The Tenet Group

University of California, Berkeley
and International Computer Science Institute

Recent and Current Research

April 1994

1. Introduction

The Tenet Group was formed in September of 1989 at the University of California at Berkeley and at 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 design and development of real-time communication services, and on network support for continuous-media applications.

In its real-time communication work, the group emphasizes mathematically provable (but not necessarily deterministic) performance guarantees, contractual relationships between client and service, general parameterized user-network interfaces with multiple traffic and QOS (quality of service) bounds definable over continuous ranges, and large heterogeneous packet-switching networking environments. The Tenet solution to the guaranteed-performance problem was one of the very first approaches based on admission control and rate control that were proposed and fully specified. To our knowledge, it is the first solution implemented and fully operational on a wide-area network.

The Tenet Group is participating in three major multi-institution research projects: BLANCA, Sequoia 2000, and BAGnet. 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 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. The initial version of Xunet 3 West, which connects the U.C. Berkeley campus to the Lawrence Berkeley Laboratory, is based on HIPPI instead of ATM, due to the earlier availability of HIPPI equipment. Most 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 allow them to visualize and browse sequences of digitally represented satellite maps and other large sets of data to be used in earth science research. 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 2000 network (S2Knet) connects at T3 speed (45 Mbps) the Berkeley campus to the Davis, 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 testbed for the third project, BAGnet (the Bay Area Gigabit network), is currently being installed. BAGnet will connect fifteen top computing research organizations located in the San Francisco Bay area; the organizations have obtained support for ATM SONET services from CalREN, a research support foundation established by Pacific Bell. These services are to be used for a number of collaborative multimedia applications, the first of which to be implemented will be a "teleseminars" facility. The current bandwidth is 155 Mbps, and may be upgraded to 622 Mbps later.

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

2. Real-Time Communication Services

2.1. Evaluation of Suite 1

The first Tenet real-time protocol suite, Suite 1, was presented and demonstrated during the 16th Annual Industrial Liaison Program Conference of the College of Engineering at the University of California at Berkeley on March 10, 1994.

Measurement of Suite 1 on the Sequoia 2000 testbed. Anindo Banerjea, Ed Knightly, Fred Templin and **Hui Zhang** have performed extensive experiments with the Tenet Real-Time Protocol Suite 1 on the Sequoia 2000 testbed. This study is unique in that it used real multimedia applications on a production wide-area network infrastructure to compare the performance provided by RMTP/RTIP to that provided by UDP/IP. The measurements showed that the end-to-end throughput and packet inter-arrival times remain unaffected for RMTP/RTIP connections even during periods of extremely high network congestion, while the performance of multimedia traffic carried by the UDP/IP protocol under the same network loading conditions degraded substantially. The Mean Opinion Score (MOS) of participants in an opinion poll was used to evaluate the perceived quality of video transported by the two protocol suites (UDP/IP and RMTP/RTIP). This investigation confirmed that the perceived quality of real-time video transported via UDP/IP degrades when the network load increases, while the quality of the RMTP/RTIP connection is largely unaffected. The measurements also showed that, when there are multiple RMTP/RTIP connections simultaneously active, not only are the RMTP/RTIP connections protected from congestion in the network due to IP traffic, but they are also protected from each other by admission control and priority scheduling in the hosts and routers. Finally, these experiments demonstrated that the latency involved in performing the admission control tests is negligible with respect to the link delays and signaling overheads.

Comparison of Tenet protocols with alternative real-time protocols. Pasquale Di Genova compared Tenet real-time protocols to other protocols already implemented (ST-II, SRP) or under development (RSVP). In his study, he compared solutions to the congestion control problem ranging from congestion avoidance to predicted service, to guaranteed levels of service. He also considered a transport layer protocol in the Tenet real-time suite, the Continuous Media Transport Protocol (CMTP), and simulated the simple algorithm on which CMTP is based, which significantly reduces buffer usage by smoothing out traffic.

2.2. Enhancements and extensions of the Tenet approach

Several improvements and extensions are being investigated for the original Tenet approach. The enhancements of Scheme 1 are easily portable to Scheme 2 (described in Section 2.4).

New admission tests for deterministic service. Hui Zhang has developed new admission control algorithms for providing deterministic service. In the previous version of Scheme 1, the deterministic test ensures that the sum of the peak rates of all deterministic connections is less than the link speed. This condition will result in a low average link utilization if the peak-rate to average-rate ratio is high for real-time traffic. The new study shows that deterministic service can be provided even when this condition is not met. When the traffic is bursty, the new admission control algorithm results in a multifold increase in the number of admitted real-time connections.

