The Tenet Group

University of California, Berkeley
and International Computer Science Institute

Recent and Current Research

May 1995

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 has participated in three major multi-institution research projects, two of which have recently completed their mission: BLANCA, Sequoia 2000, and BAGNet. The BLANCA project was part of the Gigabit Testbed Initiative, which was 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 included AT&T Bell Laboratories, the universities of California, Illinois, and Wisconsin, Rutgers and Columbia University, and a few research institutes, among which is ICSI. The goal of BLANCA was 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. Xunet 3 West, which connected the U.C. Berkeley campus to the Lawrence Berkeley Laboratory, was based on HIPPI instead of ATM, due to the earlier availability of HIPPI equipment. Most applications were of the scientific visualization type and required the transmission of image sequences and/or video streams. The sequences or streams were 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. Work is now continuing on Xunet 2, and the port of the Tenet Suite 1 to it has almost been completed.

Sequoia 2000, the second major initiative in which the Tenet Group has participated, aimed at building high-speed storage and communication facilities for global-change scientists at remote locations, allowing these scientists 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 took place at U.C. Berkeley and U.C. San Diego. The Sequoia 2000 network (S2Knet), which is still in existence, 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), has only recently been installed. BAGNet connects fifteen top computing research organizations located in the San Francisco Bay Area; the organizations have obtained support for ATM SONET services from CalREN (California Research and Education Network), a 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 infrastructure at both U.C. Berkeley and ICSI has a local SynOptics ATM switch with 16 155-Mbps fiber ports, providing ATM all the way to the workstations.

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

2. Real-Time Communication Services

2.1. Tenet Protocol Suite 1 source code release

The Tenet Group at UC Berkeley and ICSI announces the availability of the source code of its Real-Time Protocol Suite 1. The suite consists of three protocols (RMTP and RTIP for data delivery, RCAP for guaranteed-performance channel establishment and teardown), which have been designed to coexist with the Internet protocols. The source code can be freely used for educational and research purposes without a license; its commercial exploitation requires obtaining a license from the Regents of the University of California, who own the copyright to it. While RCAP runs in user mode, RMTP and RTIP are part of the kernel; the code being distributed can be used on Ultrix 4.2A, Irix 4.0.5f, and BSD/OS 2.0. An OSF-1 version is in preparation. The code is still being tested, and is therefore distributed as is, without any promises it will run in your environment. It can be obtained at URL ftp://tenet.cs.berkeley.edu/pub/src/tenet-suite1-0.8.tar.Z Licensing and installation information is included in the READ_ME file that accompanies the code of the protocols. For any questions or comments, please send e-mail to: suite1-comments@tenet.cs.berkeley.edu

2.2. Evaluations, 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 2.3).

A new approach to statistical guarantees. Hui Zhang (a former Tenet group member, now at Carnegie-Mellon University) 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.

A new traffic model (D-BIND) for deterministic service. **Ed Knightly** and Hui Zhang have developed a new deterministic traffic model called the Deterministic Bounding Interval Dependent (D-BIND) traffic model. This model has the advantage that it captures a source's burstiness over different time-interval lengths, thereby providing a more accurate traffic characterization than a model based only on a source's peak and average rate. The new traffic model, together with tight analysis techniques for deriving delay bounds, provides the potential for achieving high network utilization even for bursty traffic sources requiring deterministic guarantees on throughput and delay. As an example of the effectiveness of the new model, several long traces of MPEG-compressed video have been used to show that considerably high network utilizations are achievable even for such bursty video sources.

A new approach (RED VBR) to supporting VBR video in packet switched networks. Previous approaches to supporting VBR video on packet-switched networks include deterministic service, statistical service, predicted service, and feedback-based schemes. These schemes represent different tradeoffs in QOS, achievable network utilization, and method of dealing with overload. Hui Zhang and **Ed Knightly**

have proposed a new service that attempts to strike an efficient balance with the above tradeoffs. The approach, termed Renegotiated Deterministic VBR (RED VBR), is based on deterministic guarantees with client controlled renegotiation of QOS parameters and graceful adaptation during overload periods. Evaluation of the scheme with two traces of MPEG-compressed video showed that, even with simple renegotiation polices and relatively low renegotiation frequencies, high network utilization in the range of 50% to 80% can be achieved. For traffic that is bursty over long intervals, this represents a 100% to 150% improvement in network utilization compared to deterministic service. Compared to statistical and predicted service, this approach allows more graceful and client-controlled QOS degradation during overload period.

