The Tenet Group

University of California, Berkeley
and International Computer Science Institute
October 20, 1992

Introduction

The Tenet Group was formed in September of 1989 at the University of California at Berkeley and the International Computer Science Institute (ICSI, also in Berkeley) to perform research in high-speed computer networking. The focus of the group is on the development of distributed network control and real-time communication services, and support for continuous-media applications.

The Tenet Group is participating in two major multi-institution research projects: BLANCA and Sequoia 2000. The BLANCA project is part of the Gigabit Testbed Initiative, which is sponsored by the Corporation for National Research Initiatives and supported by the National Science Foundation and by the Defense Advanced Research Projects Agency. The BLANCA project officially started in May 1990 and currently includes AT&T Bell Laboratories, the universities of California, Illinois, and Wisconsin, Pacific Bell, and a few research institutes, among which is ICSI. The goal of BLANCA is to design and build two wide-area ATM-based testbeds called Xunet 2 (at 45 Mbps) and Xunet 3 (at 622 Mbps), and demonstrate on them several applications requiring high-speed networking. Most of applications are of the scientific visualization type, and require the transmission of image sequences and/or video streams. The sequences or streams are to be browsed, stopped, and played forward or backward at speeds slower or faster than the nominal one, by scientists at sites far from those where the sequences or streams are generated or stored.

Sequoia 2000, the second major initiative in which the Tenet Group is participating, aims at building high-speed storage and communication facilities for global-change scientists at remote locations to visualize and browse sequences of digitally represented satellite maps. Project Sequoia 2000, the sequel to Project Athena, was awarded in July 1991 by the Digital Equipment Corporation to the University of California system; its network design and research activities are taking place at U.C. Berkeley and U.C. San Diego. The Sequoia network connects at T1 speed the Berkeley campus to the Santa Barbara, Los Angeles, and San Diego campuses, as well as to the Scripps Institution of Oceanography in San Diego and the State's Department of Water Resources near Sacramento. The network's bandwidth will be increased to 45 Mbps in March 1993.

The contributions of the Tenet Group to both projects are concerned primarily with the technological foundations, rather than with the end-user applications, and can be classified into three partially overlapping areas: (a) real-time communication services, (b) continuous-media networking technologies, and (c) network control policies.

1. Real-Time Communication Services

The principles of the Tenet approach. The strategy followed by the Tenet Group consists of designing a set of algorithms through which a network or a class of networks can support real-time (i.e., guaranteed-performance) communication, then designing protocols that realize those algorithms, then implementing those protocols for one or more specific networking environments, then experimenting with these implementations under realistic applications, and finally using the results of the experiments to design the next set of algorithms, protocols, and so on. At the basis of all of this work there has been, since the outset, a philosophy that was never clearly distinguished from the simplifications and compromises that had to be made to facilitate its initial incarnations. **Domenico Ferrari, Anindo Banerjea**, and **Hui Zhang** have tried in a recent report to separate this philosophy from the non-essential aspects of the initial Tenet schemes and of the first Tenet protocol suite. This clarification effort is intended to lead to a better understanding of the group's work and its relationships to the work done by other research groups.

1.1. The Prototype Tenet Protocol Suite

The Real-time Channel Administration Protocol. **Anindo Banerjea** and **Bruce Mah** have recently completed a prototype of the Real-time Channel Administration Protocol (RCAP). This software performs establishment, teardown, and management of real-time channels at the network and transport layers. Although the RCAP software has been completed and tested, it must be configured for proper operation in the network where it is to be used. In particular, the instance of the RCAP software running on each the various network nodes needs some characterization of that node's performance. Examples of such performance parameters include the link speeds of each of the attached network interfaces and packet processing delays incurred in the node itself. Towards this end, Banerjea and Mah are now working with **Hui Zhang** and **Tom Fisher** to determine the values of the necessary parameters, so that RCAP can be run on the Tenet Group's experimental FDDI testbed. The information and techniques gained during this effort will aide the determination of these parameters for different networks (for example, Xunet 3) where RCAP will eventually be installed.