A new approach to statistical guarantees. Hui Zhang and **Ed Knightly** have developed a new, efficient, and general approach for providing end-to-end statistical guarantees in packet-switching networks. This is achieved by modeling a traffic source with a family of interval-dependent bounding random variables and by using a rate-controlled service discipline inside the network. The traffic model stochastically bounds the number of bits sent over time intervals of different lengths. The model captures different source behaviors over different time scales by making the distribution an explicit function of the interval length. The RCSP service discipline has the priority queuing mechanisms necessary to provide statistical performance guarantees in integrated services networks. In addition, RCSP provides the means for efficiently extending the results from a single switch to a network of arbitrary topology.

Modifying the Ethernet for real-time traffic. Brian Whetten is testing out a collision resolution algorithm (CRA) as an alternative to CSMA/CD for the new 100 Mbps Ethernet controllers being developed. With simulations, he has shown that CSMA/CD exhibits the Packet Starvation Effect (PSE) at offered loads above 60-70%. The PSE causes a significant percentage of packets to experience delays up to 50 times that of the average packet and some packets to starve completely. The PSE makes CSMA/CD unsuitable for real-time traffic at offered loads above 60%. His proposed alternative, called Fair Dual Distributed Queuing (FDDQ), provides utilization and average latency comparable with those of CSMA/CD, but puts packets into one of two global distributed queues when the network becomes congested. In this way, it provides two-priority FCFS access to the network, which is sufficient to obtain performance guarantees. Whetten is currently finishing the characterization of FDDQ through simulations, and hopes to build one or more controllers in the near future.

CPU scheduling for real-time networking. Masaya Miyazaki is developing a new CPU scheduling algorithm. Current real-time network protocols, including the Tenet real-time protocol Suite 1, do not fully take CPU resources into account. A few CPU scheduling algorithms take network tasks into account, but their models of real-time tasks are too simple to handle real-time network protocols . The goal is to provide a new mechanism to schedule both real-time network protocols and other real-time tasks such as data decompression tasks and data displaying tasks. Miyazaki is currently looking into the applicability of the new admission test algorithm recently proposed by Hui Zhang to the CPU schedulability problem. He is also looking into the Sporadic Server algorithm proposed by the CMU real-time research group to schedule real-time tasks on a fixed-priority preemptive kernel. He plans to implement and evaluate this algorithm in the DEC OSF/1 real-time kernel (which conforms to POSIX 1004.3).

Applying Tenet admission control to ATM networks. Rebbie Moon is studying the possibility of applying the Tenet approach to ATM networks. She is currently performing an analysis to compare the services supported under the ATM Forum User-Network Interface, Version 3.0, with those available in Tenet Scheme 2 (described in Section 2.4).

Real-time traffic over ATM networks. Pillalamarri Seshasayi is looking into the problem of providing real-time communication with guaranteed quality of service (QOS) in ATM networks. Guaranteeing QOS in an ATM network requires that the traffic on a particular channel be characterized inside the network. While such a characterization is well understood at the source, it is not clear how it changes after the channel has traversed many switches in a loaded network. He is investigating by simulation the characteristics an efficient real-time service should have in ATM networks.

2.3. Porting and extending the prototype Tenet real-time protocol suite (Suite 1)

Scheduling real-time transmissions over HIPPI switches. Riccardo Bettati is investigating means of supporting real-time traffic over circuit-switched networks. In order to give performance guarantees to connections while achieving a high utilization of network resources, he is studying a time division multiplexing approach. In such an approach, the admission test for a new connection consists of the explicit construction of a transmission schedule. The problem of optimally determining transmission schedules for even simple network topologies is non-tractable. Simulation results have shown, however, that simple EDD-based heuristics achieve very high success rates in generating such schedules. Work is under way (together with **Anisoara Nica**) to implement this approach as part of a real-time communication service for the Xunet 3 West HIPPI testbed connecting the University of California at Berkeley campus and the Lawrence Livermore Laboratory.

Integrating RCAP with the Xunet and Fore Systems signaling protocols. Ed Knightly is collaborating with the networking group at Sandia National Laboratories to investigate providing real-time performance guarantees on several interconnected network testbeds, including Xunet 2, a Fore Systems ATM LAN, and a heterogeneous network composed of FDDI rings connected with a DEC Gigaswitch. The project currently has several applications running on it, including the DAVE (Distributed Audio Video Environment) conferencing tools and an MPEG video distribution system. In addition, several PVM (parallel virtual machines) distributed applications with demanding network constraints are being tested. The next step to providing end-to-end real-time guarantees to these applications involves integrating RCAP (running on all workstations and routers) with the Xunet and Fore Systems signaling protocols. **Fukiko Hidano** is currently installing RCAP on top of the Xunet signaling protocol, and interfacing this protocol to RTIP, so that the performance guarantees now available in Xunet 2 between routers can become available between hosts.