Fundamental limits and tradeoffs of providing deterministic guarantees to VBR video traffic. A network service that supports delay-sensitive video imposes many problems since compressed video sources are variable bit rate (VBR) with a high degree of burstiness. **Ed Knightly** and Hui Zhang, together with Jorg Liebeherr (former Tenet group member) and Dallas Wrege from the University of Virginia, have investigated a network service that can provide deterministic guarantees on the minimum throughput and the maximum delay of VBR video traffic. A common belief is that due to the burstiness of VBR traffic, such a service will not be efficient and will necessarily result in low network utilization. This conventional wisdom is shown to be untrue by first investigating the fundamental limits to providing deterministic guarantees to VBR video by considering an ideal network environment that employs the most accurate deterministic, time-invariant video traffic characterizations, Earliest-Deadline-First packet schedulers, and exact admission control conditions. It then considers fundamental tradeoffs by investigating the utilization achievable in a network environment that takes into account practical constraints, such as the need for fast policing mechanisms, simple packet scheduling algorithms, and efficient admission control tests.

Effects of smoothing and traffic shaping on network utilization. Ed Knightly and **Paola Rossaro**, are currently investigating how smoothing VBR sources such as MPEG-compressed video affects achievable network utilization. On one hand, smoothing individual sources makes their traffic less bursty, or more like constant bit rate traffic, so that higher network utilization can be achieved. On the other hand, smoothing also introduces a smoothing delay at the source. Thus, while smoothing decreases the bound on queueing delay, it increases the bound on smoothing delay, or the time that packets spend in the traffic shaper before entering the network. In this sense, smoothing can be considered a tradeoff between buffering at the source and buffering inside the network. Using the D-BIND traffic model and considering deterministic performance guarantees, preliminary analysis indicates that, for a single hop, smoothing or traffic shaping at the source is an ineffective means for achieving higher network utilization. That is, the savings in queueing delay achieved by transmitting smoother traffic are most likely outweighed by the smoothing delay incurred at the source. Preliminary investigations into the effects of smoothing in the case of multi-hop networks indicate that smoothing may be an effective means of reducing total end-to-end delay.

Independently and in parallel, **Amit Gupta** is investigating the same problem using sigma-rho models for traffic.

Modifying the Ethernet for real-time traffic. Brian Whetten is testing 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 plans to build one or more interface boards in the near future.

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. This scheme has been realized in RMP/RMCP, a prototype real-time communication service for XUNET-3 West, the HIPPI testbed connecting the U.C. Berkeley campus and Lawrence Berkeley Laboratory. In order to accomodate the considerable number of platforms (SunOS, Solaris, SGI/Irix, Raid-II) attached to the testbed, RMP/RMCP is implemented in user space. Nevertheless, experiments have shown that rates of over 100Mb per second can be sustained between a pair of Sparc-10s using this implementation. Additional experiments were run to measure the effect of contention for network resources. The results show that timing guarantees are below a few milliseconds.

Integrating RCAP with the Xunet signaling protocol. Fukiko Hidano has installed RCAP (the channel establishment protocol in Suite 1) on top of the Xunet signaling protocol, and is now interfacing this version of the protocol to the one running on the routers and the hosts, 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)

Riccardo Bettati, Domenico Ferrari, Amit Gupta, Wendy Heffner, Rebbie Moon, Mark Moran, Rolf Oppliger, Kurt Rothermel, Paola Rossaro, and **Raj Yavatkar** are continuing to work on Scheme 2, the next Tenet scheme for real-time communication, and are implementing the scheme in a protocol suite, Suite 2. The main focus of this work is to extend the basic real-time communication service provided by Suite 1 in two important respects: to provide abstractions and techniques for efficient multi-party communication and to make the client-service interface more flexible. The key networking abstraction is the *target set*, which is similar to the *host group* in IP multicast. Receivers join target sets depending on their interest in a certain transmission (e.g., video from a given distributed conference), and channels are established from data sources to the members of the respective target sets. 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 *resource sharing*. To improve the flexibility of the client-service interface, the values of the 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.

