

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

**Architecture and Protocols for  
High-Speed Networks**

Dagstuhl-Seminar-Report; 71  
30.08.-03.09.93 (9335)

ISSN 0940-1121

Copyright © 1993 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.-Ing. José Encarnação,  
Prof. Dr. Hans Hagen,  
Dr. Michael Laska,  
Prof. Dr. Thomas Lengauer,  
Prof. Dr. Wolfgang Thomas,  
Prof. Dr. Reinhard Wilhelm (wissenschaftlicher Direktor)

Gesellschafter: Universität des Saarlandes,  
Universität Kaiserslautern,  
Universität Karlsruhe,  
Gesellschaft für Informatik e.V., Bonn

Träger: Die Bundesländer Saarland und Rheinland-Pfalz

Bezugsadresse: Geschäftsstelle Schloss Dagstuhl  
Universität des Saarlandes  
Postfach 1150  
D-66041 Saarbrücken, Germany  
Tel.: +49 -681 - 302 - 4396  
Fax: +49 -681 - 302 - 4397  
e-mail: office@dag.uni-sb.de

## **Architecture and Protocols for High-Speed Networks (August 30 - September 3, 1993; Dagstuhl, Germany).**

organized by:

**André Danthine**  
University of Liège

**Wolfgang Effelsberg**  
University of Mannheim

**Otto Spaniol**  
Techn. University of Aachen

The first seminar on "Architecture and Protocols for High-Speed Networks" was held with the purpose of discussing important questions concerning the next generation of digital computer networks. The main topics of the seminar were:

- Switched networks, in particular ATM
- Local and Metropolitan Area Networks
- New Network and Transport Layer protocols
- Network applications, in particular multimedia applications
- Protocol implementation on multiprocessors, and
- Formal description techniques.

Those topics are in the center of interest for research and implementation of communication systems. This is proven also by the fact that in the very same week the "Second IEEE Workshop on the Architecture and Implementation of High Performance Communication Subsystems" was held in Williamsburg, Virginia (USA). The exact duplication of topics of interest in both events made it difficult or impossible for some of the invited US experts to attend the Dagstuhl workshop. Nevertheless, 32 participants came from many European countries (including Eastern Europe), from the USA, Canada and Australia.

It was the purpose of the seminar to bring together telecommunications engineers and computer scientists, two groups of people who not very often have a chance to talk with each other. One of the hot topics during the seminar week was the status and future of ATM (Asynchronous Transmission Mode). Analytical and simulation models for the performance of ATM switches were presented, the status of ATM demonstrators in the framework of the RACE project in Europe was reported, and the ATM public service offering by PTTs was the subject of much debate. Although ATM was initially designed to provide a wide-area high-speed telecommunications infrastructure, all the practical experience reported by the participants was with ATM switches in a local environment.

The new generation of applications in high-speed networks will contain multimedia data streams, i.e. digital audio and video. All of the speakers on upper layer issues agreed on this point. However, continuous media streams require transmission with guaranteed performance, in particular guaranteed bandwidth and bounds for delay and jitter. In addition, many multimedia applications will require peer-to-multipeer communication. Guaranteed performance can only be provided with resource reservation in the network, and efficient multipeer communication must be based on multicast support in the lower layers of the network. In several sessions the current approaches to resource reservation and multicast in computer networks were discussed.

In the area of communication systems architecture several speakers reported their research work on internal structures for high-speed communication nodes. There was general agreement that the performance bottleneck is currently in the end systems, upper layers and applications rather than in the MAC adapters, on the links or in the switch fabrics. Parallel implementation of protocols on multiprocessors was considered as a possible solution, with very promising measurement results reported from a parallel implementation of TCP and IP.

A special highlight of the workshop was the presentation by Z. Budrikis, the co-inventor of the DQDB (or QPSX) protocol. He introduced a new medium access control protocol for a LAN/MAN called DQDT (Distributed Queue - Dual Tree). It is based on tree topology of interconnected dual buses and maintains the DQDB advantages of short latency and high utilization with insensitivity to the size and speed of the network.

In two evening discussion rounds the participants came back to open questions on ATM. The provision of "statistical performance guarantees" was subject of a lively debate; isn't a "statistical guarantee" an inherent contradiction? How much would you be willing to pay for such a service offer? Another hot issue was whether the networking community is thoroughly spoiled by the OSI reference model, under the motto "layering considered harmful". Of course no agreement could be reached on this issue.

Some more general comments about Dagstuhl workshops:

At first, the organisers were a little sceptical whether the participants would be willing to come for a full week since that duration is in sharp contrast of other meetings or conferences which usually are no longer than three days (and where most participants only attend the presentations during one or two days). The experience of the workshop, however, was that everybody liked the very creative atmosphere - and almost everybody attended all the presentations including the evening sessions. The discussions were extremely lively and intensive. Even the 45 minute time slots per presentation were far from being sufficient for solving all of the numerous questions. We may proudly state that a similar learning effect has never been observed in a 'classical' conference.

Moreover, the fact that the participants were real experts in the field and that they knew each other very well (at least after the first two days) became a guarantee for very (!) critical comments; however, nobody was upset about "nasty questions" and, therefore, the global outcome was extremely positive.

All of those positive things were supported and greatly facilitated by the wonderful setting of the workshop in the castle of Dagstuhl. A library, Unix workstations, musical instruments and a well-stocked wine cellar were at the disposition of the participants, and everybody enjoyed the casual and open-minded atmosphere.

A second workshop is planned for June 1995.

**Dagstuhl, September 1993**

André Danthine, Wolfgang Effelsberg, Otto Spaniol

## **Workshop programme**

### **Monday (Aug. 30):           MAC, ATM, Switched Networks,...**

J.W. Mark:

Performance Modelling and Buffer Allocation in ATM Networks

P. E. Boyer:

Evolution Towards Switched Services in ATM Networks.

A. Pombortsis:

A contribution to the problem of avoiding and controlling congestion in multistage networks for ATM switching

P. Tran-Gia:

Discrete-time analysis of Usage Parameter Control Functions in ATM Systems

P.J. Kühn:

Quality of Service in ATM-Networks - Key Issues and Demonstrator Verification

Z. L. Budrikis:

Size and Speed Insensitive Distributed Queue Network

### **Tuesday (Aug. 31)           Network and Transport Layer,...**

O. Casals:

A Cell Based MAC Protocol with Traffic Shaping and a Global FIFO Strategy

M. Becker:

Models of ATM switches.

H. Stüttgen:

An Internetworking Architecture for Multimedia Communication over Heterogeneous Networks

W. Kalfa:

Receiver initiated multicast in the ST-II protocol

D. Ferrari:

The Tenet Approach to Real-Time Communication and the Tenet Real-Time Protocol Suite

E. W. Biersack:  
The TP++ protocol

R. G. Herrtwich:  
Reservation Mechanisms for Internetworks

O. Spaniol:  
High Speed Interconnection of Workstations: Concepts, Problems and Experiences  
(not presented due to lack of time; manuscript was distributed)

Discussion: ATM - success or disaster?

