biersack

Institut EURECOM, Sophia-Antipolis, France

Scaleable Video Server Architectures

Scalability is an important requirement. It is necessary to allow to adapt to an increase in (i) the number of videos stored and (ii) the number of clients concurrently demanding service.

We propose a video server architecture that distributes each video over all server nodes (application level striping). Our approach achieves perfect load balancing independent of the request pattern and addresses scalability.

An important aspect to resolve is reliability. In order to avoid that the failure of a server node affects the replay of the video at the clients, we propose to use redundant information. Redundancy can also be used to deal with transmission problems in the network such as data being lost or delayed.

Doug Shepherd

Computer Centre, Lancaster University, Lancaster, Great Britain

Support of Multimedia Servers in Heterogeneous Environments

This talk addresses the issues of scalability and support for heterogeneous clients within a networked multimedia storage architecture that employs server replication as a technique to increase both storage and I/O bandwidth. Two conventional load balancing techniques, file replication and network striping, are examined with the conclusion that neither are truly scalable. A hybrid load balancing technique is proposed that combines the ideas of network striping and file replication that exploits the characteristics of multi-resolution compression to divide files into a number of compressed components, each of which can be stored and replayed independently. This approach has the advantage that it also provides storage support for client heterogeneity as data stored in multiple resolutions can be combined to meet the QoS requirements of clients, hence reducing load and increasing the number of streams that can be serviced.

Wolfgang Effelsberg

Computer Science IV, University of Mannheim, Mannheim , Germany

Chaos with QoS

Current networks do not support Quality of Service concepts (QoS). There were many recent proposals from the research community how QoS parameters should be specified, and how guaranteed bandwidth, delay, jitter and loss could be provided, both from the Internet community and from the ATM standardization bodies. However, these schemes are incompatible, and many basic questions are still unsolved: How could application QoS parameters be mapped to network QoS parameters? Should the QoS interface be based on the contract model, or should applications be adaptive to the network load? Can media scaling be combined with multicast filtering? It is argued that an overall QoS architecture is needed rather than just interface parameters, and that the most appropriate model is a contract with [lower bound, upper bound] parameters.

Raj Yavatkar

Dept. of Computer Science, University of Kentucky, Lexington, KY, U.S.A.

An End-System Architecture for QoS Guarantees over High-Speed Networks

A considerable amount of work is in progress for providing QoS guarantees for network traffic. However, relatively little work has been done in the area of extending QoS guarantees to applications running on an end-system. We introduce a framework for managing end-system resources (CPU, NIA, bus/I/O bandwidth, --) in an integrated fashion. The architecture called AQUA (Adaptive Quality-of-Service Architecture) is designed to provide average rate guarantees on rate of progress by individual applications and provides mechanisms for application-level adaptation in the presence of changing resource requirements and arrival of new applications.

Gurudatta Parulkar

Dept. of Computer Science, Washington University, St. Louis, MO, U.S.A.

QoS Guarantees within Endsytms

We have introduced a framework for providing QoS guarantees for multimedia applications within the endsystems. The framework consists of four components: 1. QoS specification at a high level by applications using a small number of parameters; 2. QoS mapping which is a library (automated mechanism) that maps application QoS specifications into resource requirements; 3. QoS enforcement which ensures that an application gets its share of resources; 4. A protocol implementation model which works with QoS framework.

We have created Real-Time Signals in NetBSD OS which allows QoS enforcement in a Standard Unix OS. RT-Signals also ensure that the number of context switches is reduced.

Ahmed Tantawy

IBM T.J. Watson Research Center, Yorktown Heights, NY, U.S.A.

Building End-to-End Multimedia Systems

Most characteristics of multimedia traffic and services are still unknown. The best way to understand and quantify them in a realistic manner is to build experimental testbeds. We have built an end-to-end multimedia system, including all the major components from the server (or a real-time MPEG-2 video source) to the client (PC- or set-top-box). We have adopted a modular design approach in order to enable the comparison of different approaches to the design and implementation of the various functional components. We have focused on the use of existing cable TV plants as a delivery medium and we have identified several issues in this area.