2.4. The next Tenet scheme (Scheme 2) and suite (Suite 2)

The Scheme 2 paradigm. **Domenico Ferrari, Riccardo Bettati, Amit Gupta, Wendy Heffner,** and **Mark Moran** have completed the final design of the second-generation Tenet real-time protocol suite, Suite 2. The main focus of this work is on extending the basic real-time communication service provided by Suite 1 in two important respects: to provide support for efficient multi-party communication, and to make the client-service interface more flexible. In Scheme 2, the set of algorithms and abstractions on which Suite 2 is based, support for multi-party communication includes networking abstractions and optimizations to improve the utilization of network resources for multi-party communication. The key networking abstraction is *target sets*, which are similar to *host groups* in IP multicast. Receivers join target sets depending on their interest in a certain class of data (e.g., video from a given distributed conference). Sources then establish channels to target sets, causing real-time channels to be attempted to all destinations in the target set. Join and leave primitives support dynamic membership in target sets, as well as the associated change in multicast channels. Efficient utilization of network resources is achieved by the use of true multicast channels and *sharing groups*, which allow the network to use application-specific information to share network resources between channels while maintaining the client-specified performance guarantees. To improve the flexibility of the client-service interface, the values of input parameters may be specified as a pair, representing the client's desired value and the worst acceptable value. This provision reduces the need for negotiations. The initial design was subjected to a design review by internal and external judges in October 1993. Feedback from the design review has been incorporated into the design document, along with some implementation strategies. **Rainer Lienhart** and **Rebbie Moon** have also contributed to the final design.

Implementation of Scheme 2. **Raj Yavatkar** and **Wendy Heffner** are planning the implementation of Suite 2. A timetable for the development has been created, and modules are being assigned to the implementors: **Amit Gupta**, **Riccardo Bettati**, Raj Yavatkar, and Wendy Heffner. Implementation will begin this summer. Wendy Heffner and **Pasquale Di Genova** have begun to prepare a testbed on which to test the Suite 2 protocols. For this purpose, a network of 486-based PC systems running BSD/386 Unix from BSD*i* has been installed. Various pieces of the network have been procured and tested. Each of the hosts has four ports for network interfaces; pairs of interfaces are connected through Ethernets (if there is only one sender on the Ethernet, then the Ethernet link is a good approximation of a direct unicast link). The testbed provides a simple but real and working platform for the testing of routing algorithms, as well as of admission control algorithms for multicast real-time channels, and for the implementation of the resource sharing and the advance reservations mechanisms.

Routing multicast real-time channels. **Ron Widyono** and **Mark Moran** have developed and begun to evaluate techniques for routing real-time multicast channels. Recent contributions include a prototypical cost function to represent available network resources and several heuristics for computing constrained Steiner (min-cost) trees. Widyono has performed preliminary simulations to evaluate these techniques. The cost function reflects the availability of throughput and scheduling priority (which affects delay) on each link. Preliminary simulations have indicated that this cost function performs better than one incorporating a hop count, schedulability or throughput alone. More simulations will be performed in the near future to determine the correlation between the relative cost of trees produced by various algorithms (which is not affected by the suitability of the cost function) and their relative performance in the simulation experiments. A high correlation would indicate that the cost function is providing an accurate comparison of available network resources. The Delay-Constrained Bellman Ford algorithm (DCBF) finds a set of minimum cost paths from a single source to a set of destinations where each path meets the delay constraint for the given destination. Analysis and simulations confirm that this algorithm is more efficient than the Floyd-Warshall All Pairs Shortest Path algorithm we had been using. These constrained min-cost paths must be merged to produce a multicast tree that approximates a constrained Steiner tree. Two techniques have been developed to increase path overlap during the merge: clustering nodes before routing, and a greedy approach in which the next path to be merged is chosen so as to add the minimal incremental cost to the partial tree. Preliminary simulations indicate that both of these techniques improve the algorithm for certain scenarios. Many more simulations will be conducted in this area.