The Real-time Message Transport Protocol and the Real-Time Internet Protocol. **Hui Zhang** has implemented RMTP/RTIP within Ultrix on DECstation 5000's and within HP/UX on HP9000/700's, and is in the process of moving the implementation to the Sun 4/280 and SGI IRIS-4D workstations. The RMTP/RTIP protocol stack coexists with the TCP/IP protocol stack in the kernel, and provides the familiar socket interface to communication clients. Hui Zhang and **Tom Fisher** are measuring the performance of the RTIP and RMTP implementation. The measurements are performed in a local testbed consisting of a couple of DECstation 5000's connected by an FDDI ring. Preliminary results show that, under no other loads, the throughputs and delays obtained by using RMTP/RTIP are comparable to those obtained by using raw IP or UDP/IP without checksums. When extra load (more IP traffic) is introduced, however, the performance of raw IP and UDP/IP declines considerably, while the performance of RMTP/RTIP is almost unaffected. Another set of measurements are intended to identify the bottlenecks in the implementation and the architectural defects with respect to high speed networking in the current DECstation/Ultrix environment.

The Continuous-Media Transport Protocol. The Continuous-Media Transport Protocol, CMTP, is oriented towards continuous-media applications. Such applications are characterized by periodic frames, possibly of varying sizes (as in the case of compressed video). A prototype of CMTP, completed by **Mark Moran** and **Amit Gupta**, is currently operational on the local testbed, although it has yet to be tested in concert with the other real-time protocols. Moran and Gupta have written a paper with **Bernd Wolfinger** and **Domenico Ferrari** that briefly describes the Continuous Media Transport Service (CMTS), the interesting problems addressed in its prototype implementation, and simulation results that verify some of the gains in efficiency over message-based transport services for delivering continuous-media data. The proposals described in the paper will be studied via simulation experiments.

Porting the Tenet suite to the Xunet 3 testbed. Work continues on facilities for the segment of the Xunet 3 testbed between the University of California at Berkeley campus and the Lawrence Berkeley Laboratory. The fiber-optic links between the two sites have been installed and tested; researchers at UC Berkeley and LBL have begun experimenting with using some of the fibers to carry FDDI network traffic. Various pieces of HIPPI network equipment have been procured and installed. A HIPPI Serial Extender from Broadband Communications Products has been tested over short distances in the laboratory. Two host interfaces for VME-based machines (in this case, Sun 4 workstations and SPARCstations) have been installed and tested in loopback configurations. **Bruce Mah** and **Tom Fisher** have begun work on the port of the real-time protocol suite to the HIPPI network. This consists of producing a working SunOS kernel from source code and then adapting and integrating into it the existing real-time protocol code. Of particular interest are the changes in the operating system and machine architecture (the prototype suite was developed on DECstation 5000's running DEC Ultrix, while the Xunet 3 hosts are SPARCstations and Sun 4's running SunOS). It is hoped that these differences, with respect to the protocol suite, will be minor (the networking code in the two operating systems is almost the same).

Real-time and distributed operating systems. **Newton Faller** is investigating how real-time operating systems similar to UNIX could provide a suitable environment for continuous-media applications. **Tom Fisher** has modified the ULTRIX kernel in several ways, making the kernel more predictably responsive to both internal and external events. The changes include the addition of preemption points and a scheme for data structure locking. He has also begun kernel modifications to implement a mechanism that provides a mapping between user and kernel address spaces for the purpose of eliminating the copying of data that takes place during network communication between these addresses spaces on both the send and the receive side. **Fred Templin**, a Digital Equipment Corporation employee, has been acting as an operating systems and network I/O systems consultant for the Tenet researchers. He is also investigating

WAN topologies, and porting the Tenet protocol suite to the Sequoia 2000 network.

1.2. The Next Tenet Scheme and Protocol Suite

Multicasting. Research has been started to serve as the foundation for the design of a new version of the Tenet real-time protocol suite. The current Tenet protocol suite implements a simple real-time communication abstraction: the simplex unicast channel. However, a large set of real-time communication applications require that the network support multicast communications, i.e., communication from a single source to multiple destinations. Examples of multicast applications are teleconferencing, distributed classrooms, and tele-shopping. To support these services, the abstraction must be extended to include multicast communication. Several challenging problems, however, have to be attacked and solved before the simplex multicast channel can be taken as the basic abstraction of the new suite. **Amit Gupta** and **Colin Parris** are investigating the new routing mechanism that needs to be devised for multicast communication.