Implementation of Scheme 2. Raj Yavatkar and **Wendy Heffner** are coordinating Suite 2 implementation. **Riccardo Bettati, Amit Gupta,** Wendy Heffner, **Winnie Howe, Mark Moran, Quyen Nguyen,** and Raj Yavatkar have continued their work in implementing the Suite 2. This implementation provides a framework for experimenting with multi-party real-time network applications in heterogeneous internetworks: it supports multiple traffic specification models and packet scheduling disciplines; it also supports location independence, which allows third-party clients to initiate all actions. All of these facilities are administered by the RCAP daemon. In Suite 2, the RCAP daemon consists of a set of objects (e.g., Target Set, Channel, Establishment, and the Location and Naming Server) and two levels of message dispatchers. In the current prototype, which has been successfully tested, channels can be established to multiple destinations, and destinations can dynamically join and leave a target set and its associated channels. This prototype also supports advance reservation of real-time connections, as described in a later section. A prototypical routing module has been implemented that uses a heuristic to approximate min-cost, delay-constrained multicast trees. These routes are likely to be successfully established with delay bounds that meet the specified constraints. Additionally, a client library has been constructed that contains primitives to establish a connection to the daemon and to send and receive messages between clients and RCAP. Riccardo Bettati designed an agent-based architecture for monitoring and managing the network state. This work is being implemented as part of the existing graphical network management tools. We are currently implementing the code for sharing resource allocations among related channels (described in the resource sharing section below) and RTIP is being modified to transfer data over multicast channels.

Scaling in Suite 2. **Wendy Heffner** has evaluated the scalability of both the system design and the prototype implementation of Suite 2 of the Tenet Protols. The scalability of the design has been analyzed on several levels: with regard to its support for large internetworks, many multi-party connections, and large numbers of receivers in a single connection. In addition, the prototype implementation has been examined to see where decisions have been made that reduce the scalability of the initial system design. Where problems have been found, she has proposed implementation alternatives that are more scalable. Finally, she compared the scalability of system design to those of ST-II and of RSVP.

Testing of Suite 2. **Wendy Heffner** is continuing her work with a network of six 486-based PC systems running BSD/OS Unix from BSD*i*. Each host contains three Ethernet interface cards. Pairs of interfaces are connected to create a mesh network of degree three. This testbed has been set up to supply a variety of topologies and to allow for rapid reconfiguration; it will be used to debug and test the Suite 2 implementation. The testbed will also facilitate the evaluation of protocol performance at each phase of the implementation. The current prototype of RCAP2, the Suite 2 signaling protocol, has been ported to this testbed. This summer, she will work on developing RTIP2, the Suite 2 data delivery network layer, on this platform.

Routing multicast real-time channels. **Mark Moran** is developing scalable techniques for routing real-time multi-party channels. These techniques take into account dynamic information on the availability of network resources and produce multicast routes in which the path from the source to any destination meets the delay constraint of the destination. Resource availability is incorporated as a link cost. A Delay-Constrained Bellman Ford (DCBF) algorithm then finds a set of minimum-cost paths from a source to each destination. Analysis and simulations confirm that this algorithm is more efficient than the Floyd-Warshall All Pairs Shortest Path algorithm. These constrained min-cost paths are then merged to produce a multicast tree that approximates a constrained Steiner tree. Moran is currently conducting simulations to evaluate various cost functions and merge techniques. Since the distribution of routing update messages may limit the scalability of a routing architecture, techniques to limit update distribution are also being developed. Techniques proposed to limit the distribution of routing updates are based on the introduction of routing *hierarchies*. However, important information is lost when abstracting routing state from a lower level to a higher level. Instead of imposing a hierarchy, Moran is developing techniques to limit the *scope* of routing updates so that routing servers have general information about remote nodes and paths, and more detailed information about local nodes and paths. He is also building a simulator to evaluate these techniques. In the near future, the DCBF routing algorithm will be extended to enhance resource sharing performance by routing related channels over shared sub-paths.