Scheme 2 simulator. **Amit Gupta**, **Mark Moran**, and **Ron Widyono** have continued their work on a simulator for the algorithms and techniques that comprise Scheme 2. The previous work included support for establishment and routing for multicast (1-to-N) channels in a network that supports resource partitioning (described in a later paragraph in this section). The new work includes support for five more routing algorithms in the simulator, and support for resource sharing and advance reservations (also described later). We also have added mechanisms for comparing different schemes by randomly generating different call arrival and holding times, and using the same arrival and holding time patterns for simulations with alternative algorithms and mechanisms. We are running these simulations on a topology identical to that of the current NSFNet backbone. We have also started working on a completely redesigned simulator for running the simulations.

Resource sharing for multi-party applications Resource sharing mechanisms allow the network to use application specific information to share network resource allocations between channels while maintaining client-specific performance guarantees. **Amit Gupta**, **Mark Moran** and **Domenico Ferrari** have continued their work on resource sharing (with help from two other students, Winnie Howe and Quyen Nguyen). The previous work included mechanisms for doing resource sharing between different channels and the channel group abstraction, which provides the applications with a unifying mechanism for specifying a variety of inter-channel relationships. The new work includes double bucket, a new traffic specification model that is well-suited for specifying resources for groups of channels. Our analysis shows that the resource sharing gains depend on the network topology, and the routing policies followed. In particular, the gains improve dramatically when the routing mechanisms try to route sharing channels through common sub-paths. We have also run a number of simulations that verify the above conclusions about the gains due to resource sharing. The simulation results show that resource sharing results in very significant increase in channel acceptance rates (gain factors of 2 to 4 were observed with a simulated workload on the NSFNet backbone topology).

Resource partitioning in real-time networks. Domenico Ferrari and **Amit Gupta** have continued their investigation of the problem of partitioning the resources of an integrated-services network. An efficient solution to this problem has many potential applications: for example, the construction of virtual private networks, the fast establishment of real-time channels, the re-routing of channels in mobile communication, and the advance reservation of real-time channels. The solution proposed by Ferrari and Gupta works within the context of the Tenet schemes: per-partition tests and computations have been shown to be derivable from those that apply to the entire network in each node; therefore, the subnetworks defined by partitions of each node's resources can be treated during channel establishment and teardown as though they were independent and isolated networks. The amount of a resource assigned to a given partition in a node may differ from the amounts assigned to the same partition in all other nodes. One of the partitions may be dedicated to non-real-time traffic; its resources will not usually be available to real-time traffic. Per-partition tests have been derived for many scheduling disciplines, including FCFS, RCSP, and variants of EDD. Simulations have shown that resource partitioning reduces the computational overhead associated with running admission control tests. Also, under certain circumstances, resource partitioning can lead to higher overall acceptance rate for multi-party real-time channels.

Advance reservation of real-time channels. In Suite 1, the network provides the clients with quality-of-service guarantees during data transfer. However, for such a service to be truly usable in multi-party applications, the clients should be able to reserve performance-guaranteed channels in advance. **Domenico Ferrari**, **Amit Gupta**, and **Giorgio Ventre** (a former Tenet Group member from the University of Naples, Italy) have designed mechanisms for supporting advance reservations in real-time communication networks. These mechanisms are based on the resource partitioning work described above. They have run simulations on many different networks, with different call traffic, call arrival and call holding time patterns. The preliminary results indicate that advance reservations interact very well with the resource sharing mechanisms, and that, during periods of heavy real-time network load, significantly higher channel acceptance rates are obtained with advance reservation facilities than without them. Under these conditions, advance reservations reduce the per-channel computational overhead (for resource reservation and admission control); these gains are primarily due to the manner in which the underlying resource partitioning mechanisms interact with resource sharing.

Dynamic Connection Management. Colin Parris is currently implementing the DCM scheme that provides a network with the ability to dynamically modify the traffic and performance parameters as well as the routes of real-time connections. These modifications are subject to a *modification contract* which specifies the degree of disruption that a client can accept during these modifications. Dynamic Connection Management consists of the DCM scheme and the DCM policies. The DCM scheme is a collection of algorithms that permit the modification of the parameters and routes, whereas the DCM policy is a collection of rules that determine whether modifications should be attempted, and provide the appropriate traffic and performance parameters and the routing constraints to the DCM scheme. The DCM scheme is being implemented using the Simple Network Management Protocol (SNMP) v1 as the management protocol, with RMTP and RTIP as the data transfer protocols. In this implementation, connection establishment and modification can be initiated by either the client or the network. The client can initiate actions through a sequence of Unix socket calls, which inform the DCM policy manager of the client's requests. The policy manager then directs the SNMP agents. In the network-initiated case, the DCM policy manager uses the network state to determine if policy actions are to be taken and then directs the SNMP agents. The DCM policy manager establishes or modifies a connection by doing an SNMP `set` on the traffic and performance objects related to the appropriate link. SNMP `get` commands can also be used to read the connection state information in the nodes. SNMP will also be used to monitor RMTP/RTIP, thus providing DCM policies with full monitoring capabilities for real-time traffic, as well as full control capabilities for real-time connections through the DCM schemes. Experiments will be conducted to compare the performance of this SNMP-based implementation with that of RCAP.