Ralf Steinmetz,

IBM European Networking Center, Heidelberg, Germany

Reservation of Resources in Advance

Today's resource reservation schemes allot to make a reservation "from now until an undetermined point in time". However, many applications (like conferencing) and situations in which we are involved, require semantics like reservation styles as "from 3 p.m. until 4:15 p.m. every Monday". This presentation aims to provide a first application-driven framework for such a generic reservation scheme. A set of so far unsolved research issues arises; examples are: How shall a reservation be extended if the applications/users want to make use for a duration longer than the one specified? How do we cope with failure if something already reserved is not any longer "available"?

Heinrich Stüttgen

*Broadband Multimedia Communication, IBM European Networking Center
Heidelberg, Germany*

Resource Allocation vs. Traffic Adaptation in Multimedia Communication

The paradigms are used to provide traffic or flow control, namely resource allocation and traffic adaptation. The presentation compares and clarifies the existing and novel traffic control techniques according to their suitability for real-time traffic like variable rate video streams. In particular we show the delay penalty introduced by ATM traffic sharing and first results of an adaptive, loss traffic control technique called media scaling.

Georg Carle

*High Performance Networking Group, Institute of Telematics, University of Karlsruhe
Karlsruhe, Germany*

Towards High Performance Group Communication Services in ATM Networks

The efficient provision of reliable multipoint services remains a major unresolved issue. A novel framework for support of reliable group communication in ATM networks is presented, consisting of two adaptation layer protocols in combination with a Group Communication Server. The Reliable Lightweight Multicast Protocol (RLMCP) is a simple and efficient protocol for the Service Specific Convergence Sublayer of AAL 5, featuring frame-based retransmission in selective repeat and go-back-N mode. The Reliable Multicast ATM Adaptation Layer (RMC-AAL) additionally offers cell-based retransmissions and cell-based forward error correction FEC). Results of a performance evaluation are presented which allow to identify the scheme most appropriate for a given scenario.

Serge Fdida

Laboratoire MASI, Université Paris 6, Paris, France

Multicast Transport Protocol for Multimedia Cooperative Applications

Besides Throughput and Delay issues, new functionalities are required for multimedia cooperative applications. We address first the reliability semantics associated with multicast conversations and emphasize the need for a new semantic named "Statistical Reliable Multicast". This solution allows to relieve the sender from monitoring the state of each single receiver, enable to scale properly with large groups size and perform efficiently on nonbroadcast networks.

We then present the key features associated with the algorithm designed to provide a statistically reliable multicast, adding also mechanisms to cope with AGI (Active Group Integrity) requirements. Finally, issues such as the influence of the underlying network topology and the interest of a hierarchical approach (such as Group Communication Servers) are discussed.

Andres Albanese

ICSI - International Computer Science Institute, Berkeley, CA , U.S.A.

Priority Encoding Transmission (PET)

The key idea is an encoding scheme that will recover the most important information of a message no matter which packets are lost or intentionally deleted in a transmission. By the judicious use of ensure encoding in the application, PET achieves greatly improved performance, enabling statistical multiplexing of video images, and graceful acquisition or degradation. This is done without adding excessive overhead in terms of network capacity (20%) or delay (a group of pictures), and effectively decreasing transmission cost per channel. For more, tune in URL <http://www.ICSI.Berkeley.edu/PET> .

Peter Martini

*Fachbereich 17 - Mathematics and Computer Science, University of Paderborn
Paderborn, Germany*

Real-Time Communication in IEEE 802 High Speed LANs & MANs

The presentation compares the performance of the IEEE 802 protocols DQDB (IEEE 802.6) and Demand-Priority LAN (IEEE 802.12) in terms of the price to be paid by normal-priority traffic for the support of high-priority traffic. Additionally, it includes results for FDDI.

It turns out that the Demand-Priority LAN fails for large distances and/or small frames. In contrast to this, DQDB and FDDI yield excellent results for all scenarios studied in our simulations. This holds both for the average delays and for the frame loss characteristics of normal priority traffic.