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 still maintaining client-specific performance guarantees. **Amit Gupta** and **Mark Moran**, in association with ex-members Winnie Howe and Quyen Nguyen, have continued their work on resource sharing. The previous work included mechanisms for sharing resources among different channels as well as the *channel group* abstraction, which provides the applications with a single mechanism for specifying a variety of inter-channel relationships. Our analyses show that the gain from resource sharing depends 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. A number of simulations have been run to verify the above conclusions regarding the gains due to resource sharing. The simulation results show that resource sharing results in a 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. The generality of the resource partitioning techniques has been demonstrated by deriving the per-partition tests for many scheduling disciplines, including FCFS, RCSP, WFQ, 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 results in higher overall acceptance rates of multi-party real-time channels.

Advance reservation of real-time channels. When running the Tenet protocols, the network provides real-time 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 continued their research on providing advance reservation service for networks running the Tenet protocols. The mechanisms used to support this service are based on the resource partitioning work described above. Simulations have been run on many different networks, with different call traffic, call arrival, and call holding-time patterns. 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.

Rerouting real-time connections for mobile computing. Real-time communication is expected to play an important role in supporting distributed multimedia applications. Recent advances in small, portable computers and wireless networks are pushing distributed computing towards mobile, ubiquitous computing. A key problem in mobile computing concerns rerouting connections when a mobile computer moves into a different transmitter's range. **Amit Gupta** and **Raj Yavatkar** have designed techniques for supporting fast rerouting of real-time connections for mobile computing; these techniques are based on the resource partitioning mechanisms described above. They are currently designing simulation experiments to evaluate the performance of these techniques.

Fault recovery and fault tolerance for multi-party real-time communication. Amit Gupta and **Kurt Rothermel** are working on improving tolerance to faults in the Tenet multi-party protocols (Scheme 2). The objective is to re-establish connectivity for already established connections and to permit setting up new connections to existing conferences in the face of faults that may make parts of the network inaccessible. The protocols should gracefully degrade in face of network faults; e.g., in the presence of network faults, the routes selected may not be as good, the connection set up may take a little more time, or resource allocation may be less efficient. Current investigation includes distributing the state information (e.g., about Target Sets, Channels, and Sharing Groups) in keeping with the *fate-sharing* principle; an important design goal is to come up with a distribution that is independent of the underlying routing mechanism, or at least find distributions that are applicable for classes of routing algorithms. For each of these distributions, the plan is to extensively study the respective modifications and extensions to the Tenet protocols. Gupta and Rothermel will also study using forward error correction (FEC), and dispersing the FEC'ed data among disjoint routes.

Dynamic resource migration. Riccardo Bettati and **Amit Gupta** are investigating dynamic, adaptive resource allocation schemes for real-time multicast connections over internetworks. They have devised a mechanism (called Dynamic Resource Migration (DRM)) that allows nodes on the same connection to dynamically renegotiate previously allocated resources; this mechanism does not violate QoS guarantees both during and after its renegotiation and reallocation steps, and this allows DRM to be used as the basic mechanism for fully distributed, adaptive global reallocation schemes. Specifically, they are analyzing distributed, gradient-based algorithms for this purpose. They are also studying the stability issues, especially when these mechanisms interact with the routing subsystem.

Security issues. Rolf Oppliger and **Amit Gupta** are investigating several security issues in the context of the Tenet Scheme 2. They are developing security extensions that use cryptographic techniques to provide authentication and authorization, as well as data confidentiality and integrity services. The objective is to design a scheme that (a) scales well, (b) is robust to faults in the network, and (c) resists to various kinds of attacks.

Dynamic Traffic Modification. Rebbie Moon will be implementing Dynamic Traffic Modification (DTM) in Suite 2. DTM, which is based on the Dynamic Connection Management work by Colin Parris and Hui Zhang, allows changes to be made to the traffic characterization of an established channel. DTM will calculate the effects of a changed traffic characterization, and reduce or augment the resources allocated to the established channel. If a reduction is indicated, DTM will free up the excess resources at the affected nodes. If increased resources are needed, the DTM module will perform admission control tests at the affected nodes to see if the requested augmentation can be accommodated. DTM takes into account the existing allocation of resources during admission tests, and will not duplicate allocations at any time.