### **Wednesday (Sept. 1) : FDT, Multiprocessor Architecture,...**

P. Gunningberg:  
Locking Effects in Multiprocessor Implementations of Protocols

K. Fabian, (I. Fellner):  
A Multiprocessor Communication Adapter for High-speed Networks

A. Wolicz:  
An integrated approach to communication protocols specification and performance analysis.

B. Wolfinger:  
Formal description technique to support load modelling for innovative communication systems

Afternoon: Excursion to Trier (Trèves)

### **Thursday (Sept. 2): Multimedia, Applications, QoS,...**

B. Heinrichs:  
Towards a High Performance and Configurable Multipeer Transfer Service

K. Rothermel:  
A Configurable Architecture for Distributed Multimedia Applications

W. Tawbi:  
QoS Issues for Distributed Multimedia Applications

R. Steinmetz:  
Media Synchronization

A. Danthine:  
From best effort to enhanced QoS

O. Drobnik:  
Towards an integrated configuration and performance management for distributed applications

H. Schlichter:  
High Bandwidth Communication for Computer Supported Cooperative Work (CSCW)

D. Shepherd:  
MICRO-Kernel Support for continuous Media Communications

Discussion: Most important open questions

## **Friday (Sept. 3):                      Traffic Control,...**

D. Baum:  
Combined Flow and Error Control Analysis for the European Space Data Network.

H. Löffler:  
Topology-Bandwidth Considerations of Gbit/s-Networks

E. Raubold:  
Deadlock avoidance in networks (Lessons learned from packet switching for HPNs)

W. Effelsberg:  
Digital Movies in High-Speed Networks

O. Spaniol:  
Concluding remarks (and "the poem")

## **Abstracts of presentations:**

**Dieter Baum, Universität Trier, Trier (Germany)**

### **Combined Flow and Error Control Analysis for the European Space Data Network**

The impact of window flow control and GBN-ARQ error control mechanisms is investigated for the proposed Space Data Network chain consisting of an on board LAN, a satellite connection, and an ISDN terrestrial WAN. Two architectural alternatives for protocol hierarchies lead to two different basic approaches for error and flow control implementation: End-to-End versus Step-by-Step realization. These approaches correspond to the types of network service used to perform subnetwork interconnection: connection oriented and connectionless. The packet delay behaviour under variable load is analyzed for each scheme considering the combined influence of flow and error control. Analysis is based on standard techniques being adjusted to the concrete more complex system. Results prove the superiority of Step-by-Step approaches. They have been validated by simulation.

**Monique Becker, Institut National des Telecommunications, Evry Cedex (France)**

### **Models of ATM switches**

In order to design a three stage Clos ATM switch two performance criteria are studied:

- cell loss rate
- average delay.

Two kinds of switches are studied:

- switches with input queues and
- switches with output queues.

The performance criteria are compared for a given value of input links in the switch (which is equal to the number of output links) and for a given value of the global buffer size for the whole switch.

The design parameters are:

- the number of switching elements in the two first stages
- the capacity of the queues
- the input load on each link
- the destination of the calls.

Approximate analytical models are solved and validated by simulations for some parameter values (there is a problem of simulation of rare events). The model is solved for symmetrical input traffic and for some cases of unsymmetrical input traffic load.

**Ernst Biersack, Institut Eurecom, Valbonne (France)**

### **The TP++ protocol**

TP++ is a new transport protocol that is being designed to support multimedia applications at one Gigabit per second across high bandwidth-delay product networks. This talk presents an overview of the TP++ project and discusses the goals of the project as well as a description of the environment for which TP++ is designed. The basic design decisions made for TP++ then are described.

**Pierre Boyer, CNET - Lannion, Lannion Cedex (France)**

### **Evolution Towards Switched Services in ATM Networks.**

Even if "switching" usually refers to the basic interconnecting capability of network elements, the definition of "switched" services is basically a matter of signalling procedures and call admission control: call set-up is performed within a few seconds in a "switched" network - opposed to cross-connected networks where call set-up is a matter of days.

At the beginning, operators will use ATM to build networks of Virtual Paths. Functionally similar to leased lines, these paths will be cross-connected and allocated a peak cell rate value. Compared to N-ISDN, the fine rate granularity and flexible QOS are sufficiently attractive to motivate initial subscription.

However, the specifications of telecommunication services will be subject of competition between operators. In this perspective, the introduction of switched services enables a shorter network response time.

At the same time, traffic management is becoming more complex. Pricing will become crucial which entails an optimized usage of network resources. A large variety of applications offer bursty traffics but cannot give the accurate description which is requested to perform simple rate-based multiplexing. Applications may even exceed the negotiated activity and the excess traffic has to be transferred as well as possible.

For the operator, a "risk" is coming with the introduction of "switched" services :

1. the network may deny or delay admission of a candidate switched connection;
2. the network may not meet the Quality-Of-Service requested by a newly accepted switched connection; the admission of a new switched connection may jeopardize other calls already in progress - which is even worse.

Starting with a cross-connected virtual path network, it is proposed to introduce switching gradually into the network operation :

1. Routes could still be cross-connected throughout the network - by means of classical telecommunication network Operating Systems and workstations - while bandwidth could be allocated "on-demand" within a very short amount of time by means of Fast Resource Management - e.g. FRPs.
2. Further on, when a signalling system will be finalized and agreed upon in standardizing bodies, "switched" services could initially be offered within a peripheral network area, between subscribers and servers. These servers could still be interconnected by long-haul cross-connected connections.

Finally the paper details how this evolution towards switched services translates into subscriber traffic contracts.

**Z.L. Budrikis, QPSX Communication LTD, West Perth 6005 (Australia)**

#### **Size and Speed Insensitive Distributed Queue Network**

A Distributed Queue Dual Bus (DQDB) Network is outstanding among known LANs by its high speed x size capability and by its approximation to a FIFO multiplexer. Nonetheless, a DQDB network is limited to several hundred Mbps x km before its efficiency and good multiplexing properties begin to deteriorate. A distributed queue network employing a dual tree topology, i.e. a distributed queue dual tree (DQDT) network, is described that is effectively insensitive to size and speed. A DQDT network with a size x speed product of 10 km x Gbps shows no noticeable reduction in efficiency or fairness. It is an obvious candidate for shared LANs and MANs at Gigabit/s rates.

**Olga Casals Torres, Universidad Polit cnica de Cataluna, Barcelona (Spain)**

#### **A Cell Based MAC Protocol with Traffic Shaping and a Global FIFO Strategy**

A Medium Access Control (MAC) protocol for a broadband network access facility, using a passive optical network (PON) with a tree structure will be presented. This protocol will be implemented in the demonstrator of the RACE project R2024. The MAC protocol is cell based and uses a request/permit mechanism to access to the shared medium. Each Network Termination (NT) declares its required bandwidth by sending requests to the master of the protocol located in the Line Termination (LT). A request contains information about the number of cells waiting for transmission at the NT. The requests can be issued in two ways:

- Coupled to upstream ATM cells
- As part of a Request Block.