A test application for DCM. **Rainer Lienhart** is working on the design and implementation of a network-based Lossless Digital Image Browser (LDIB) which will extensively use the DCM algorithms. LDIB is a new class of applications proposed by Colin Parris and Hui Zhang. It allows the user to retrieve image sequences from remote sites across the network, viewing them at different speeds and in either direction. It differs from other networked video applications because it is lossless. Instead of dropping frames or changing image resolution to accommodate the changes in the frame rate, a LDIB uses the DCM algorithms to renegotiate the parameters of its connection.

Fault recovery mechanisms. **Anindo Banerjea** has developed rerouting schemes to allow real-time networks to recover from link faults. These schemes work over the Dynamic Channel Management protocol developed by Colin Parris. Recent experiments using a simulator developed in cooperation with Colin Parris have shown that these schemes can allow very fast recovery of a significant portion of the traffic affected by single or multiple link faults, with reasonably efficient use of network resources. Experiments are being performed to explore the influence of various factors such as varying topology, application mix and network load. Preliminary results show that the schemes are stable under a wide variety of conditions, and degrade gracefully under increasing load. Experiments are also being conducted on the use of dispersity routing to provide fault tolerance for real-time channels. These experiments will be used to design schemes for real-time channels which deliver uninterrupted service in the presence of single network faults.

3. Continuous-Media Networking Applications

Measurements and observations of IP multicast traffic. A recent development on the Internet has been the advent of Internet-wide IP multicasting services, as supported by the MBONE (multicast backbone). The most frequent applications seen on the MBONE are the so-called "multimedia" applications, such as the `vic`, `nv`, `vat`, and `wb` conferencing tools. These programs frequently generate network data whose characteristics are very different from traditional bulk data and interactive applications seen on the Internet, such as ftp and telnet. Understanding the network usage of these applications is important in order to control and design networks. This need is particularly compelling in light of the increasing popularity of multimedia applications. To better understand the network usage of these applications, **Bruce Mah** has been collecting and analyzing traces of IP multicast traffic in and around U.C. Berkeley.

Multiplexing IP datagrams on ATM virtual circuits. **Bruce Mah** is investigating the transmission of IP datagrams over ATM virtual circuits. When an IP datagram is sent across an ATM subnet, it must be mapped onto a virtual circuit for transmission. Current IP over ATM implementations typically establish a single virtual circuit for all traffic between any pair of hosts. However, this strategy does not support per-conversation quality of service. By examining the network- and transport-layer headers of an incoming datagram, an IP gateway at the edge of an ATM cloud can switch that datagram to use the appropriate ATM virtual circuit. Various characteristics such as well-known TCP ports or the IP TOS field can be used to help determine an appropriate quality of service for new virtual circuits. In effect, this amounts to ‘‘synthesizing’’ a connection for datagrams while they traverse the ATM subnet, in order to associate them with appropriate virtual circuits and qualities of service. A mechanism to implement this multiplexing will be built into the routers for Xunet 2, and various multiplexing policies will be tested and evaluated.

MBONE coordination tool. Hoofar Razavi is working on a coordination tool (`ct`) for the Internet multicast backbone. The tool, which is still under development, will provide efficient means for audio and video synchronization, statistics collection for a conference, as well as floor control functionality for large seminars.

MBONE VCR tool. Wieland Holfelder is working on the design of a VCR-like application for use on the MBONE. This application uses a VCR interface to control some of the sessions that are sent over the MBONE. It will allow users to record and play back these sessions. It will also allow users to browse forward (FF) or backward (REW) through a recorded session. Holfelder is also investigating efficient storage mechanisms for the multimedia data, synchronization aspects of this data, and session management and indexing. One target use of this application is to use it in conjunction with the teleseminars application to be deployed on BAGnet.