A model for client-network interactions. **Jean Ramaekers** and **Giorgio Ventre** have developed a model of client-network interactions for the second version of the Tenet real-time protocol suite. The model includes new mechanisms for the establishment and runtime management of real-time connections; it is expected to reduce the complexity and the time required to establish a real-time connection, and to increase network utilization. New classes of real-time communication clients with flexible quality-of-service requirements have been introduced as a basis for the new client-network interaction model. In two recent reports, Ramaekers and Ventre presented a new mechanism for the establishment of real-time connections in a quality-of-service network developed for the Tenet real-time protocol suite. The solution proposed is based on three ideas. The first is to increase the amount of data that a client obtains from the network when the request for a real-time channel is rejected. The second is an improved procedure for channel establishment. The third idea is the introduction of more flexible types of real-time communication service. Ramaekers and Ventre, together with **Domenico Ferrari**, are now designing a new scheme for the dynamic modification of the real-time channel characteristics. In addition to improving the utilization of network resources, this scheme should represent a first step toward the implementation of efficient mechanisms for load balancing and fault tolerance.

Dynamic Channel Management (DCM). **Colin Parris** and **Hui Zhang** have designed a Dynamic Channel Management (DCM) scheme that will be used to increase the availability and flexibility of a real-time network while maintaining the integrity of all performance contracts made to the clients with real-time channels. One method for achieving this goal requires that the network be able dynamically to modify the performance parameters and traffic characteristics of real-time channels as well as the route traversed by the channel. These modifications will be done by the DCM scheme in a manner that is non-disruptive to the client. In most cases the modifications will be transparent to the client, unless the network deems it necessary for the client to know. This need to know is determined on the basis of the DCM policy that is in effect. DCM consists of the DCM scheme and the DCM policies. The DCM scheme is the collection of algorithms that permits the network to dynamically modify the channel parameters and routes. The modification of a channel follows a procedure whereby a real-time channel with the new performance parameters (referred to as the alternate channel) is established, the client's traffic is moved from the current real-time channel (referred to as the primary channel) to the alternate channel, and the primary channel is removed. A DCM policy is the collection of rules that determine whether a real-time channel is to be modified and the selected parameters that are to be modified. These rules may examine the network or client state data to determine whether the modification should take place. DCM policies are entirely dependent on the application, and are independent of the management algorithms. The DCM scheme consists of essentially two algorithms: (1) a channel administration algorithm, which determines whether a real-time channel can be accepted along a specified alternate route and, if so, reserves resources along the alternate route so that all of the client's traffic and performance requirements are met; the algorithm also ensures that on links that are common between the primary and alternate routes there is no duplication of resources reserved; (2) a channel transition algorithm, which ensures that the transition from the primary to the alternate channel does not violate the client's performance contract. This is especially important in the case of the delay and delay jitter bounds guaranteed in that contract. A third algorithm that is part of the DCM scheme, in a supportive rather than essential capacity, is a real-time channel routing algorithm that determines the most probable route that would satisfy the client's traffic and performance requirements. This routing algorithm considers the resources currently reserved for the primary channel when determining the alternate route. These algorithms will be incorporated into the next Tenet scheme.

Service disciplines for integrated-services networks. Hui Zhang is studying issues and tradeoffs in the design of service disciplines for integrated-services networks. The first step of the research was a comparison study of several recently proposed service disciplines, including Virtual Clock, Fair Queueing, Delay-Earliest-Due-Date, Jitter-Earliest-Due-Date, Stop-and-Go and Hierarchical Round Robin. The study examined the similarities and differences among the disciplines, and showed why each discipline can or cannot provide bandwidth, delay, and delay jitter guarantees. A follow-up study then proposed a new service discipline called Rate-Controlled Static-Priority (RCSP) queueing. RCSP overcomes the limitations of the previously proposed solutions and has such desired features as bounded delay, bounded delay jitter, decoupling of delay and bandwidth allocation, evenly distributed buffer space reservation to prevent packet loss, and easy implementation in high-speed switches. The key idea of RCSP is to separate rate control and delay control in the design of a server: the rate controller does not only allocate bandwidth, but also isolates the interactions between adjacent switches, and controls jitter; the scheduler services packets with different priorities to meet a variety of delay requirements. Generalizing this idea by having combinations of different rate controllers and schedulers, a class of service disciplines is obtained, of which RCSP is an instance. Currently, Zhang is (a) designing with **Massimo Maresca**, a former Tenet Group member who has now returned to the University of Genoa, Italy, a VLSI implementation of RCSP to demonstrate that RCSP can be implemented efficiently, (b) developing deterministic and probabilistic models to analyze the performance of various disciplines, and (c) simulating various disciplines to determine how best-effort traffic is affected by real-time traffic under each discipline.