If an NT can only reveal its bandwidth needs when it is allowed to send an ATM cell, it may not be able to react fast enough to traffic changes. Therefore, the protocol is provided with the second type of request. The so called Request Blocks contain requests originating from several NTs, not coupled to upstream ATM cells. They are issued during idle periods. The combination of these two mechanisms leads to a fast reaction time.

The MAC protocol allocates the available bandwidth to the NTs according to a Bandwidth Allocation Algorithm based on the received information from the requests. The MAC protocol uses permits to inform the NTs about the obtained bandwidth. For each permit received, an NT can send one cell.

The Bandwidth Allocation Algorithm used by the MAC protocol has the following characteristics:

- Enforce the peak bit rate per NT by spacing the cells
- Approximate a global FIFO discipline (over all NTs) for fairness reasons and to minimize the delay variance
- Use the available buffer capacity in the NTs to store the cells waiting for transmission (distributed buffering)

The peak bit rate of each NT is calculated as the sum of the peak bit rates of each of its connections (VCI/VPI). The MAC protocol counts the number of new cell arrivals (deduced from the request) and assigns the necessary permits to the NT. These permits are then put into A FIFO queue with permits for the other NTs, with the additional constraint that a minimal distance between two consecutive permits for the same NT is enforced. The minimal distance is determined by the inverse of the peak bit rate. In this way, the actual queueing takes place in the NT, while the central control of the PON maintains a FIFO queue, by which the transmission instants and the order in which the different NTs are emptied are governed.

**André Danthine, Université de Liège, Liège (Belgium)**

### **From Best Effort to enhanced QoS**

This technical report is entirely devoted to the problem of the QoS in the connection-mode services. After presenting a framework for the study of the QoS parameters including the QoS negotiation scheme, the QoS semantics as well as the constraints of the service user or service provider, we discuss how the QoS is handled in the network and transport layers for the ISO on one hand and for the Internet on the other hand. Then we present the Enhanced QoS which has been developed in the framework of the OSI 5 Transport Services. In the last part of this technical report we compare our proposals to other new transport services: ETS, HSTS and BIS.

**Oswald Drobnik, Universität Frankfurt, Frankfurt / Main (Germany)**

**Towards an integrated configuration and performance management  
for distributed applications**

A new approach aiming at a synthetical construction and management of distributed applications will be presented. A specification technique has been developed to support the different aspects of a distributed application, namely application functionality, management properties, and communication requirements. It is not restricted to client server applications and can be used for a variety of cooperation paradigms including peer to peer and group cooperation. A classification of management properties allows for obtaining systematically a tailor-made management architecture.

The constructs of the specification language can be mapped onto hierarchies of application-, communication- and management-oriented classes (in C++). Implementation of a distributed application including management functionality is then achieved by selection and configuration of appropriate objects. Guidelines have been developed to improve the runtime efficiency of the resulting software architecture by using threads.

Our approach has been validated by a prototype implementation of a Message Handling System application including an interactive configuration management facility (as part of an external research project with Digital Equipment Munich).

At present, the integration of performance monitoring into this approach is under implementation. Concepts for decentralized management architectures are being developed in order to support efficient configuration and performance management for large distributed applications.

**Wolfgang Effelsberg, Uni Mannheim, Mannheim (Germany)**

**Digital Movies in High-Speed Networks**

The University of Mannheim is developing a system called XMovie for storing, transmitting and presenting digital films in a computer network. The hardware used in the system is standard hardware, as found in typical workstations today; no special hardware is required. The movies are shown in a window of the X Window System. This allows full integration with the classical components of computer applications, such as text, color graphics, menus and icons.

The emphasis of the research project is on transmission protocols for digital video data streams. Compared with classical data communications the transmission of video streams has very demanding requirements for high bit rates and for isochronous transfer, i.e. transfer with a small delay jitter. We show why these requirements are not met with traditional communication protocols (such as TCP/IP or OSI), and how they can be achieved with the new generation of high-speed networks. We also propose an application layer protocol for the transmission of movies.

**Ivan Fellner, Slovak Academy of Sciences, Banska Bystrica (Slovak Republic)**

### **A Multiprocessor Communication Adapter for High-speed Network**

High speed data link transfer of the new networks can not be sufficiently utilized by present implementations of network and transport protocols. In the paper we present new architecture of multiprocessor adapter for high speed network, which implements the light-weight XTP protocol. We show the reasons led us to the fine grain multiprocessor, which combine processor per function principle in the simple hardware blocks with the implementation of the less urgent functions on software basis. At the end we give performance evaluation of this adapter.

**Domenico Ferrari, University of California at Berkeley, Berkeley, CA 94720 (USA)**

### **The Tenet Approach to Real-Time Communication and the Tenet Real-Time Protocol Suite**

The Tenet approach to the design of real-time communication services is the foundation of one of the first protocol suites built to support multimedia and other real-time transmissions in integrated-services internetworks. The same approach is now being used in the design of another, more advanced protocol suite, based on a multicast guaranteed-performance channel abstraction.

In this talk, the principles of the Tenet approach and the main features of the four protocols that constitute the first Tenet suite will be described. Some of the initial experiences with using the suite in some of the several packet-switching and cell-switching testbeds to which it is being ported will also be presented, and the most important characteristics of the second Tenet suite will be outlined.

**Per Gunningberg, Swedish Institute of Computer Science, Stockholm-Kista  
(Sweden)**

### **Locking Effects in Multiprocessor Implementations of Protocols**

Many workstations will consist of a multiprocessor system with a small number of processors, typically less than ten. We investigate in how to exploit these processors for parallel protocol processing. A "processor-per-message" paradigm is used to partition the work over processors. Locks are used to protect protocol state and data among processors which need to be shared exclusively. Mutual exclusion by locking can be costly if the parallel protocol code frequently accesses shared state and data. This paper addresses the effect of locking on performance.

We present a multiprocessor implementation of the x-kernel protocol environment from University of Arizona and a simulation tool based on a parallel protocol execution model for predicting locking effects on performance. The x-kernel is implemented in user address space on a Sequent Symmetry under Dynix Operating System.

Performance results from multiprocessor implementations of UDP/IP and TCP/IP are given and we compare them to performance predictions from simulations. The measured speedups for the parallel implementations relative to the sequential implementations are more than 12 times for UDP and 3 times for TCP. These limits are set by locking and machine contention.

**Bernd Heinrichs, RWTH Aachen, Aachen (Germany)**

### **Towards a High Performance and Configurable Multipeer Transfer Service**

This paper highlights the need for sophisticated multicast mechanisms to be provided by transfer protocols in order to support group communication. This need becomes even more evident in the light of the special QOS requirements imposed by multi-media applications. Today's protocols do not provide the broad range of functionality required by these upcoming new applications. From a number of sample applications we derive functionality and performance requirements which have to be provided by transfer systems.