Evaluation of a multimedia application on Suite 2. Paola Rossaro is evaluating the behavior of the Tenet protocol suite Suite 2 from the end-user multimedia application perspective. In particular, she is considering the video conferencing tool *Vic* (described in a later section). *Vic* employs an extensible architecture to support multimedia, multi-party communication and its behavior in a real time environment should provide interesting and useful insights in the capabilities of the Tenet protocols.

3. Continuous-Media Networking Applications

Priority Encoding Transmission. Andres Albanese Andres Albanese and **Bernd Lamparter**, in collaboration with Michael Luby, Johannes Bloemer, and Jeff Edmonds from the ICSI Theory Group, has introduced a novel method for sending messages simultaneously over lossy networks. The new method, called PET (Priority Encoding Transmission), allows a user to encode with selective redundancy each part of a message to be transmitted over a lossy medium. Based on the given level of priorities, the sender uses the system to break 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 resilient video transmission over the Internet. Bernd Lamparter and Malik Kalfane implemented PET within the frame work of vic. vic is a video conferencing tool deloped by Steve MCCanne, it is used in the MBone, the multicast network on top of the internet. With vic MPEG streams are send PET protected over the network. Because of the fragile nature of MPEG the added redundancy is a great improvement to the picture quality.

An MBone VCR tool. Wieland Holfelder designed and developed a VCR application for use on the Internet's Multicast Backbone (MBone). This application uses a VCR interface to control recording and playback of sessions sent over the MBone. VCR stores these session by synchronizing different media streams based on information provided by RTP (the Real Time Transport Protocol). It also allows indexing and random access within recorded sessions. A command language enables limited programming features e.g. to schedule recordings or playbacks at a later point of time. The VCR tool is currently alpha tested and will be available via anonymous ftp from icsi.berkeley.edu soon.

An 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.

Application characterization and related performance guarantees. Ed Knightly is collaborating with the Distributed Computing Group at Sandia National Laboratories to investigate applications requiring quality of service guarantees from the network. The applications include a video conferencing system based on DAVE (Distributed Audio Video Environment) and an MPEG video distribution system. The study of the traffic characteristics and performance requirements of real applications provides a better understanding of the multiplexing behavior of actual applications running on the testbeds. For example, this work seeks to quantify the gains in network utilization achievable with the recently proposed extensions to the Tenet deterministic and statistical models. Development and experimentation with the applications and with network support for the applications are under way over both Xunet 2 and a local area ATM network at Sandia.

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 tool called *vic*, which employs an extensible architecture to facilitate experimentation with new video coding schemes and conference control protocols. *Vic* is in production use in the MBone and is available via anonymous ftp from tenet.cs.berkeley.edu.

Toolkit development of applications over the Tenet network suites. Created by the Berkeley Plateau Multimedia Research Group led by Prof. Lawrence A. Rowe, the Continuous Media Toolkit (CM Toolkit) is a flexible toolkit which facilitates development of local and distributed continuous media applications. **Peter Staunton** has added new objects to the toolkit which allow the developer to employ the guaranteed network services of Tenet Suite 1. CM Player, a movie playback application, has been extended to make use of these new toolkit objects. There are plans to add further toolkit objects for Tenet Suite 2, so that multi-party continuous media applications can be created.

Private conferencing over the MBone. Rolf Oppliger and **Andres Albanese** have studied the security requirements of multi-party communication and conferencing systems in general, and applications for the Internet Multicast Backbone (MBone) in particular. They have come up with a scheme to support both public conferencing with participants registration and authentication, and private conferencing over the MBone. The scheme is being implemented in a session registration tool (SRT) for the MBone.

4. Internet Studies

Quality of service and Asynchronous Transfer Mode in IP Internetworks. 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 or routers. However, this strategy does not support the notion of providing different qualities of service to different IP conversations. By examining the network- and transport-layer headers of an incoming datagram, an IP gateway at the edge of an ATM cloud can route that datagram onto a per-conversation ATM virtual circuit. Various headers, such as well-known TCP ports, can be used to help determine the parameters for new virtual circuits. In effect, this amounts to synthesizing an appropriate QOS-guaranteed connection for a stream of datagrams while they traverse the ATM subnet. While this arrangement is not intended to provide IP conversation with end-to-end QOS guarantees, it may improve the performance of IP conversations by giving them ATM virtual circuits well-suited to their requirements. The use of quality of service at the ATM layer also has implications for the strategies used for multiplexing and virtual circuit management. A simulation to investigate QOS, multiplexing, and virtual circuit management policies is being constructed, to evaluate different strategies under various workloads.