Priority Encoding Transmission. Andres Albanese, in collaboration with Michael Luby, Johannes Blomer, and Jeff Edmonds from the ICSI Theory group, and Madhu Sudan from the IBM Research Center in Yorktown Heights, have introduced a novel method for sending messages simultaneously over high and low-speed networks. The new method, called Priority Encoding Transmission, allows a user to specify a different priority on each part of the message to be transmitted over a lossy heterogeneous medium. Based on the given priorities, the sender uses the system to encode the message into packets for transmission. The system ensures a smooth rate of recovery of the message by receivers as a function of the received fraction of the information transmitted. This work has immediate applications to multimedia multicast communication and high speed networks with heterogeneous resources.

Extending TCP/IP to support real-time transmissions in mobile computing. Pietro Manzoni is working on issues related to mobile computing. He has proposed changes to the IP protocol to support mobility in the Internet. He is now investigating how the TCP/IP protocol suite should be extended to support real-time and continuous-media traffic in a mobile environment. The goal is to understand how simple modifications can improve IP’s behavior in the face of congestion. In addition to a new approach for admission control of incoming packets, support for mobile computing requires that guarantees given at connection establishment be maintained even during a route change. Manzoni is currently studying the applicability of an advanced reservation scheme to solve this problem.

Application characterization and related performance guarantees. Ed Knightly is collaborating with **Hui Zhang** and Robert Mines of Sandia National Laboratories to analyze the traffic characteristics and performance requirements of various real-time applications along with their respective multiplexing characteristics. For example, this work seeks to quantify the gains in network utilizations in the recently proposed extensions to the Tenet deterministic and statistical models. This investigation seeks a better understanding of the multiplexing characteristics of actual applications running on the testbeds, including applications such as the DAVE (Distributed Audio Video Environment) conferencing tools, an MPEG video distribution system, several PVM (parallel virtual machine) distributed applications, and scientific visualization applications.

Multimedia conferencing tools development and deployment. Steven McCanne has been investigating schemes for efficient and robust transport of video over packet-switching networks. He has built a video conferencing prototype called *vic*, which employs an extensible architecture to facilitate experimentation with video coding schemes and session protocols. *Vic* is in production use on the Sequoia 2000 and Xunet 2 testbeds.

Characteristics of wide-area TCP connections. Vern Paxson has been studying the characteristics of wide-area network traffic on today’s Internet. He has investigated growth-trend patterns at the Lawrence Berkeley Laboratory (LBL), and developed analytic models for characterizing the bulk transfer properties of wide-area TCP connections. Together with Sally Floyd of LBL, he performed a study assessing the validity of assuming that wide-area traffic has Poisson arrival processes; they found that in general the assumption grossly underestimates burstiness, and that other models reflect the traffic much more faithfully. He is currently developing tools for dynamically measuring network characteristics.

4. Tenet Group Members

The Tenet Group consists of students and faculty of the Department of Computer Science at U.C. Berkeley as well as visiting scholars and post-doctoral researchers. The group is led by **Domenico Ferrari** and includes the following graduate students: **Anindo Banerjea**, **Amit Gupta**, **Wendy Heffner**, **Fukiko Hidano**, **Edward Knightly**, **Steven McCanne**, **Bruce Mah**, **Rebbie Moon**, **Mark Moran**, **Colin Parris**, **Vern Paxson**, **Brian Whetten**, and **Ron Widyono**. In addition, the group includes an undergraduate student, **Hoofar Razavi**, and a U.C. Berkeley Computer Science Reentry Program student, **Anisoara Nica**.

The following visiting scholars and post-doctoral researchers are also members of the Tenet Group: **Riccardo Bettati**, a Swiss post-doctoral fellow at ICSI; **Stephan Boecking**, from the Corporate R&D Division, Siemens, Munich, Germany; **Pasquale Di Genova**, a Junior Specialist at U.C. Berkeley; **Wieland Holfelder**, from the University of Mannheim, Germany; **June Kato**, from the NTT Software Laboratories, Tokyo, Japan; **Rainer Lienhart**, from the University of Mannheim, Germany; **Pietro Manzoni**, from the Politecnico di Milano, Milano, Italy; **Masaya Miyazaki**, from the Core System Development Center (CSDC), Matsushita Electric Industrial Co., Ltd., Osaka, Japan; **Pillalamarri Seshasayi**, from Siemens, Munich, Germany; **Fred Templin**, from the External Research Program of Digital Equipment Corporation, Maynard, MA, U.S.A.; **Hirohisa Wakai**, from the Information Technology Research Laboratories, Sharp Corporation, Nara, Japan; **Raj Yavatkar**, from the University of Kentucky; **Hui Zhang**, from the Lawrence Berkeley Laboratory; and the co-leader of ICSI's Networks Group, **Andres Albanese**.