Although various transport protocols are capable of providing some basic functionality used for multicast applications, they do not address all of the requirements of multicast applications. Therefore, some of the unsolved or very difficult issues for the provision of multicast services, like application specific error control, acknowledgement implosion, support of different grades of reliability, guarantee of different QOS requirements, scalability etc. will be discussed.

After comparing different ISO and Internet approaches the XTP multicast mechanism is illustrated in more detail. We show how some of the identified demands are met by this protocol. Furthermore, necessary enhancements are suggested. The usefulness of these enhancements is proved by simulations and measurements.

Although XTP provides an interesting approach it is not possible to enhance the protocol by integrating all the desired functionality without changing the primary semantics of XTP. Thus we specify our own approach, called AMTP (Adaptive Multipeer Transfer Protocol), a new multicast transfer protocol capable to provide high performance and flexibility.

**Ralf-Guido Herrtwich, IBM ENC, Heidelberg (Germany)**

### **Reservation Mechanisms for Internetworks**

In conjunction with providing guaranteed or predictable service for multimedia applications, several reservation schemes for network bandwidth have been proposed. Among the most prominent examples are ST-II and RSVP which originate from the Internet community. Other examples include SRP and RCAP from Berkeley and CBSRP from CMU. In this talk, we compare the different approaches and identify their respective merits and disadvantages in respect to multimedia applications that shall benefit from them. We also take a critical look at the usefulness and limitations of reservation schemes in general.

**Winfried Kalfa, TU Dresden, Chemnitz (Germany)**

### **Receiver initiated multicast in the ST-II protocol**

In the specification of the the ST-II only the source of the multicast stream is able to connect a new receiver to the stream. This needs an alternate communication connection to inform the source about an new receiver. A better way consists in the use of the ST-II protocol themselves. Two versions are presented by addition of new features to the protocol. The arised problem of authorization are discussed.

## **Quality of Service in ATM-Networks: Key Issues and Demonstrator Verification**

The Asynchronous Transfer Mode (ATM) is the basis of the future Broadband-ISDN which allows for the integration of arbitrary constant and variable bitrate services. The international standardization is almost finished. The feasibility of the technology has been shown in key research and development initiatives in the USA, Europe and in Japan. Especially, the European RACE programme has resulted in broadband islands with ATM switch technology (RACE demonstrators). The major network operators start their pilot service in 1994/95.

Despite all these advances, ATM bears still some largely unsolved problems. Due to the inherent resource sharing problem of the cell-based, fast packet switching principle, cell losses and cell delay variations depend strongly on the statistical (bursty) behaviour of the source traffic, such as packetized voice, video or high speed data applications. The economic effect of the multiplexing gain is strictly limited by the Quality of Service (QoS) requirements of such services, which are indicated in terms of cell loss probability, and cell delay variation of the network performance. For an economic use of the expensive network resources, flexible and efficient traffic control schemes have to be applied, such as source policing (usage parameter control, UPC), connection admission control (CAC), congestion control and fast reservation protocols.

Traffic performance and engineering methods have been developed during the recent years. The underlying source models are mostly based on multi-level (multi-rate) cell generators such as finite state machine modulated deterministic or stochastic processes (GMDP, MMPP, MMDP, MAP). The parameters of these models have been validated for known classical services such as voice or LAN data. The sound knowledge of the individual and of aggregated traffic streams is a prerequisite for traffic engineering. Therefore, source traffic and network performance have to be validated for a reliable network planning and network operation.

The contribution is organized as follows: Chapter 1 briefly introduces the ATM network architecture and services. In Chapter 2, the nature of typical ATM traffics is characterized revealing highly correlated discrete time point processes and some popular modelling approaches. In Chapter 3, the network performance relevant QoS metrics are addressed, which are subject to negotiation in the future ATM networks. Some examples of performance evaluation are given to show the parametric influence of traffic parameters, namely the finite buffer loss probability and the end-to-end cell delay variation within a serial queueing chain considering a reference virtual channel connection. Also, the inadequacy of traditional simulation techniques for high speed networks is addressed.

Within the RACE project EXPLOIT traffic measurements and traffic control experiments are in progress using the just available first ATM traffic generator and traffic measurement equipment. In Chapter 4, a short introduction is given in the ATM traffic generator and measurement equipment and techniques. The RACE EXPLOIT Experimental testbed is shortly explained in Chapter 5 which consists of a Local and Transit ATM Exchange, a Remote ATM Exchange, a Business and a Residential B-NT2 (broadband BPX), various VBR and CBR B-ISDN Sources, LAN and ISDN Interworking Units and Bulk Traffic Generators. This testbed is operable in October 1993 and subject of the ATM Traffic Experiments. From these results, major insights are expected in the nature of traffic and the effectiveness of the various traffic control schemes which will have feedback for standardization, network planning, network operation and the development of new applications, such as multi-media communication.

H., Löffler, TU Dresden, Dresden (Germany)

### Topology-Bandwidth Considerations of Gbit/s-Networks

At given work load, the performance of Gbit/s-networks depends on a set of factors: efficiency of the use of offered bandwidth of the fiberoptic medium, medium access protocols, topology and others. This paper is concerned with topology-bandwidth relationships. We start with definitions of aggregate bandwidth of the network, access rates of the nodes connected to the network, network channel efficiency, access efficiency of an end node etc. Then we present a set of characteristic criteria for selected topologies. In order to evaluate topologies of high performance networks, the following parameters are of importance: scalability, normalized aggregate bandwidth, normalized effective bandwidth, mean access efficiency of end nodes, mean path length, number of electro-optical elements etc. We demonstrate that for selected topologies. We also present some results concerning shared medium architectures (multiple access networks). It will be shown that "classical topologies" such as ring, folded bus, dual bus do not meet all our requirements for high performance networks. Results of load-separating architectures demonstrate that some of them fulfill the requirements with respect to efficiency and scalability. It is remarkable that the improvement of the efficiency due to load separating is not as high as expected. Probably, moreover, we found that architectures with multiple shared bandwidths which use WDM (wave-length division multiplexing) may be an approach for effective topologies of high-performance networks.

Jon W. Mark, University of Waterloo, Waterloo Ontario N2L 3G1 (Canada)

### **Performance Modelling and Buffer Allocation in ATM Networks**

An ATM network is a mesh connection of switches and multiplexers, in which switching and multiplexing are cell-oriented. ATM supports connection-oriented services; a call must be admitted and a connection established before communication can take place. Call admission requires a suitable definition of traffic descriptors and an effective mechanism to allocate resources to support the connection. It is thus important to understand the end-to-end network performance and the effects of admitting a new call on the performance of in-progress calls. To contribute to this understanding, we consider the queue management and cell loss at a generic switching node, on the premise that an end-to-end network partition can be modelled by chaining generic nodes.

The generic node is endowed with a finite amount of buffer space and a service discipline. We will discuss the modelling and buffer allocation, and the evaluation of cell loss and cell delay measures at the generic node. Based on the performance characteristics of the generic node, conjectures regarding the performance of an end-to-end network partition will also be discussed.