A hierarchical structure for real-time channels. Guaranteed-quality services are a key aspect of the Tenet approach to real-time communication. Qualities are defined as characteristic parameters of real-time channels as maintained by the Real-time Channel Administration Protocol (RCAP). However, the notion of Quality of Service (QoS) is subject to an end-to-end argument. Certain applications, for example those involving multipoint and multimedia conversations, require more complicated channel setups than those currently supported by RCAP. **Clemens Szyperski** is currently investigating a model called Berkeley Structured Channels (BSCs). A BSC structures and presents all QoS parameter settings to all layers of the system. It is hoped that this will allow the next Tenet scheme to avoid the "layer ignorance" problem encountered in deeply layered protocol stacks. The task includes isolating further relevant QoS dimensions and selecting useful parameters that are both sufficiently independent and implementable.

Dynamic bandwidth control for real-time channels. **Makiko Yoshida** is studying dynamic bandwidth control methods by which a network can provide efficient service to bursty traffic requiring large amounts of bandwidth in ATM environments. This kind of bandwidth management would be useful, for example, in large-size file transfers. Dynamic bandwidth control algorithms are being investigated in the context of the Real-time Channel Administration Protocol (RCAP). Some experiments will be performed on high speed networking testbeds to evaluate the effects of dynamic bandwidth control.

A simulator for networks offering real-time services. **Eckhardt Holz, Edward Knightly**, and **Giorgio Ventre** have designed and implemented an object oriented simulator for real-time communication networks and protocols. This simulator has two main features that distinguish it from the vast number of existing network simulators. The first is the ability to do quality-of-service simulation with an emphasis on multimedia applications. The main goal of the real-time network simulator is to show the effect of network management decisions on the traffic to verify both the correctness and efficiency of resource allocation schemes. In addition, work is in progress to integrate the simulator with multimedia devices and applications. For example, the simulator supports the sending of actual video packets across the simulated network. The simulator will show a "before" and an "after" picture of a video frame (or sequence of frames) in the network to facilitate the quantitative evaluation of protocols. Second, the simulator is object-oriented. This feature effectively exploits the modular and hierarchical nature of communication network protocols, and facilitates their rapid prototyping and testing.

2. Continuous-Media Network Technologies

Network performance requirements of image browsing applications. **Colin Parris, Hui Zhang**, and **Giorgio Ventre** are studying the communication needs of a client browsing high density, lossless, still images. In this application, the client cannot tolerate any loss of frames, although some marginal compression may be permitted. One of the many situations in which this occurs is in meteorological studies, where scientists wish to examine the changes in cloud formations on large satellite images, and must view all of the frames in order to observe even the minutest changes. It is assumed that the client may cause an increase by a factor *K* in the throughput it receives from the network when it requests an enhancement to the quality of service of its real-time channel. This request is received by the DCM algorithms, which determine whether there are sufficient resources in the network to support the request. If so, an alternate route is set up (with the higher throughput bound but the same delay and loss parameters), and a transition is made from the primary to this alternate route. The new enhanced-QOS real-time channel can then be maintained for the remaining lifetime of the channel or for a client/network-specified duration. This increase in throughput will facilitate fast-forwarding or fast-rewinding of the still images. A control channel from the receiver to the sender is implemented as a real-time channel whose performance parameters provide response times of the same order as those experienced in the use of a commercial video recorder.

Network support for multi-party communication. **Mark Moran** is investigating the support required for multi-party, interactive multimedia (MIM) communication between distributed parties (e.g., computer supported collaborative work or CSCW). This work has involved gaining experience with a commercial desktop conferencing system (DECspin), and experimenting with a platform for development of multimedia applications (XMedia for DECstations). From this experience and an extensive literature search, he has characterized the requirements of canonical MIM applications and has developed a coherent design strategy for providing system and network support for MIM communication. Areas of support include time-based media, dynamic access to continuous-media I/O devices and compression modules, and real-time, multicasting communication services. This work has been done in collaboration with **Riccardo Gusella** at Hewlett Packard Laboratories.