5. Tenet Group Publications

The Tenet Group publications since July 1992 are listed below. Unless otherwise noted, the publications are available by anonymous FTP to tenet.berkeley.edu and through the World Wide Web at URL http://tenet.berkeley.edu. Publications are also 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*. Inquiries for further information should be directed to the Tenet Group's leader, Professor Domenico Ferrari, Computer Science Division, EECS Department, University of California, Berkeley CA 94720. Electronic mail may be sent to ferrari@cs.berkeley.edu.

A. Banerjea, E. Knightly, F. Templin, and H. Zhang ''Experiments with the Tenet Real-Time Protocol Suite on the Sequoia 2000 Wide Area Network'', submitted to ACM Multimedia 94; also Tech. Rept. TR-94-020, International Computer Science Institute, Berkeley, CA, April 1994.

A. Banerjea and S. Keshav, ''Queuing Delays in Rate Controlled Networks'', Technical Report TR-92-015, International Computer Science Institute, Berkeley, CA, February 1992; also in *Proc. INFOCOM'93*, San Francisco, CA, March-April 1993.

A. Banerjea, C. Parris, and D. Ferrari, ''Recovering Guaranteed Performance Service Connections from Single and Multiple Faults'', Technical Report TR-93-066, International Computer Science Institute, Berkeley, CA, November 1993.

R. Cáceres, ''Multiplexing Traffic at the Entrance to Wide-Area Networks'', Ph.D. Dissertation, Technical Report UCB/CSD 92/717, University of California, Berkeley, CA, December 1992.

P. Di Genova, ''Real-Time Communication on Computer Networks: Introduction and Analysis of the Tenet Approach of Berkeley'', Tesi di Laurea, University of Bologna, Italy, December 1993.

W. Effelsberg and E. Müller-Menrad, ''Dynamic Join and Leave for Real-Time Multicast'', Technical Report TR-93-056, International Computer Science Institute, Berkeley, CA, October 1993.

D. Ferrari, ''A New Admission Control Method for Real-Time Communication in an Internetwork'' in S. Son, Ed., *Real-Time Systems.* Prentice-Hall, Englewood Cliffs, NJ, 1994.

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’’, Technical Report TR-91-056, International Computer Science Institute, Berkeley, CA, October 1991; also in *Journal of Internetworking: Research and Experience*, vol. 4, n. 1, pp. 1-20, 1993.

D. Ferrari, ‘‘Real-Time Communication in an Internetwork’’, *Journal of High Speed Networks*, vol. 1, n. 1, pp.79-103, 1992; also Technical Report TR-92-001, International Computer Science Institute, Berkeley, CA, January 1992.

D. Ferrari, A. Banerjea, and H. Zhang, ‘‘Network Support for Multimedia - A Discussion of the Tenet Approach’’, Technical Report TR-92-072, International Computer Science Institute, Berkeley, CA, October 1992; to appear in *Computer Networks and ISDN Systems, special issue on Multimedia Networking, 1994.*

D. Ferrari and A. Gupta, ‘‘Resource Partitioning for Real-Time Communication’’, *Proc. IEEE Symp. on Global Data Networking,* Cairo, Egypt, December 1993.

D. Ferrari, A. Gupta, M. Moran, and B. Wolfinger, ‘‘A Continuous Media Communication Service and its Implementation’’, *Proc. GLOBECOM ’92*, 8A.02., Orlando, FL, 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*, pp. E1-1–E1-14, 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.

A. Gupta, W. Heffner, M. Moran, C. Szyperski, ‘‘Network Support for Realtime Multi-Party Applications’’, Fourth International Workshop on Network and Operating System Support for Digital Audio and Video, Lancaster, England, November, 1993.

A. Gupta and M. Moran, ‘‘Channel Groups: A Unifying Abstraction for Specifying Inter-Stream Relationships’’, Technical Report TR-93-015, International Computer Science Institute, Berkeley, CA, March 1993.

K. Keeton, B. Mah, S. Seshan, R. Katz, and D. Ferrari, ‘‘Providing Connection-Oriented Network Services to Mobile Hosts’’, *Proc. of the 1993 USENIX Symposium on Mobile and Location-Independent Computing*, Cambridge, MA, August 1993.

E. Knightly and G. Ventre, ‘‘Galileo: a Tool for Simulation and Analysis of Real-Time Networks’’, Technical Report TR-93-008, International Computer Science Institute, Berkeley, CA, March 1993; also in *Proc. IEEE 1993 International Conference on Network Protocols,* San Francisco, CA, October 1993.

