delim $$

Multimedia network protocols: where are we?

Domenico Ferrari

The Tenet Group
University of California at Berkeley
and International Computer Science Institute

ABSTRACT

Protocols for multimedia communication are needed to integrate into a single network services intended to satisfy the different requirements of multiple types of traffic. An essential prerequisite for designing those protocols is that a service model be chosen; in other words, the services to be offered by the network must be selected and specified in detail. This paper presents the service models proposed, or being developed, by two groups within the Internet community, by the ATM community, and by the Tenet Group; it also compares them, focusing on their common characteristics and their possible convergences; while doing this, it reveals some of the characteristics of the services the first integrated-services networks are likely to offer. The services referred to in this paper are those at the network and transport layers, which support the services to be offered to the system's end users.

Keywords: multimedia, protocols, integrated-services networks, Internet, ATM, Tenet, service models.

1. Introduction

One of the most active areas of networking research in the last few years has been that of protocols and services designed to transmit continuous-media (e.g., audio, video, image sequence) traffic. The less accurate, but more popular, term in current use for them is *multimedia protocols.* This term is often applied also to designate protocols that, by virtue of some of their properties, are used to carry continuous-media traffic even though they have not been designed to meet its quality requirements; such protocols include, for example, XTP, RTP, and IP multicast. In this paper, we will not discuss any of these protocols, in spite of their usefulness and importance in many current multimedia applications. We shall restrict our attention to those protocols that are intended to satisfy the Quality of

Service (QoS) needs of continuous-media transmission in packet-switching or cell-switching networks, and to the architectures of which they will be a part.

Our objective is to try to detect the current trends in the design of multimedia protocols, especially in the service models they implement. Section 2 discusses the notion of service model for integrated-services networks, where these protocols will have to operate. We examine the service models proposed by three communities of designers, and their initial instantiations, in Section 3. To a discussion of the approaches that seem to be prevailing within the Internet community and the ATM Forum, we add a description of the service model adopted by a much much smaller group, the Tenet Group, which has the advantage of having completed more than one year ago the implementation of its first suite of multimedia protocols and of being close to demonstrating its second suite. Experience with some prototype multimedia protocols is therefore already being gathered; this activity will become more intensive as new protocols are deployed, and will provide us with the real-world feedback we need to design the integrated-services networks of the future. A brief summary of the currently discernible trends in service models is given in Section 4.

2. Service Models for Integrated-Services Networks

An *integrated-services network* is a network designed to carry various types of traffic, and to offer various services in an attempt to satisfy the different requirements of those types. Different types of traffic have different QoS requirements; for example, an interactive voice packet is to be delivered with a relatively short delay and a small delay variation; a file data needs reliable transmission and a reasonable transfer time; and so on. Thus, an integrated-services network must be governed by protocols that are sensitive to the QoS requirements of the various types of traffic and try to satisfy each one of them as much as possible.

The *service model* of such a network describes the services the network is able to offer in terms of the requirements these services can satisfy. Along with the QoS characteristics referred to above, other requirements may be included, such as the ability to build a private virtual network, the ability to reserve services in advance (particularly important for multi-party applications such as a multimedia conference or lecture), the requirement of understandable and affordable pricing policies, and so on.

The services and service models we discuss in this paper are those at the network and transport layers; many of them are only indirectly useful to the end users, but they are very important nevertheless. For example, multicast is only an efficient way of implementing point-to-multipoint connections, which could be obtained, though much more wastefully, by multiple unicasts; however, if the network offers a multicast service, all end users will indirectly benefit, as this will allow the network to support many more of them, and will therefore drastically reduce the probability of rejection of access requests. Other services of this type include sender-initiated setup, receiver-initiated setup, and third-party-initiated setup, given any one of which any connection can be established, though more or less easily depending on whether the setup method is the "natural" one for the application at hand; *media scaling,* the ability by the network to request and obtain that sources reduce their rates to make room for more connections; *fault recovery,* which for multimedia traffic has even more stringent requirements than for more traditional types of traffic.

All the service models discussed in this paper recognize that an integrated-services network will at the very least provide two extreme types of service, which may be called *guaranteed deterministic* and *best effort,* respectively. The former type of service offers deterministic guarantees on given performance indices; these guarantees are usually in the form of hard upper or lower bounds, and are satisfied even in the worst case, as long as there are no failures in the network. Best-effort service, on the other hand, does not offer any guarantees; while it is possible to build services guaranteeing reliability on top of best-effort services, no performance guarantees can be obtained by such an approach. A best-effort service need not restrict access to the network, whereas a guaranteed service may have to incorporate an *admission control* algorithm to make congestion impossible. The applications that may need, or benefit from, guaranteed bounds are collectively called *real-time;* a subclass of real-time applications is that of *multimedia applications. Non-real-time applications* do not have stringent performance requirements.