Phuoc Tran-Gia

*Distributed Systems, Institute of Computer Science, University of Würzburg
Würzburg, Germany*

Performance Aspects of Video Transmission in ATM Systems

MPEG is a standard for video in ATM systems, under discussion in different bodies of standardization. The presentation discusses performance aspects like cell blocking estimation, J-frame alignments, multiplexing, etc., concerning video the transmission in ATM systems using MPEG encoding techniques. It is shown that in some cases exact queuing system analysis can be applied to investigate cell loss accurately and simple traffic characterization methods can be efficiently used to describe video traffic in ATM environments.

Demetres Kouvatso

Dept. of Computing, University of Bradford, Great Britain

Performance Modeling and Analysis of ATM Networks

The indices of dispersion for counts and intervals are widely recognized as appropriate measures for characterizing bursty, correlated ATM traffic. However, the analysis of such systems is complex and is often found to be intractable. The talk described a novel approach to study of a $GI^G/D/1/N$ queue with general correlated traffic, in that the arrival process is not modeled upon any particular type of traffic source. The analysis exploits properties of the batch renewal process, which is the least biased choice given only measures of indices of dispersion which might be expected to result from actual traffic measurement. Closed form expressions for basic performance distributions, such as queue length and waiting-time distributions as well as blocking probability and burst-length distributions were derived. Numerical examples were presented for studying the effect of correlation of the arrival process on the performance metrics of a multibuffered ATM switch architecture. Comments on implication of the results as "building blocks" for the analysis of general ATM networks were given.

August Casaca

INESC, Lisbon, Portugal

An ATM-based Customer Premises Network for the Residential Environment

An architecture for the Customer Premises Network (CPN) based on ATM and applicable to residential and small business environments is described. The proposed solution consist of a distributed architecture and it is based on a shared medium with an access mechanism characterized by a Move headed First Served discipline, allowing also the establishment of access priorities. This discipline suits well the variation on the sources traffic parameters, typical of ATM, keeping on the other hand, the level of simplicity required by the target environments.

In more complex applications, the traffic concentration in each network node is achieved by ATM mergers, which have a decentralized contention mechanism to allow the evaluation of which input has the highest priority before each cell transmission.

Otto Spaniol

Department of Computer Science 4, Aachen University of Technology, Aachen, Germany

High Speed and Mobility

"High Speed" and "Mobility"...do these things fit together when we mean by "mobility" the more or less continuous communication made by constantly moving locations and effected by wireless (radio) communication [but not the so-called "nomadic" mobility which is characterized by plugging at different locations]. It is shown that indeed there is an inherent conflict - and there remain a lot of things to do in research and development in order to reach a certain compromise. A lot of design decisions, e.g. protocol stacks, will need to be rather completely redesigned.

Stephen Pink

Swedish Institute of Computer Science, Stockholm, Sweden

A Flexible Network Protocol for High Speed and Wireless Links

Today's network protocols are either datagram or virtual-circuit based, have fixed size headers, are based on unicast routing paradigms and see the stationary end-system as predominant. IP6 (IP Next Generation) could produce a hybrid VC/datagram architecture, but at the penalty of very large headers. We propose a new protocol (NP++) that allows the sender to send large addresses as frequently as necessary, trading off robustness for bandwidth on, for example, low speed wireless links. The header is made up of "chunks" that can be sent independently! Also, routing in NP++ uses multicast as primary with unicast as a special case. Mobile has to "join multicast groups" that have only one member; this provides unicast service.

Adam Wolisz

Institute for Telecommunication, Technical University Berlin, Berlin, Germany

Wireless Access to High Speed Networks

The need for (constrained) mobility causes increasing interest in wireless LANs, a pico-cell based kind of radio networks. Capacity (processing power, memory, battery) constrained terminals need relatively high bandwidth for asymmetric communication with higher use at downlink. Wireless LANs are to be seen as an access network to high speed (ATM) networks. One of the promising approaches is tunneling either ATM traffic of the UNI services over the WLAN as contrary to IP type of connectivity. Research topics concerning MAC requirements, link layer, mobility support have been addressed. Simulation experiments show that emerging WLAN will not support properly that traffic of short ATM-like cells. Also forwarding mechanism is not suited to extend the coverage in the radionet. Thus tunneling of UNI services, and using multicast to several base stations to support quick handover, together with asymmetric, partially FEC based link protocols may be part of the direction.