3. Network Control Policies

Multiplexing policies at the entrance of a wide-area network. **Ramón Cáceres** has evaluated different policies for multiplexing application-level datagram streams, or conversations, at the entrance to wide-area virtual circuit networks. Datagrams are attractive for end-to-end internetwork applications, while virtual circuits hold many advantages for the provider of integrated-services wide-area networks. When a datagram arrives at the entrance to a wide-area network, it must be routed onto an appropriate virtual circuit. It is important to provide each conversation with good performance while making efficient use of network resources. This multiplexing problem has arisen in the past and reappears with the advent of Asynchronous Transfer Mode (ATM) networks. Cáceres studied also how to map a set of datagram streams onto a possibly smaller set of virtual circuits and how to choose the queueing discipline for scheduling datagrams onto these virtual circuits. His results show that the network should separate traffic types by establishing at least one virtual circuit for each type of traffic flowing between two network points of presence; conversations of the same type can share a virtual circuit. In addition, the network should use a round-robin queueing discipline to service conversations sharing a virtual circuit, and virtual circuits sharing a communication line. This policy combines the performance advantages of allocating dedicated virtual circuits to each conversation and the resource advantages of sharing fewer virtual circuits.

The members of the Tenet Group, and their recent publications since the beginning of 1991, are listed below. Unless otherwise noted, the publications are available by electronic mail from the Tenet file server **(file-server@tenet.berkeley.edu)**. Instructions on how to use the file server may be obtained by sending it a mail message with the one-word subject: *Help*. The same publications are also available by anonymous ftp to tenet.berkeley.edu. Inquiries for further information should be directed to the Tenet Group leader, Professor **Domenico Ferrari**, Computer Science Division, 571 Evans Hall, University of California, Berkeley CA 94720. Electronic mail may be sent to **ferrari@cs.berkeley.edu**.

The current members of the Tenet Group are the following graduate students of the Department of Electrical Engineering and Computer Science of U. C. Berkeley: **Anindo Banerjea**, **Ramón Cáceres**, **Amit Gupta**, **Edward Knightly**, **Bruce Mah**, **Mark Moran**, **Colin Parris**, and **Hui Zhang**; the programmer **Tom Fisher**; the following post-doctoral visiting scholars and researchers: **Newton Faller**, from the Núcleo de Computação Eletrônica and the Instituto de Matemática of the Universidade Federal do Rio de Janeiro, Rio de Janeiro, Brazil; **Eckhardt Holz**, from the Fachbereich Informatik of Humboldt-University, Berlin, Germany; **Jean Ramaekers**, from the Institut d'Informatique of the Facultés Universitaires Notre-Dame de la Paix, Namur, Belgium; **Wolfgang Schroeder-Preikschat**, from the German National Research

Center for Computer Science (GMD-FIRST), Berlin, Germany; **Clemens Szyperski**, from the Swiss Federal Institute of Technology (ETH), Zurich, Switzerland; **Fred Templin**, from the External Research Program of Digital Equipment Corporation, Maynard, Massachusetts, U.S.A.; **Giorgio Ventre**, from the Dipartimento di Informatica e Sistemistica of the Universitá di Napoli, Naples, Italy; **Makiko Yoshida**, from the C&C Systems Research Laboratories, NEC Corp., Kawasaki, Japan; and the leader of the Tenet Group, **Domenico Ferrari**, who is with the Department of Electrical Engineering and Computer Science of U. C. Berkeley and the International Computer Science Institute.

List of Publications (January 1991 to present)

A. Banerjea and S. Keshav, ‘‘Queuing Delays in Rate Controlled Networks’’, Tech. Rept. TR-92-015, International Computer Science Institute, Berkeley, CA, February 1992.

A. Banerjea and B. Mah, ‘‘The Real-Time Channel Administration Protocol’’, *Proc. Second International Workshop on Network and Operating System Support for Digital Audio and Video*, Springer-Verlag, Heidelberg, Germany, pp. 160-170, November 1991.