Andreas Pombortsis, Aristotele University Thessaloniki, Thessaloniki (Greece)

### **A contribution to the problem of avoiding and controlling congestion in multistage networks for ATM switching**

Self-routing, packet-switched Multistage Interconnection Networks (MINs) are well known from parallel computer architectures. More recently, they have been proposed as building blocks for the implementation of an ATM (Asynchronous Transfer Mode) switching node supporting Integrated Broadband Communication Networks (IBCNs). In MINs the presence of unbalanced traffic loads, in which packets tend to follow fixed paths through the network, has an important impact on the performance characteristics of the network. More precisely, the MINs suffer from the serious problem of rapid and significant performance degradation when there is heavy traffic to one or a set of output ports. Within the network such a traffic may create several "communities of interest". Several methods have been proposed to solve the problem. The major objective of this presentation is to propose and analyze various combinations of feedback schemes and switching strategies, at the packet level in MINs. The proposed flow control procedures are aiming to detect the possibility of congestions, as soon as possible, and using proper switching strategies, to prevent the creation of saturated trees, by utilizing the knowledge about the traffic load distribution. Generally, the feedback mechanism is based on the queue size of each Switching Element, in the final network stage or in stages within the network.

Eckart Raubold, GMD Darmstadt, Darmstadt (Germany)

**Deadlock avoidance in networks**  
(Lessons learned from packet switching for HPNs)

From '75 to '80 we studied at the GMD the behaviour of packet switching networks under high loads. Our requirements had been

- no congestion
- fair sharing of bottleneck resources under high load conditions.

Because most of the research at that time had been devoted to minimization of packet loss probabilities we concentrated on methods which avoid deadlocks within the network by construction and allow therefore to build pure wait systems without any packet loss. After development of a general proof technique we looked into two mechanisms for controlled resource sharing with deadlock prevention guaranty: Buffer allocation restrictions and routing restrictions. The first one resulted in ordered buffer classes, the second in networks with some prohibited paths in order to prevent circular waiting conditions. Some results could be gained on the properties of both mechanisms. Because the problems of optimal resource sharing under the constraints of service guaranties are still in principle the same in now-a-days high speed networks, a review of our results might be of interest and use to todays network designers.

Kurt Rothermel, Universität Stuttgart, Stuttgart (Germany)

**A Configurable Architecture for Distributed Multimedia Applications**

The CINEMA Project aims at the development of a platform for configurable distributed multimedia applications. Distributed applications consist of a number of components, which are sources or sinks of multimedia streams, or are functions that manipulate one or more data streams. Basic components can be used to compose more complex ones, and components can be linked with each other to configure arbitrary complex applications. To transfer multimedia information channels with end-to-end characteristics can be established. A logical clock abstraction is provided to express synchronization requirements between streams. The major goal of the CINEMA platform is to ease the development of distributed multimedia applications by providing a simple configuration mechanism and hiding most aspects of resource management, scheduling and synchronization.

The project is in an early stage. The major concepts and the architecture of CINEMA are rather stable. Some components have been prototyped.

Hans Schlichter, TU München, München (Germany)

### **High Bandwidth Communication for Computer Supported Cooperative Work**

The evolution of high bandwidth network technology, such as broadband ISDN or ATM, will have significant impact on the way how teams of people cooperate and how they are supported by computer systems. These emerging technologies offer the potential for a variety of newly designed cooperative applications. For example, high performance networks will provide opportunities for multimedia conferencing between geographically remote sites; data, audio and video are integrated into a single information stream supporting real-time interaction between group members.

The impact of high bandwidth on CSCW must be examined from two perspectives: the application perspective and the system support perspective. In the first case new forms of cooperation will be identified, e.g. how can virtual reality be applied to support cooperation across wide area networks. The second perspective will examine the distributed system support for CSCW in the context of high bandwidth communication and its impact on policies and mechanisms for group work, e.g. identifying suitable concurrency control strategies and notification mechanisms to support group awareness.

Doug Shepherd, Dept. of Computing, Lancaster LA1 4YR (United Kingdom)

### **MICRO-Kernel Support for continuous Media Communications**

Operating system support for distributed applications primarily involves the co-ordinated real-time processing and communication of continuous media such as audio and video. Continuous media types are demanding on distributed systems not only because of their high bandwidth requirements, but also because they demand a flexible range of qualities of service (QoS) at varying levels of commitment. A considerable amount of research has already been carried out in communications support for continuous media over high speed networks such as FDDI, DQDB and ATM. However, most experimental implementations of communications systems so far reported have not been successfully integrated into conventional end-system software environments. Typically, implementations have either been embedded in conventional operating system environments and suffered from poor performance, or have been implemented in specialised hardware/ software environments unable to support general purpose applications. Research at Lancaster is attempting to combine the advantages and eliminate the disadvantages of these two implementation strategies by efficiently supporting state of the art continuous media networking and communications support in a real-time, micro-kernel based environment which is also capable of running conventional UNIX applications. This talk will describe how we are extending the Chorus micro-kernel to incorporate high speed connection oriented communications, QoS configurability and integrated communications driven thread scheduling.

Otto Spaniol, RWTH Aachen, Aachen (Germany)

## High Speed Interconnection of Workstations: Concepts, Problems and Experiences

High speed networking technologies with transmission capacities in the range of hundreds of megabits per second (or even gigabits per second) together with very low error rates are difficult to operate due to a performance gap between the enormous amount of raw capacity when compared to the end-to-end performance which remains available for applications. One of the reasons for that effect is the limited processing power of endstations. Although the performance of those stations is rapidly increasing there is a mismatch between end station capacities and network capacities. This is due to network I/O interfaces, system architectures, host interprocess communication and - last but not least - to inadequate functionality and implementation of transport protocols.

In this contribution we describe different methods which may be applied in order to increase the end to end performance of high speed networks. After illustrating the bottlenecks which may limit the performance we propose different approaches for high performance network interfaces. The results are illustrated by a comparison of implementations for TCP/IP as well as for a lightweight XTP protocol within a SUN OS kernel.

Ralf Steinmetz, IBM ENC, Heidelberg (Germany)

## Media Synchronization

Multimedia synchronization comprises the definition as well as the establishment of temporal relationships among audio, video and other data. The presentation of "in synch" data streams processed at computers equals a human perception in a non-technical environment, this is natural. If data is "out of synch" the human perception tends to identify this error as the data being artificial and annoying. The goal of any multimedia system is to present all data without a synchronization error. The achievement of this goal requires a detailed knowledge on the notion of synchronization errors at the user interface.

This paper outlines the results of a series of experiments leading to a first order guideline for the definition of a synchronization quality of service. The results show that a skew between related data streams may still lead to perceived "in synch" data and, that the notion of a synchronization error highly depends on the related types of media. A quality of Service parameter is derived from this knowledge of synchronization leading to a micromanagement of resources by a network layer protocol.

## **An Internetworking Architecture for Multimedia Communication over Heterogeneous Networks**

Distributed and interactive multimedia computer applications require communication systems transporting audio and video in realtime. In order to guarantee timely delivery of packetized multimedia data three basic components are needed:

- networks capable of providing fixed and low delay communication like Token Ring, FDDI, ISDN and ATM based networks,
- low overhead ("light weight") communication protocols,
- mechanisms to communicate and coordinate quality of service values between terminals, routers, connection and resource managers.