Benoit Macq

*Lab. de Telecommunications et Teledetectio, Université de Louvain
Louvain-la-Neuve, Belgium*

Conditional Access and Copyright Protection

Service Providers (SPv) intend to promote contracts with Users (U) through a mechanism which gives them certification of the U validity (the U will pay) and which allows them to address a wide market. U, from his side, would like to avoid unsolicited services and would like to preserve his anonymity (he does not want anyone to know what he is looking for). These contradictory requirements can be implemented by protocols using smart cards and Trusted Third Parties (TTP). The copyright issues of SPv may also make use of the TTP for monitoring pirate copies, if the original ones are labeled by a resistant watermark.

Dagstuhl-Seminar 9525:

Andres Albanese
Int. Computer Science Institute
1947 Center Street
Berkeley CA 94704
USA
aa@ICSI.Berkeley.EDU
tel.: +1-510-642-4274X134

Ernst Biersack
Institut Eurecom
B. P. 193
F-06904 Sophia Antipolis
France
erbi@eurecom.fr
tel.: +33-9300-2611

Olivier Bonaventure
Université de Liège
Institute Montefiore B28
B-4000 Liège
Belgium
bonavent@montefiore.ulg.ac.be
tel.: +32-41-66 26 99

Z. L. Budrikis
Curtin University of Technology
ATRI
G.P.O. Box 19 87
Perth 6001
Australia
zig@atri.curtin.edu.au
tel.: +61-9-351-32 41

Georg Carle
Universität Karlsruhe
Institut für Telematik
Zirkel 2
D-76128 Karlsruhe
Germany
carle@telematik.informatik.uni-karlsruhe.de
tel.: +49-0721-608-4027

Augusto Casaca
INESC
Apartado 13 069
Rua Alves Redol 9
1000 Lisboa
Portugal
augusto.casaca@inesc.pt
tel.: +351-1-31 00-2 33

List of Participants

Jon Crowcroft
University College London
Computer Science Department
Gower Street
London WC1E 6BT
Great Britain
j.crowcroft@cs.ucl.ac.uk
tel.: +44-171-3807296

Andre Danthine
Université de Liège
Institute Montefiore B28
B-4000 Liège
Belgium
danthine@vm1.ulg.ac.be
tel.: +32-41-66 26 91

Wolfgang Effelsberg
Universität Mannheim
Lehrstuhl Praktische Informatik IV
68131 Mannheim
Germany
effelsberg@pi4.informatik.uni-mannheim.de
tel.: +49-621-292-3131

Serge Fdida
Université Pierre et Marie Curie
Laboratoire MASI
4 Place Jussieu
F-75005 Paris
France
Serge.Fdida@masi.ibp.fr
tel.: +33-1-44.27.30.58

Domenico Ferrari
University of California at Berkeley
Computer Science Division
1947 Center Street
Berkeley CA 94720
USA
ferrari@icsi.berkeley.edu
tel.: +1-510 642 3806

Nicolas D. Georganas
University of Ottawa
Faculty of Engineering
EE Dept.
161 Louis Pasteur
Ottawa Ontario K1N 6N5
Canada
nngpb@acadvm1.uottawa.ca
tel.: +1--613-562-5800 Ext. 6225

Per Gunningberg
University of Uppsala
Dept. of Computer Science
P.O. Box 325
S-751 05 Uppsala
Sweden
per.gunningberg@docs.uu.se
tel.: +46-18-18 31 71

Demetres Kouvatso
University of Bradford
Department of Computing
Phoenix Building
Bradford BD71DP
Great Britain
demetres@comp.brad.ac.uk
tel.: +44-274-383-941