R. Cáceres, ‘‘Efficiency of Asynchronous Transfer Mode Networks in Transporting Wide-Area Data Traffic’’, Tech. Rept. TR-91-043, International Computer Science Institute, Berkeley, CA, July 1991; also submitted to *Computer Networks and ISDN Systems*, December 1991.

R. Cáceres, ‘‘Multiplexing Datagrams on Virtual Circuits at the Entrance to a Wide-Area Network’’, Ph.D. Dissertation, University of California, Berkeley, CA, October 1992. (not available from the file server)

R. Cáceres, D. P. Danzig, S. Jamin, and D. J. Mitzel, ‘‘Characteristics of Wide-Area TCP/IP Conversations’’, *Proc. ACM SIGCOMM ’91*, Zurich, Switzerland, pp. 101-112, September 1991.

P. Danzig, S. Jamin, R. Cáceres, D. Mitzel, and D. Estrin, ‘‘An Empirical Workload Model for Driving Wide-Area TCP/IP Network Simulations’’, to appear in the *Journal of Internetworking Research and Experience*, 1992. (not available from the file server)

N. Faller, ‘‘Measuring the Latency Time of Real-Time UNIX-like Operating Systems’’, Tech. Rept. TR-92-037, International Computer Science Institute, Berkeley, CA, June 1992.

D. Ferrari, ‘‘Design and Applications of a Delay Jitter Control Scheme for Packet-Switching Internetworks’’, *Proc. Second International Workshop on Network and Operating System Support for Digital Audio and Video*, Springer-Verlag, Heidelberg, Germany, pp. 72-83, November 1991; also in *Computer Communications*, vol. 15, n. 6, pp. 367-373, July-August 1992.

D. Ferrari, ‘‘Distributed Delay Jitter Control in Packet-Switching Internetworks’’, Tech. Rept. TR-91-056, International Computer Science Institute, Berkeley, CA, October 1991; also to appear in the *Journal of Internetworking Practice and Experience*, 1992.

D. Ferrari, ‘‘Real-Time Communication in an Internetwork’’, Tech. Rept. TR-92-001, International Computer Science Institute, Berkeley, January 1992; also to appear in the *Journal of High Speed Networks*, 1992.

D. Ferrari, A. Banerjea, and H. Zhang, ‘‘Network Support for Multimedia - A Discussion of the Tenet Approach’’, Tech. Rept. TR-92-072, International Computer Science Institute, Berkeley, CA, October 1992.

D. Ferrari, A. Gupta, M. Moran, and B. Wolfinger, ‘‘A Continuous Media Communication Service and its Implementation’’, *Proc. GLOBECOM ’92*, Orlando, Florida, 8A.02., December 1992.

D. Ferrari, J. Ramaekers, and G. Ventre, ‘‘Client-Network Interactions in Quality of Service Communication Environments’’, *Proc. 4th IFIP Conf. on High Performance Networking*, Liege, Belgium, December 1992.

T. Fisher, ‘‘Real-Time Scheduling Support in ULTRIX-4.2 for Multimedia Communication’’, *Proc. Third International Workshop on Network and Operating System Support for Digital Audio and Video*, San Diego, CA, November 1992.

M. Gilge, ‘‘Visuelle Kommunikation ueber paketvermittelte Netze mittels netzwerkadaptiver Bildcodierung’’, 4. Dortmunder Fernsehseminar, Dortmund, Germany, October 1991. (not available from the file server)

M. Gilge and R. Gusella, ‘‘Motion Video Coding for Packet Switching Networks -- An Integrated

Approach'', *SPIE Visual Communications and Image Processing '91*, Boston, Mass. USA, November, 1991.

R. Gusella, ''Characterizing the Variability of Arrival Processes with Indices of Dispersion'', Tech. Rept. TR-90-051, International Computer Science Institute, Berkeley, CA, September 1990; also in the *IEEE Journal on Selected Areas in Communications*, vol. 9, n. 2, pp. 203-211, February 1991.

S. Keshav, ''Congestion Control in Computer Networks'', Ph.D. Dissertation, University of California at Berkeley, September 1991; also, Tech. Rept. UCB/CSD 91/649, University of California, Berkeley, CA, September 1991. (not available from the file server)

S. Keshav, ''On the Efficient Implementation of Fair Queueing'', *Journal of Internetworking Research and Experience*, vol. 2, pp. 157-173, September 1991.