Improving World Wide Web latency. Venkata Padmanabhan has been working with Jeffrey Mogul of DEC's Western Research Laboratory on improving WWW latency. They have designed and implemented a modified HTTP protocol that cuts down the latency perceived by the client to less than half in some cases. Their client and server interoperate fully with unmodified implementations. They have also investigated predictive prefetching techniques to further reduce the latency perceived by clients. Their results show that latency can be reduced substantially at the cost of a moderate increase in network traffic.

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, this assumption grossly underestimates burstiness, and that fractal models reflect the traffic much more faithfully. He is currently developing tools for dynamically measuring network characteristics.

5. Tenet Group Members

The Tenet Group consists of students and faculty of the Department of Electrical Engineering and Computer Science at U.C. Berkeley as well as visiting scholars and post-doctoral researchers, both at UCB and at ICSI. The group is led by **Professor Domenico Ferrari** and includes the following graduate students: **Jacqueline Chang**, **Rahul Garg**, **Amit Gupta**, **Wendy Heffner**, **Fukiko Hidano**, **Edward Knightly**, **Steven McCanne**, **Bruce Mah**, **Rebbie Moon**, **Mark Moran**, **Venkata Padmanabhan**, **Vern Paxson**, **Peter Staunton**, and **Brian Whetten**. In addition, the group includes an undergraduate student, **Hoofar Razavi**, and a member of the U.C. Berkeley research support staff, **Keith Sklower**.

The following visiting scholars and post-doctoral researchers are also members of the Tenet Group: **Riccardo Bettati**, a Swiss post-doctoral fellow at ICSI; **Wieland Holfelder**, from the University of Mannheim, Germany; **Kazuhiro Honda**, a UCB Visiting Industrial Fellow from Sony Corporation; **Bernd Lamparter**, a German post-doctoral fellow at ICSI; **Tetsuro Mano**, a UCB Visiting Industrial Fellow from the Software Development Center, Hitachi, Ltd., Yokohama, Japan; **Rolf Oppliger**, a Swiss post-doctoral fellow at ICSI; **Hiroki Rikiyama**, a UCB Visiting Industrial Fellow from NEC Corporation, Kanagawa, Japan; **Kurt Rothermel**, from the University of Sttutgart, Germany; **Paola Rossaro**, a post-doctoral fellow at ICSI from the University of Pavia, Italy; **Doug Shepard**, from Lancaster University, U.K.; **Rainer Storn**, a German post-doctoral fellow at ICSI; **Hirohisa Wakai**, from the Information Technology Research Laboratories, Sharp Corporation, Nara, Japan; **Raj Yavatkar**, from the University of Kentucky; and the leader of ICSI's Networks Group, **Andres Albanese**.

6. Tenet Group Publications

Further information about the Tenet group's research may be obtained through the World Wide Web at URL **http://tenet.berkeley.edu**. This site contains pointers to the group's projects, testbeds, and publications as well as the home pages of group members. The Tenet Group publications since January 1994 are listed below. Unless otherwise noted, the publications are available by anonymous FTP to **tenet.berkeley.edu** and through the WWW site above. 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 an e-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 94720-1776. Electronic mail may be sent to ferrari@cs.berkeley.edu.

A. Albanese, J. Bloemer, J. Edmonds, and M. Luby, ‘‘Priority Encoding Transmission’’, Technical Report TR-94-039, International Computer Science Institute, Berkeley, CA, August 1994.

A. Banerjea, ‘‘Fault Management for Realtime Networks’’, Ph.D. dissertation, University of California, Berkeley, CA, December 1994.