B. Mah, ‘‘A Mechanism for the Administration of Real-Time Channels’’, M.S. Report, Technical Report UCB/CSD 93/735, University of California, Berkeley, CA, March 1993.

B. Mah, S. Seshan, K. Keeton, R. Katz, and D. Ferrari, ‘‘Providing Network Video Service to Mobile Clients’’, *Proc. Fourth Workshop on Workstation Operating Systems*, Napa, CA, October 1993.

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.

E. Müller-Menrad, ‘‘Dynamic Membership in Real-time Multicast Groups’’, M.S. report, Technische Universitat Muenchen, Munich, Germany, 1993.

C. Parris and A. Banerjea, ‘‘An Investigation into Fault Recovery in Guaranteed Performance Service Connections’’, Technical Report TR-93-054, International Computer Science Institute, Berkeley, CA, October 1993.

C. Parris, G. Ventre, and H. Zhang, ‘‘Graceful Adaptation of Guaranteed Performance Service Connections’’, *Proc. IEEE GLOBECOM‘93,* Houston, TX, November 1993.

C. Parris, H. Zhang, and D. Ferrari ‘‘A Dynamic Management Scheme for Real-Time Connections’’, *Proc. INFOCOM ‘94*, Toronto, Canada, June 1994.

C. Parris, H. Zhang, and D. Ferrari, ‘‘A Mechanism for Dynamic Re-Routing of Real-Time Channels’’, Technical Report TR-92-053, International Computer Science Institute, Berkeley, CA, August 1992.

C. Parris, H. Zhang, and D. Ferrari, ‘‘Dynamic Management of Guaranteed Performance Multimedia Connections’’, *ACM/Springer-Verlag Multimedia Systems*, 1994.

V. Paxson, ‘‘Empirically-Derived Analytic Models of Wide-Area TCP Connections: Extended Report’’, submitted to *IEEE/ACM Transactions on Networking*, 1994.

V. Paxson, ‘‘Growth Trends in Wide-Area TCP Connections’’, to appear in *IEEE Network*, 1994.

V. Paxson and S. Floyd, ‘‘Wide-Area Traffic: The Failure of Poisson Modeling’’, to appear in *Proc. ACM SIGCOMM ’94 Conf.*, August-September 1994.

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

C. Szyperski and G. Ventre, ‘‘A Characterization of Multi-Party Interactive Multimedia Applications’’, *High-Performance Network Research Report 1/93*, and Technical Report TR-93-006, International Computer Science Institute, Berkeley, CA, January 1993.

C. Szyperski and G. Ventre, ‘‘Efficient Multicasting for Interactive Multimedia Applications’’, Technical Report TR-93-017, International Computer Science Institute, Berkeley, CA, March 1993.

C. Szyperski and G. Ventre, ‘‘Efficient Group Communication with Guaranteed Quality of Service’’, *Proc. Fourth IEEE Workshop on Future Trends in Distributed Computing Systems*, Lisboa, Portugal, September 1993.

R. Widyono, "The Design and Evaluation of Routing Algorithms for Real-Time Channels", M.S. Report, University of California, Berkeley, CA, May 1994.

H. Zhang and E. Knightly, ‘‘Providing End-to-End Statistical Guarantees Using Bounding Interval Dependent Stochastic Models’’, *Proc. ACM SIGMETRICS’94 Conf.*, Nashville, TN, May 1994.

H. Zhang and D. Ferrari, ‘‘Improving Utilization for Deterministic Service In Multimedia Communication’’, *Proc. 1994 IEEE International Conference on Multimedia Computing and Systems*, Boston, MA, May 1994.

H. Zhang, ‘‘Service Disciplines For Integrated Services Packet-Switching Networks’’, Ph.D. dissertation, University of California, Berkeley, CA, November 1993.

H. Zhang and D. Ferrari, ‘‘Rate-Controlled Static Priority Queueing’’, Technical Report TR-92-003, International Computer Science Institute, Berkeley, CA, January 1992; also in *Proc. INFOCOM’93,* San Francisco, CA, March-April 1993.

H. Zhang and D. Ferrari, ‘‘Providing Deterministic Guarantees For Bursty Traffic’’, *Proc. Workshop On the Role of Real-Time in Multimedia/Interactive Computing Systems, IEEE Real-Time Systems Symposium’93*, Raleigh-Durham, NC, November 1993.

H. Zhang and T. Fisher, ‘‘Preliminary Measurements 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.