This contribution is primarily concerned with a signalling architecture that is applicable to both, the traditional broadcast type local area networks, as well as switched ATM networks.

To support multimedia communication over LANs interconnected by serial links we have implemented a communication architecture based on the ST II (revised internet stream) protocol. ST II is a network layer protocol that is capable of transporting flow specs (quality of service descriptions) accross heterogeneous networks. Each ST II agent interfaces with various resource managers like a subnetwork bandwidth manager, a router's buffer manager or others. In a switched network resource management is typically handled by a switch controller. Terminals communicate with this controller using signalling protocols like Q.931 for ISDN or various flavors of Q.93B for ATM networks.

The ST II protocol itself can be split in two parts, one being the light weight data part which provides label switching based routing of data and control packets. The second being the SCMP (streams control message protocol) providing connection control functions. Thus SCMP is a signalling protocol in its own rights. As the most important evolving network especially suited for multimedia communication is ATM, it is important to understand how the ATM signalling protocol Q.93B relates to SCMP.

This presentation analyses the functional requirements of a signalling architecture suitable for multimedia communication over interconnected LAN and ATM networks.

Three alternatives for integration of Q.93B and SCMP are compared on the basis of

- functional requirements,
- efficiency of the interworking scheme,
- adaptation to the evolving networking infrastructure.

Finally an architecture solution is presented and discussed.

Wassim Tawbi, Laboratoire MASI, Paris Cedex 05 (France)

### **Management of QoS for multimedia applications**

The contribution addresses QoS issues in distributed multimedia systems and more precisely the management of multimedia applications QoS. A framework that consists of four levels ("users", "programmers", "system", and "network and resources") identifies important QoS aspects that should be considered and studied for distributed multimedia applications. The work focuses on the QoS management at the system level. At this level, an applications' QoS manager (AQOSM), residing at each site, ensures the link between the applications and the available network and local system resources and manages the QoS of these applications. AQOSMs communicate by means of a management protocol called HLQNP (High Level QoS Negotiation Protocol). HLQNP allows QoS negotiation between the applications prior to applications' execution. Dynamic negotiations and renegotiations are also supported by the protocol and deal with any dynamic changes in the applications or any problems in reserved resources. The protocol maintains relevant information for applications QoS management and contributes in the management of the underlying resources.

Phuoc Tran-Gia, Universität Würzburg, Würzburg (Germany)

### **Discrete-time analysis of Usage Parameter Control Functions in ATM Systems**

The design of the User Network Interface (UNI) in accordance with the incorporated Usage Parameter Control (UPC) plays an important role in the current ATM development and standardisation process. Due to the discrete-time nature of ATM cell traffic and the control functions of UNI, queueing models operating in discrete-time domain can be used in a quite direct way.

In this talk, two models with UNI/UPC functions will be presented:

- i. a queueing model to analyse the cell traffic shaping using a spacer where a discrete-time algorithm for the spacer output process is developed and
- ii. a queueing model for the Generic Cell rate Algorithm (GCRA) with which dimensioning aspects of the Cell Delay Variation (CDV) are discussed, in accordance with a versatile discrete-time algorithm.

**Bernd Wolfinger, Universität Hamburg, Hamburg (Germany)**

**Formal description technique to support load modelling for innovative communication systems**

In order to evaluate the performance of computer and communication systems under realistic boundary conditions, adequate load models are urgently required. It is desirable to derive those models directly from load measurements. In this presentation, Petri net based models for users of innovative communication systems are introduced. First, a survey of basic requirements to a formal load description technique is given. A starting point of our technique is an elementary user modeled by a Petri net being an automaton with a time factor. Then basic compositions of elementary users such as sequential, alternative, and iterative compositions are elaborated. Finally, the parallel compositions of elementary users are considered. In order to define the parallel composition of many users with similar behaviour, equivalence relations between users are defined.

**Adam Wolisz, GMD-FOKUS, Berlin 12 (Germany)**

**An integrated approach to communication protocols specification and performance analysis.**

It will be discussed how a properly extended Formal Description Technique can be used as a basis for solving different problems in evaluating communication protocols and comparing alternatives in their implementation. First short characteristic of Timed Interacting Systems - a performance enhanced LOTOS-like formal description Technique - will be presented. TIS supports in comparison with LOTOS a more complex gate semantics: preparing for synchronization requires time, synchronization itself requires time and in addition synchronization may also require the availability of predefined resources. Thus TIS makes it possible to specify not only the protocol under study but also supports adding later some aspects of the execution environment. Afterwards we shall discuss a programming package for deriving simulation runs directly out of TIS specifications. We shall address the use of the obtained in this way performance estimates for optimizing the execution environment. Finally we shall point out our work on methodology to verify deterministic performance bounds for TIS specifications. We shall also comment on our recent experiences in using this approach for simple examples and for the XTP transport protocol.

And finally "the poem":

### **Dagstuhl 9335 (Architecture and Protocols for High Performance Networks)**

We were so proud since we could see  
so many freaks in Germany.  
The experts came - we said hurrah -  
from countries such as Canada.  
East Germans, still more from the western  
part and from UK, Lancaster(n).

All those arrived at castle Dagstuhl  
which sounds in German such as "Kackstuhl".  
That word again (what we admit)  
does nothing mean but "chair of shit".  
This may have been one of the reasons  
why in spite of the autumn season  
some US experts were in favour  
of Williamsburg's tobacco flavour  
since there was held what shouldn't be  
the same workshop by IEEE.  
Nevertheless but natural  
we were quite international.  
Participants from many nations  
were active with some presentations.

**Pierre** and **Doug** liked ATM,  
some others couldn't feel with them.

**André's** request for guarantees  
was taken up by Msieu' **Tawbi**.  
And **Olga's** view was to support this  
as strong as **Andreas Pombortsis**.

The same was said without much dreami'  
by big man **Sunday Lamborghini**,  
pardon - **Domenico Ferrari** -  
with some Italian Larifari.

The talk provided by **Jon Mark**  
did none of us leave in the dark.

The Tango dance was not invented  
but multimedia-presented  
by our well beloved **Ralf Steinmetz**  
as serious as seven dead cats.

**Hans Schlichter** did CSCV  
the reason was not clear to see.  
[I know: CSCW; (dabbeljuh!)]  
I couldn't find a rime for you!]

In order to increase confusion  
**Paul Kuehn** began with the conclusion.

**Budrikis** acted as translator  
in the nearby museum of Wädern. (*Wow!*)  
He left the old DQDB  
and invented DQDT.  
My question: Is it really better  
to change only a single letter  
instead of trying first to sell  
the old products (but heavens hell!)?

I didn't like the talk of **Adam**  
since everybody, man or madam,  
was really forced to stay awake;  
that was a dirty piece of cake!