A. Banerjea, D. Ferrari, B. Mah, M. Moran, D. Verma, and H. Zhang, ‘‘The Tenet Real-Time Protocol Suite: Design, Implementation, and Experiences’’, Technical Report TR-94-059, International Computer Science Institute, Berkeley, CA, November 1994.

A. Banerjea, E. Knightly, F. Templin, and H. Zhang, ‘‘Experiments with the Tenet Real-Time Protocol Suite on the Sequoia 2000 Wide Area Network’’, *Proc. ACM Multimedia'94*, San Francisco, CA, pp. 183-192, October 1994; also Tech. Rept. TR-94-020, International Computer Science Institute, Berkeley, CA, April 1994.

R. Bettati, D. Ferrari, A. Gupta, W. Heffner, W. Howe, M. Moran, Q. Nguyen and R. Yavatkar, ‘‘Connection Establishment for Multi-party Real-time Communication’’, in *Proc. IEEE Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV) '95*, Durham, NH, pp. 255-266, April 1995.

R. Bettati and A. Nica, ‘‘Real-Time Networking over HIPPI’’, *Proc. Third Workshop on Parallel and Distributed Real-Time Systems*, Santa Barbara, CA, pp. 76-83, April 1995.

J. Domingo-Pascual, A. Albanese, and W. Holfelder, ‘‘Emulation of Traffic Congestion on ATM Gigabit Networks’’, Technical Report TR-95-005, International Computer Science Institute, Berkeley, CA, February 1995.

D. Ferrari, ‘‘A New Admission Control Method for Real-Time Communication in an Internetwork’’ in S. Son, Ed., *Advances in Real-Time Systems. Ch. 5, pp. 105-116, Prentice-Hall, Englewood Cliffs, NJ, 1995.*

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; also in *Computer Networks and ISDN Systems*, vol. 26, pp. 1267-1280, July 1994.

D. Ferrari, A. Gupta and G. Ventre, ‘‘Distributed Advance Reservations of Real-time Connections’’, in *Proc. IEEE Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV) '95*, Durham, NH, pp. 15-26, April 1995.

R. Garg, ‘‘Characterization of Video Traffic’’, Technical Report TR-95-007, International Computer Science Institute, Berkeley, CA, January 1995.

A. Gupta and D. Ferrari, ‘‘Resource Partitioning for Multi-Party Real-Time Communication’’, Technical Report TR-94-061, International Computer Science Institute, Berkeley, CA, November 1994; to appear in *IEEE/ACM Transactions on Networking*, October 1995.

A. Gupta, W. Howe, M. Moran, and Q. Nguyen, ‘‘Evaluation of Resource Sharing Benefits’’, Technical Report TR-94-051, International Computer Science Institute, Berkeley, CA, October 1994.

A. Gupta, W. Howe, M. Moran, and Q. Nguyen, ‘‘Resource Sharing for Multi-Party Real-Time Communication’’, *Proc. IEEE INFOCOM'95*, Boston, MA, April 1995.

A. Gupta, W. Howe, M. Moran, and Q. Nguyen, ‘‘Scalable Resource Reservation for Multi-Party Real-Time Communication’’, Technical Report TR-94-050, International Computer Science Institute, Berkeley, CA, October 1994.

W. Heffner, ‘‘Scaling Issues in the Design and Implementation of the Tenet RCAP2 Signaling Protocol’’, Technical Report TR-95-022, International Computer Science Institute, Berkeley, CA, May 1995.

D. Kandlur, K. Shin, and D. Ferrari, ‘‘Real-Time Communication in Multihop Networks’’, *IEEE Trans. Parallel and Distributed Systems*, vol. 5, n. 10, pp. 1044-1056, October 1994. (not available on line)

E. Knightly, R. Mines, and H. Zhang, ‘‘Deterministic Characterization and Network Utilizations for Several Distributed Real-time Applications’’, *Proc. of IEEE WORDS'94*, Dana Point, CA, October 1994.

E. Knightly, A. Banerjea, F. Templin, H. Zhang, B. Mah, and D. Ferrari, ‘‘A Demonstration of the Tenet Protocol Suite’’, *Video Program of ACM Multimedia'94*, San Francisco, CA, October 1994.