Aurel A. Lazar
Columbia University
Department of Electrical Engineering
500 West 120th Street
New York NY 10027-6699
USA
aurel@ctr.columbia.edu
tel.: +1-212-854-1747

Jean-Yves Le Boudec
Ecole Polytechnique Federale
Lab. Reseaux de Communication
(IN) Ecublens
CH-1015 Lausanne
Switzerland
leboudec@di.epfl.ch
tel.: +41-21-693-66 31

Benoit Macq
Université de Louvain
Lab. de Telecommunications
et Teledetection
Batiment Stevin
Place du Levant 2
B-1348 Louvain-la-Neuve
Belgium
macq@tele.ucl.ac.be
tel.: +32 10 47 22 71

Peter Martini
Universität GH Paderborn
FB 17 - Mathematik/Informatik
Warburgerstr. 100
33098 Paderborn
Germany
martini@uni-paderborn.de
tel.: +49-5251-60 33 37

Gurudatta Parulkar
Washington University
Dept. of Computer Science
One Brookings Drive
St. Louis MO 63130-4899
USA
guru@arl.wustl.edu
tel.: +1-314-9354621

Michael Paterakis
Technical University of Crete
Dept. of Electronic and
Computer Engineering
73 100 Chania Crete
Greece
pateraki@ced.tuc.gr
tel.: +30-821-46566; +30-821-69451

Stephen Pink
Swedish Institute of Computer Science
P.O. Box 1263
S-16428 Stockholm-Kista
Sweden
steve@sics.se
tel.: +46 8 752 15 59 / +46 8 38 60 15

Andreas Pombortsis
Aristotle University
Department of Computer Science
GR-54006 Thessaloniki
Greece
pombortsis@olymp.ccf.auth.gr
tel.: +30-31-99-80 45

Doug Shepherd
Lancaster University
Computer Centre
Lancaster LA1 4YR
Great Britain
doug@comp.lancs.ac.uk
tel.: +44-1524-59-3827

Otto Spaniol
RWTH Aachen
Lehrstuhl für Informatik IV
Ahornstr. 55
52056 Aachen
Germany
spaniol@informatik.rwth-aachen.de
tel.: +49-241-80-21400

Heinrich Stüttgen
IBM European Networking Center
Broadband Multimedia Communication
Vangerowstr. 18
D-69115 Heidelberg
Germany
stuttgen@heidelberg.ibm.com
tel.: +49-6221-594337

Ralf Steinmetz
IBM ENC
Vangerowstr. 18
D-69115 Heidelberg
Germany
steinmetz@heidelberg.ibm.com
tel.: +49-6221-59-42 80

A.N Tantawy
IBM T. J. Watson Research Center
P.O. Box 704
Yorktown Heights NY 10598
USA
ahmed@watson.ibm.com
tel.: +1-914-784-74 59

Fouad Tobagi
Stanford University
Department of Electrical Engineering
Stanford CA 94305
USA
tobagi@ee.stanford.edu
tel.: +1-415-723-17 08

Don Towsley
University of Massachusetts
Computer Science Department
Amherst MA 01003
USA
towsley@cs.umass.edu
tel.: +1-413-545-0207

Phuoc Tran-Gia
Universität Würzburg
Informatik III
Am Hubland
D-94074 Würzburg
Germany
trangia@informatik.uni-wuerzburg.de
tel.: +49-931-888-5509

Harmen R. Van As
IBM Research Laboratory
Säumerstraße 4
8803 Rüschlikon
Switzerland
vas@zurich.ibm.com
tel.: +41-1-724-84 83

Adam Wolisz
TU Berlin
Institut für Fernmeldetechnik
Einsteinufer 25
D-10587 Berlin
Germany
wolisz@fts00.ee.tu-berlin.de
tel.: +49-30-314-22 911

Raj Yavatkar
University of Kentucky
Department of Computer Science
773 Anderson Hall
Lexington KY 40506-0046
USA
raj@dcs.uky.edu
tel.: +1-606-257-67 45