Another talk - I have to say -  
was somewhat difficult to stay.  
Thus **Helmut** thought: "I wish I had  
a better seat or - best - a bed".  
And as a consequence of thinking  
he continued in downwards sinking.

In his phantastic explanation  
**Heiner** used many abbreviations  
which may have had some sense within it  
but came in rates of twelve per minute!

Both calculations and emphasis  
of **Ivan** were on Dollar basis.

**Eckart's** talk came second last.  
Its title: "Message from the past".  
Thus here we saw less speculations  
than in most other presentations.

Who counts the countries, who the names  
which in Dagstuhl together came. (*Wow square!*)

**Tran-Gia, Heinrichs, Gunningberg,  
Wolfinger, Kiril, Effelsberg.**

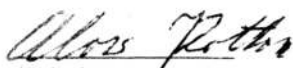
Some others have also been there  
**Baum, Rothermel, Monique Becker.**  
Finally there was more than half a  
side to be seen of **Wilfried Kalfa.**

One of the chairmen was quite lazy.  
He said: "I'm really not so crazy.  
Too many colleagues I have seen  
who have been killed by **A. Danthine**  
or by **Ralf-Guido, Oswald, Doug,**  
not to forget about **Biersáck.**  
My talk - if given here and there -  
will have a risk quite similar. (*Wow!*)  
My strategy: I'd better show 'em  
some kind of nonsense English poem  
which I may obviously present  
when the event comes to its end".

#### Conclusion:

We liked the friendly atmosphere,  
we liked the wine, we liked the beer.  
Thus everybody, what the heck,  
will gladly say: "And I'll come back!!!"

Dagstuhl, September 3, 1993

  
(Alois Potton)

alois (pardon: 'alias') Otto Spaniol



## Dagstuhl-Seminar 9335:

**Dieter Baum**  
Universität Trier  
Fachbereich IV - Informatik  
54286 Trier  
Germany  
baum@uni-trier.de  
tel.: +49-651-201-2849 /2848  
fax: +49-651-201-3923

**Monique Becker**  
INT  
Les Epinettes  
9 rue Charles Fourier  
F-91011 Evry Cedex  
France  
mbecker@int-evry.fr  
tel.: +33-160-7647-81  
fax: +33-160-7647-80

**Ernst Biersack**  
Institut Eurecom  
Sophia Antipolis  
2229 route des Crêtes  
F-06560 Valbonne  
France  
erbi@eurecom.fr  
tel.: +33-9300-2611  
fax: +33-9300-2627

**Kiril Boyanov**  
Bulgarian Academy of Sciences  
Center of Informatics &  
Computer Technology  
Acad. G. Bontchev str./Block 25a  
Bulgaria  
tel.: 00359-270-84-94 ; 00352-2-70-32-60  
Boyanov@bgearn.bitnet  
fax: 00359-2-70-72-73

**Pierre Boyer**  
CNET - Lannion  
LAA/RSM/SIM  
BP 40  
F-22301 Lannion Cedex  
France  
boyerp@lannion.cnet.fr  
tel.: +33-96052239  
fax: +33-96 482531

**Z. L. Budrikis**  
QPSX Communications LTD  
33 Richardson Street  
West Perth 6005  
Australia  
tel.: +61-9-324 1641  
fax: +61-9-324 1642

## List of Participants

**Olga Casals Torres**  
Universidad Politécnica de Cataluña  
Dept. Arquitectura de Computadores  
Campus Norte / Modulo D-6  
c/Gran Capitan s/n  
E-08071 Barcelona  
Spain  
olga@ac.upc.es  
tel.: +34-3401-6985  
fax: +34-3401-7055

**Andre Danthine**  
Université de Liège  
Institute Montefiore B28  
B-4000 Liège  
Belgium  
danthine@vm1.ulg.ac.be

**Oswald Drobnik**  
Universität Frankfurt  
Fachbereich Telematik (20)  
Robert-Mayer-Str. 11-15  
60054 Frankfurt am Main  
Germany  
dro@informatik.uni-frankfurt.de  
tel.: +49-69-798-8362  
fax: +49-69-798-8353

**Wolfgang Effelsberg**  
Universität Mannheim  
Lehrstuhl Praktische Informatik IV  
Seminargebäude A5  
68131 Mannheim  
Germany  
effelsberg@pi4.informatik.uni-mannheim.de  
tel.: +49-621-292-3131  
fax: +49-621-292-5745

**Ivan Fellner**  
Slovak Academy of Sciences  
Institute of Automation and Communica-  
tion  
UAKOM  
Severna 5  
97400 Banska Bystrica  
Slovak Republic  
fellner@uakom.sk  
tel.: +42-88-53207  
fax: +42-88-52364

**Francesco Ferrari**  
University of California at Berkeley  
Computer Science Division  
EECS Department  
551 Evans Hall  
Berkeley CA 94720  
USA  
ferrari@cs.berkeley.edu  
tel.: +1-510-642-3806 / 643-9153  
fax: +1-510-643-7684

**Per Gunningberg**  
Swedish Institute of Computer Science  
Box 1263  
S-16428 Stockholm-Kista  
Sweden  
per@sics.se  
tel.: +46 8 752 1526  
fax: +46 8 751 7230

**Bernd Heinrichs**  
RWTH Aachen  
Lehrstuhl für Informatik IV  
Ahornstr. 55  
52056 Aachen  
Germany  
heibe@informatik.rwth-aachen.de  
tel.: +49-241-80-21400  
fax: +49-241-80-21429

**Ralf Guido Herrtwich**  
IBM European Network Center  
Vangerowstr. 18  
69115 Heidelberg  
Germany  
rgh@dhdibm1.bitnet  
tel.: +49-6221-59-4218  
fax: +49-6221-59-3300

**Paul J. Kühn**  
Universität Stuttgart  
IND  
Seidenstr. 36  
70174 Stuttgart  
Germany  
kuehn@nvdv.e-technik.uni-stuttgart.dbp.de  
tel.: +49-711-121-2479  
fax: +49-711-121-2477

**Winfried Kalfa**  
Technische Universität Chemnitz  
Fachbereich Informatik  
Postfach 964  
09009 Chemnitz  
Germany  
kalfa@informatik.tu-chemnitz.de  
tel.: +49-371-668-430 / 387  
fax: +49-371-668-628

**Helmut Löffler**  
TU Dresden  
FB Informatik  
Lehrstuhl Rechnernetze  
Mommstr. 13  
01062 Dresden  
Germany  
loeffler@freia.inf.tu-dresden.de  
tel.: +49-351-4575-261  
fax: +49-351-4575-335

**Jon W. Mark**  
University of Waterloo  
Dept. of Electrical and Computer Engineering  
Waterloo Ontario N2L 3G1  
Canada  
jwmark@ccng.waterloo.edu  
or jwmark@bbcr.waterloo.edu  
tel.: +1-519-885-1211 / ext. 2532  
fax: +1-519-746-3077