The designer of an integrated-services network is faced with some fundamental initial decisions. What requirements will the network be able to satisfy? Which of them will need to be satisfied efficiently? How are the requirements going to be expressed? How are they going to be communicated to the appropriate protocols? The answers designers give to these questions reflect their philosophies and expectations about multimedia traffic; in the not-too-distant future, they should also reflect the results of the observations to be collected from the first integrated-services testbeds and production networks.

For instance, a designer could decide that the network will adopt the Internet model for data, text, and still images; then, the protocols that deal with that portion of the traffic could be TCP, UDP, and IP, or offer the same services as these protocols; however, the model could also be modified if the designer felt that this would improve some of its aspects, or make it more consistent with the other parts of the service model for the entire network.

On the continuous-media side, an important question is whether and in what form the quality of service should be guaranteed. When designing the architecture of the network, and the individual protocols, a decision is to be made about whether there ought to be a single protocol stack or multiple stacks. Also, the designer is confronted with the issue of the lowest layer that must know the type of each packet: should the network layer, for example, know that a given packet contains video information, or should it only be told about the quantitative, parametric performance requirements of such a packet?

The service models being proposed within the Internet community, the ATM community, and the Tenet Group are examined in the next section, in an effort to get a glimpse of what service models the first integrated-services networks might offer.

3. Some Proposed Service Models

3.1. Internet service models

The service model offered by the Internet is well-known: all traffic is treated as data, and transmitted with a best-effort approach; access to the network is never refused; and the communication model is connectionless and datagram-based. Since the late 1970's, however, a small subcommunity within the Internet community has been working on adding a stream-oriented paradigm to the datagram one for the purpose of providing performance guarantees to real-time, continuous-media traffic. This movement resulted in the specification of ST, the Stream Protocol, and of its successor, ST-II[1]. A *stream* is a unicast or multicast connection with performance guarantees.

These protocols are supposed to coexist with IP at the internetwork layer, and include a stream establishment protocol (SCMP). They expect suitable transport-layer protocols above them; those on top of ST were media-specific protocols (e.g., a video protocol, an audio protocol), an idea that has fallen into disfavor and has been abandoned, at least for the moment.

ST-II offers its clients a rich set of traffic and QoS parameters; the set includes most of the performance parameters proposed so far in the literature (throughput, delay, error rate) as well as others, but the protocol specification does not say what should be done with them, as resource allocation and management are left to the implementor. The danger of this rich and liberal approach is that the service model might vary with the implementation, depending on which parameters are supported and how.

Several implementations of ST-II are in existence[2,3], and other services have been added (e.g., at the transport layer) to those provided by ST-II. For example, the HeiTS (Heidelberg Transport System)[3] suite developed at the IBM European Networking Center, which includes ST-II, offers media scaling, reliability, and receiver-initiated stream setup supplementing ST-II's sender-initiated procedure.

A different service model[4,5] for a multimedia network, recently proposed within the Internet community, is based on an interesting classification of applications, which is reported in Figure 1. Non-real-time traffic is that produced by *elastic applications,* i.e., those that use the received data immediately, and wait for it (though this does not mean that they are delay-insensitive). Elastic applications can in turn be subdivided into *interactive* (with the two subclasses *burst,* representing the needs of applications like telnet, X, NFS, and *bulk,* e.g., ftp) and *asynchronous bulk,* such as e-mail. The real-time portion of an integrated-services network's traffic is due to *real-time applications,* which can be classified as being either *tolerant* or *intolerant;* the latter are those that cannot tolerate any losses of quality. Another useful distinction is that between *adaptive* and *rigid* applications; the former are the applications that can vary the playback delay at the receiver in response to changes in the network load. The playback delay is the difference between the time the received information is presented to the user and the time it should have been received in an ideal transmission.

The services to be offered by an integrated-services network according to this proposal are:

- guaranteed deterministic service for intolerant real-time applications;
- predictive service[6] for tolerant real-time applications; this service controls the admission of real-time *flows*[+] on the basis of current measured traffic levels, instead of considering at setup time the amounts of resources requested by the existing flows; thus, the needs of the incoming flow will be satisfied as long as future traffic does not exceed by too large an amount the levels measured when this flow was admitted;
- ASAP service for elastic applications; this is best-effort service with priorities, the highest to be given to interactive burst traffic, the lowest to asynchronous bulk traffic.

+ A flow is like a stream, but is not always characterized by mathematically provable guarantees.

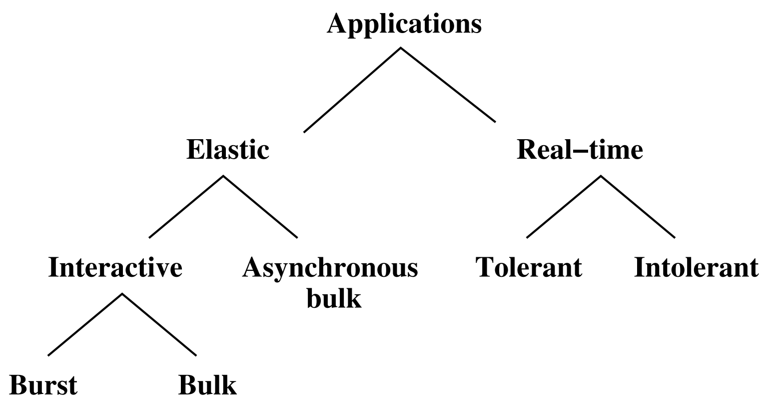

Figure 1. Classification of applications[4].