Martina Zitterbart
TU Braunschweig
Institute of Operating Systems & Computer
Networks
Büttenweg 74/75
D-38106 Braunschweig
Germany
zit@ibr.cs.tu-cs.de
tel.: +49-531-391-3288

Zuletzt erschienene und geplante Titel:

- W. Maass, Ch. v.d. Malsburg, E. Sontag, I. Wegener (editors):
Neural Computing, Dagstuhl-Seminar-Report; 103; 07.11.-11.11.94 (9445)
- G. Berry, W.P. de Roever, A. Pogné, A. Pnueli (editors):
Synchronous Languages, Dagstuhl-Seminar-Report; 104; 28.11.-02.12.94 (9448)
- B. Becker, R. Bryant, O. Coudert, Ch. Meinel (editors):
Computer Aided Design and Test, Dagstuhl-Seminar-Report; 105; 13.02.-17.02.95 (9507)
- D. Garlan, F. Paulisch, W. Tichy (editors):
Software Architectures, Dagstuhl-Seminar-Report; 106; 20.02.-24.02.95 (9508)
- W.J. Cullyer, W.A. Halang, B. Krämer (editors):
High Integrity Programmable Electronic Systems, Dagstuhl-Seminar-Report; 107; 27.02.-03.03.95 (9509)
- J. Gruska, H. Umeo, R. Vollmar (editors):
Cellular Automata, Dagstuhl-Seminar-Report; 108; 06.03.-10.03.95 (9510)
- H. Alt, B. Chazelle, R. Seidel (editors):
Computational Geometry, Dagstuhl-Seminar-Report; 109; 13.03.-17.03.95 (9511)
- W. Bibel, K. Furukawa, M. Stickel (editors):
Deduction, Dagstuhl-Seminar-Report; 110; 20.03.-24.03.95 (9512)
- B. Freitag, C.B. Jones, Ch. Lengauer, H.-J. Scheck (editors):
Object-Oriented Programming with Parallelism and Persistence, Dagstuhl-Seminar-Report; 111; 03.04.-07.04.95 (9514)
- J. Doran, N. Gilbert, U. Mueller, K. Troitzsch (editors):
Social Science Microsimulation: A Challenge for Computer Science, Dagstuhl-Seminar-Report; 112; 01.05.-05.05.95 (9518)
- R. Studer, M. Musen (editors):
Shareable and Reusable Problem Solving Methods, Dagstuhl-Seminar-Report; 113; 08.05.-12.05.95 (9519)
- J. Blazewicz, K. Ecker, L. Welch (editors):
Scheduling in Computer & Manufacturing Systems, Dagstuhl-Seminar-Report; 114; 15.05.-19.05.95 (9520)
- H. Beilner, G. Ciardo, C. Lindemann, K. Trivedi (editors):
Performance and Dependability Modeling with Stochastic Petri Nets, Dagstuhl-Seminar-Report; 115; 22.05.-26.05.95 (9521)
- M. Aigner, J. Spencer, E. Triesch (editors):
Computing with Faulty Inputs, Dagstuhl-Seminar-Report; 116; 29.05.-02.06.95 (9522)
- J.-R. Abrial, E. Börger, H. Langmaack (editors):
Methods for Semantics and Specification, Dagstuhl-Seminar-Report; 117; 05.06.-09.06.95 (9523)
- W. Effelsberg, D. Ferrari, O. Spaniol, A. Danthine (editors):
Architecture and Protocols for High Performance Networks, Dagstuhl-Seminar-Report; 118; 19.06.-23.06.95 (9525)
- Ph. Flajolet, R. Kemp, H. Prodinger, R. Sedgewick (editors):
'Average-Case'-Analysis of Algorithms, Dagstuhl-Seminar-Report; 119; 03.07.-07.07.95 (9527)
- D. Gustfield, T. Lengauer, C. Sander (editors):
Molecular Bioinformatics, Dagstuhl-Seminar-Report; 120; 10.07.-14.07.95 (9528)
- J. Chomicki, G. Saake, C. Sernadas (editors):
Role of Logics in Information Systems, Dagstuhl-Seminar-Report; 121; 17.07.-21.07.95 (9529)