**Andreas S. Pombortsis**  
Aristotle University - Thessaloniki  
Department of Computer Science  
GR-54006 Thessaloniki  
Greece  
cadz04@grtheun.bitnet  
and pombortsis@olymp.ccf.auth.gr  
tel.: +30-51-99-8045  
fax: +30-51-99-8102 or 20 6138

**Eckart Raubold**  
GMD Darmstadt  
12  
Rheinstr. 75  
64295 Darmstadt  
Germany  
raubold@darmstadt.gmd.de  
tel.: +49-6151-869281  
fax: +49-6151-869224

**Kurt Rothermel**  
Universität Stuttgart  
Institut für Parallele und Verteilte  
Hochleistungsrechner (IPVR)  
Breitwiesenstraße 20-22  
D-70565 Stuttgart  
Germany  
rothermel@informatik.uni-stuttgart.de  
tel.: +49-711-7816-434  
fax: +49-711-7816-424

**Hans Schlichter**  
TU München  
Institut für Informatik  
Postfach 20 24 20  
80290 München  
Germany  
schlicht@informatik.tu-muenchen.de  
tel.: +49-89-450552-20  
fax: +49-89-450552-22

**Doug Shepherd**  
Lancaster University  
Department of Computing  
Lancaster LA1 4YR  
Great Britain  
doug@comp.lancs.ac.uk  
tel.: +44-524-59-3827  
fax: +44-524-84-4011

**Otto Spaniol**  
RWTH Aachen  
Lehrstuhl für Informatik IV  
Ahornstr. 55  
52056 Aachen  
Germany  
spaniol@informatik.rwth-aachen.de  
tel.: +49-241-80-21400  
fax: +49-241-80-21429

**Heinrich J. Stüttgen**  
IBM European Network Center  
Vangerowstr. 18  
69115 Heidelberg  
Germany  
stuttgen@dhdibm1.bitnet  
tel.: +49-6221-594337  
fax: +49-6221-593300

**Ralf Steinmetz**  
IBM European Network Center  
Vangerowstr. 18  
69115 Heidelberg  
Germany  
steinmet@dhdibm1.bitnet  
tel.: +49-6221-594280  
fax: +49-6221-593300

**Wassim Tawbi**  
Laboratoire MASI  
Institute Blaise Pascal  
T65-66 Bureau 110  
4, Place Jussieu,  
F-75252 Paris Cedex 05  
France  
tel.: +33-1-4427-7128  
fax: +33-1-4427-6286  
tawbi@masi.ibp.fr

**Phuoc Tran-Gia**  
Universität Würzburg  
Lehrstuhl für Informatik III  
Am Hubland  
97074 Würzburg  
Germany  
trangia@informatik.uni-wuerzburg.de  
tel.: +49-931-888-55 09  
fax: +49-931-888-4601

**Bernd Wolfinger**  
Universität Hamburg  
FB Informatik  
FG Rechnerorganisation  
Vogt-Koelln-Str. 30  
22527 Hamburg  
Germany  
wolfinger@rz.informatik.uni-hamburg.dbp.de  
tel.: +49-40-54715-423  
fax: +49-40-54715-328/ 422(tel. secretary)

**Adam Wolisz**  
GMD-FOKUS  
Hardenbergplatz 2  
W-1000 Berlin 12  
Germany  
wolisz@fokus.berlin.gmd.d400.de  
fax: +49-30-254-99-202

## **Zuletzt erschienene und geplante Titel:**

3. Becker, R. Bryant, Ch. Meinel (editors):  
Computer Aided Design and Test , Dagstuhl-Seminar-Report; 56; 15.02.-19.02.93 (9307)
- iv. Pinkal, R. Scha, L. Schubert (editors):  
Semantic Formalisms in Natural Language Processing, Dagstuhl-Seminar-Report; 57; 23.02.-26.02.93 (9308)
- iv. Bibel, K. Furukawa, M. Stickel (editors):  
Deduction , Dagstuhl-Seminar-Report; 58; 08.03.-12.03.93 (9310)
- H. Alt, B. Chazelle, E. Welzl (editors):  
Computational Geometry, Dagstuhl-Seminar-Report; 59; 22.03.-26.03.93 (9312)
- H. Kamp, J. Pustejovsky (editors):  
Universals in the Lexicon: At the Intersection of Lexical Semantic Theories, Dagstuhl-Seminar-Report; 60; 29.03.-02.04.93 (9313)
- W. Strasser, F. Wahl (editors):  
Graphics & Robotics, Dagstuhl-Seminar-Report; 61; 19.04.-22.04.93 (9316)
- Beerli, A. Heuer, G. Saake, S. Urban (editors):  
Formal Aspects of Object Base Dynamics , Dagstuhl-Seminar-Report; 62; 26.04.-30.04.93 (9317)
- R. V. Book, E. Pednault, D. Wotschke (editors):  
Descriptive Complexity, Dagstuhl-Seminar-Report; 63; 03.05.-07.05.93 (9318)
- H.-D. Ehrig, F. von Henke, J. Meseguer, M. Wirsing (editors):  
Specification and Semantics, Dagstuhl-Seminar-Report; 64; 24.05.-28.05.93 (9321)
- M. Droste, Y. Gurevich (editors):  
Semantics of Programming Languages and Algebra, Dagstuhl-Seminar-Report; 65; 07.06.-11.06.93 (9323)
- Th. Lengauer, P. Quinton, Y. Robert, L. Thiele (editors):  
Parallelization Techniques for Uniform Algorithms, Dagstuhl-Seminar-Report; 66; 21.06.-25.06.93 (9325)
- G. Farin, H. Hagen, H. Noltemeier (editors):  
Geometric Modelling, Dagstuhl-Seminar-Report; 67; 28.06.-02.07.93 (9326)
- Flajolet, R. Kemp, H. Prodinger (editors):  
"Average-Case"-Analysis of Algorithms, Dagstuhl-Seminar-Report; 68; 12.07.-16.07.93 (9328)
- W. Gray, A.M. Pitts, K. Sieber (editors):  
Interactions between Category Theory and Computer Science, Dagstuhl-Seminar-Report; 69; 19.07.-23.07.93 (9329)
- D. Gabbay, H.-J. Ohlbach (editors):  
Automated Practical Reasoning and Argumentation, Dagstuhl-Seminar-Report; 70; 23.08.-27.08.93 (9334)
- A. Danthine , W. Effelsberg, O. Spaniol, (editors):  
Architecture and Protocols for High-Speed Networks, Dagstuhl-Seminar-Report; 71; 30.08.-03.09.93 (9335)
- R. Cole, E. W. Mayr, F. Meyer a.d. Heide (editors):  
Parallel and Distributed Algorithms, Dagstuhl-Seminar-Report; 72; 13.09.-17.09.93 (9337)
- V. Marek, A. Nerode, P.H. Schmitt (editors):  
Non-Classical Logics in Computer Science, Dagstuhl-Seminar-Report; 73; 20.-24.09.93 (9338)
- A. Odlyzko, C.P. Schnorr, A. Shamir (editors):  
Cryptography, Dagstuhl-Seminar-Report; 74; 27.09.-01.10.93 (9339)