S. Keshav, ''A Control-Theoretic Approach to Flow Control'', Tech. Rept. TR-91-015, International Computer Science Institute, Berkeley, CA, March 1991; also in *Proc. ACM SIGCOMM '91*, Zurich, Switzerland, pp. 3-15, September 1991.

C. Lowery, ''Protocols for Providing Performance Guarantees in a Packet-Switching Internet'', Tech. Rept. TR-91-002, International Computer Science Institute, Berkeley, CA, January 1991.

M. Moran and R. Gusella, ''System Support for Efficient Dynamically-Configurable Multi-Party Interactive Multimedia Applications'', *Proc. Third International Workshop on Network and Operating Systems Support for Digital Audio and Video,* San Diego, CA, November 1992.

C. Parris and D. Ferrari, ''A Resource Based Pricing Policy for Real-Time Channels in a Packet-Switching Network'', Tech. Rept. TR-92-018, International Computer Science Institute, Berkeley, CA, March 1992.

C. Parris, S. Keshav, and D. Ferrari, ''A Framework for the Study of Pricing in Integrated Networks'', Tech. Rept. TR-92-016, International Computer Science Institute, Berkeley, CA, March 1992.

C. Parris, H. Zhang, and D. Ferrari, ''A Mechanism for Dynamic Re-routing of Real-time Channels'', Tech. Rept. TR-92-053, International Computer Science Institute, Berkeley, CA, August 1992.

J. Ramaekers and G. Ventre, ''Quality of Service Negotiation in a Real-Time Communication Network'', Tech. Rept. TR-92-023, International Computer Science Institute, Berkeley, CA, April 1992.

J. Ramaekers and G. Ventre, ''Client-Network Interaction in a Real-Time Communication Environment'', *Proc. GLOBECOMM '92*, Orlando, FL, December 1992.

K. Umemura and A. Okazaki, ''Real-Time Transmission and Software Decompression of Digital Video in a Workstation'', Tech. Rept. TR-91-004, International Computer Science Institute, Berkeley, CA, January 1991. (not available from the file server)

D. Verma, ''Guaranteed Performance Communication in High Speed Networks'', Rept. No. UCB/CSD 91/663, University of California, Berkeley, CA, December 1991. (not available from the file server)

D. Verma and D. Ferrari, ''Evaluation of Overflow Probabilities in Resource Management'', Tech. Rept. TR-91-051, International Computer Science Institute, Berkeley, CA, October 1991; also in *Proc. ICC'92*, Chicago, IL, June 1992. (not available from the file server)

D. Verma and D. Ferrari, ''Variation of Traffic Parameters in ATM Networks'', *Proc. ICC'92,* Chicago, IL, June 1992. (not available from the file server)

D. Verma, H. Zhang, and D. Ferrari, ''Delay Jitter Control for Real-Time Communication in a Packet Switching Network'', Tech. Rept. TR-91-007, International Computer Science Institute, Berkeley, CA, January 1991; also in *Proc. TriComm '91*, Chapel Hill, NC, pp. 35-43, April 1991.

B. Wolfinger, ''Ein Transportdienst und -protokoll fuer Multimedia-Applikationen in paketvermittelten Hochgeschwindigkeitsnetzen mit Realzeit-Kommunikations-anforderungen'', *Proc. Workshop Entwicklungstenden-zen von Rechnernetzen*, Gaussig, Germany, November 1991. (not available from the file server)

B. Wolfinger and M. Moran, ''A Continuous Media Data Transport Service and Protocol for Real-Time Communication in High Speed Networks'', *Proc. Second International Workshop on Network and Operating System Support for Digital Audio and Video*, Springer-Verlag, Heidelberg, Germany, pp. 171-182, November 1991.

H. Zhang and D. Ferrari, ''Rate-Controlled Static Priority Queueing'', Tech. Rept. TR-92-003, International Computer Science Institute, Berkeley, CA, January 1992.

H. Zhang and T. Fisher, ''Preliminary Measurement of the Real-Time Internet Protocol'', *Proc. Third International Workshop on Network and Operating System Support for Digital Audio and Video,* San Diego, CA, November 1992.

H. Zhang and S. Keshav, ''Comparison of Rate-Based Service Disciplines'', *Proc. ACM SIGCOMM '91*, Zurich, Switzerland, pp. 113-122, September 1991.