E. Knightly and H. Zhang, ‘‘Traffic Characterization and Switch Utilization using a Deterministic Bounding Interval Dependent Traffic Model’’, *Proc. of IEEE INFOCOM‘95*, Boston, MA, pp. 1137-1145, April 1995.

E. Knightly, D. Wrege, J. Liebeherr, and H. Zhang, ‘‘Fundamental Limits and Tradeoffs of Providing Deterministic Guarantees to VBR Video Traffic’’, to appear in *Proc. ACM SIGMETRICS ’95*, Ottawa, Ontario, May 1995.

C. Leicher, ‘‘Hierarchical Encoding of MPEG Sequences Using Priority Encoding Transmission (PET)’’, Technical Report TR-94-058, International Computer Science Institute, Berkeley, CA, November 1994.

J. Liebeherr, D. Wrege, and D. Ferrari, ‘‘Admission Control in Networks with Bounded Delay Services’’, Technical Report TR-94-033, International Computer Science Institute, Berkeley, CA, August 1994.

P. Manzoni, ‘‘Introducing resources management in IP-based nodes’’, Technical Report TR-94-040, International Computer Science Institute, Berkeley, CA, October 1994.

V. Padmanabhan, and J. Mogul, ‘‘Improving HTTP Latency’’, *Proc. 2nd International World Wide Web Conference*, Chicago, IL, pp. 995-1005, October 1994.

C. Parris and A. Banerjea, ‘‘An Investigation into Fault Recovery in Guaranteed Performance Service Connections’’, *Proc. SUPERCOMM / ICC’94*, New Orleans, LA, pp. 175-181, March 1994; also Technical Report TR-93-054, International Computer Science Institute, Berkeley, CA, October 1993.

C. Parris, H. Zhang, and D. Ferrari ‘‘A Dynamic Management Scheme for Real-Time Connections’’, *Proc. IEEE INFOCOM‘94*, Toronto, Canada, pp. 698-707, June 1994.

C. Parris, H. Zhang, and D. Ferrari, ‘‘Dynamic Management of Guaranteed Performance Multimedia Connections’’, *ACM/Springer-Verlag Multimedia Systems*, vol. 1, pp. 267-283, 1994.

V. Paxson, ‘‘Empirically-Derived Analytic Models of Wide-Area TCP Connections’’, *IEEE/ACM Transactions on Networking*, vol. 2, n. 4, pp. 316-336, August 1994.

V. Paxson, ‘‘Growth Trends in Wide-Area TCP Connections’’, *IEEE Network*, vol. 8, n. 4, pp. 8-17, July/August, 1994.

V. Paxson and S. Floyd, ‘‘Wide-Area Traffic: The Failure of Poisson Modeling’’, *Proc. ACM SIGCOMM’94 Conf.*, London, England, pp. 257-26, September 1994.

R. Storn, ‘‘Modeling and Optimization of PET-Redundancy Assignment for MPEG Sequences’’, Technical Report TR-95-018, International Computer Science Institute, Berkeley, CA, May 1995.

B. Whetten, S. Steinberg, and D. Ferrari, ‘‘The Packet Starvation Effect in CSMA/CD LANs and a Solution’’, *Proc. of IEEE Local Computer Networks*, Minneapolis, MN, pp. 206-217, October 1994.

B. Whetten, T. Montgomery and S. Kaplan, ‘‘A High Performance, Totally Ordered Multicast Protocol’’, Theory and Practice in Distributed Systems K.P. Birman, F. Mattern, A. Schiper (Eds) Springer Verlag LNCS 938.

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, pp. 211-220, 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, pp. 295-304, May 1994.

H. Zhang and E. Knightly, ‘‘A New Approach to Support VBR Video in Packet-Switching Networks’’, in *Proc. IEEE Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV) ’95*, Durham, NH, pp. 275-286, April 1995.

H. Zhang and E. Knightly, ‘‘Comparison of Rate-Controlled Static Priority and Stop-and-Go’’, Technical Report TR-94-048, International Computer Science Institute, Berkeley, CA, August 1994; also to appear in ACM/Springer-Verlag Multimedia Systems Journal, September 1995.

H. Zhang and D. Ferrari, ‘‘Rate-Controlled Service Disciplines’’, to appear in *Journal of High Speed Networks*, accepted in February 1994.