The main QoS parameter in this approach is delay, for which an upper and a lower bound are to be specified. Bandwidth (i.e., throughput) is a traffic parameter, as well as a QoS parameter; delay jitter is not included, since it can be taken care of by buffering in the receiver if it is not sufficiently bounded by the network. Note that, here and in the rest of this paper, the term "delay" refers to end-to-end network delay.

The approach proposes also, as an additional service to be provided, *link sharing,* i.e., the possibility of partitioning the bandwidth of a link among various entities (organization, classes of traffic, protocols, and so on).

Neither a specific architecture nor particular protocols that would implement the new service model have been presented yet. The ReSerVation Protocol[7] (RSVP), which is being developed, might become part of such an implementation, though it is expected to run in conjunction with the next version of IP. Unfortunately, this new version, IPng, is unlikely to offer the services this service model proposal lists as essential.

3.2. ATM service models

The ATM community consists of the ATM committees within the International Telecommunication Union (ITU), the committees of the ATM Forum, and a large number of researchers and developers in academia, industry, and governmental laboratories. At this time, there seems to be no single, universally agreed-upon service model for ATM networks; such a model is expected to emerge gradually from the various specification and standardization activities. A number of its aspects have, however, already been specified, and some even standardized; others are now being discussed in the community; thus, an educated guess about the services that will be offered by the first integrated-services ATM networks can probably be made with reasonable accuracy.

The problem is not what services ATM is capable of offering, as ATM can in principle provide a wide range of services; the problem is which of the services it is able to support will really be supported, how, and when. The ATM protocol by itself provides connection-oriented, in-sequence, unreliable, and point-to-point (or point-to-multipoint) transport. Different services to different types of traffic are offered by various ATM Adaptation Layer (AAL) protocols running on top of ATM. The ITU (formerly CCITT) has defined four classes of applications, to which there correspond four types of traffic defined by the ATM Forum and several different AAL's:

- Class A: constant bit rate (CBR) applications;
- Class B: variable bit rate (VBR) applications;
- Class C: connection-oriented data or available bit rate (ABR) applications;
- Class D: connectionless data or unspecified bit rate (UBR) applications.

Class A traffic is to be supported by AAL1; class B by AAL2; classes C and D by AAL3/4 or by AAL5. The AAL type to be used must be specified in the SETUP message of the signalling protocol, Q.2931. A standard for this protocol has not been published yet, but most of the features its first specification will have are known, and the ATM Forum's version (Version 1) will be referred to here[8]. A comparison with the characteristics of ITU's Release 1 draft recommendation[9] can be found in a paper by Graf and Stuettgen[10]. The QoS parameters to be guaranteed for class A and class B traffic are bounds on cell delay, cell delay variation (i.e., jitter), and cell loss ratio.

ABR traffic will not be offered any guarantees, except perhaps a minimum amount of bandwidth; it will, however, be flow-controlled so that cell losses will be small even in presence of large and sudden bursts[11]. No guarantees are expected to be given to ABR or UBR data traffic, though ABR will receive substantially better treatment, while UBR traffic will essentially get a best-effort service. In principle, however, even an ABR client could get a "busy signal" because of the possible depletion of the amount of bandwidth set aside to guarantee a minimum rate to each ABR virtual circuit, or, less likely, because of a shortage of virtual circuit identifiers.

Other services to be offered by Version 1 of Q.2931 include multicast connections with dynamic joining and leaving, bidirectional unicast setup, parallel setup, and virtual paths, which can be used to partition the network's resources for, e.g., link sharing.

3.3. The Tenet service model

The Tenet Group, since the beginning of its research[12], has been concerned with providing performance guarantees in large, heterogeneous packet-switching internetworks; this emphasis is indeed the same as that of the Internet community, and differs quite noticeably from that of the ATM community. However, the group has always been interested in mathematically provable guarantees[13], which can be obtained through admission control, appropriate scheduling disciplines, and protection of guaranteed-performance connections, as well as on per-connection, per-service-type pricing. Because of this focus, the Tenet Group reached (independently and at about the same times) the same conclusions as the ATM community and the ST "camp" in the Internet community; most of these conclusions have recently been adopted also by several leading Internet researchers of the non-ST persuasion, while others are still not accepting them.

The architecture proposed by the Tenet Group for an integrated-services internetwork is shown in Figure 2. Applications are of two types: *non-real-time* and *real-time.* Non-real-time applications use the Internet protocols (or any other protocols designed for those applications); real-time ones the Tenet protocols, which in both Tenet suites (Suite 1[14] and Suite 2[15]) consist of the Real-Time Channel Administration Protocol (RCAP), the Real-Time Message Transport Protocol (RMTP), and the Real-Time Internetwork Protocol (RTIP).

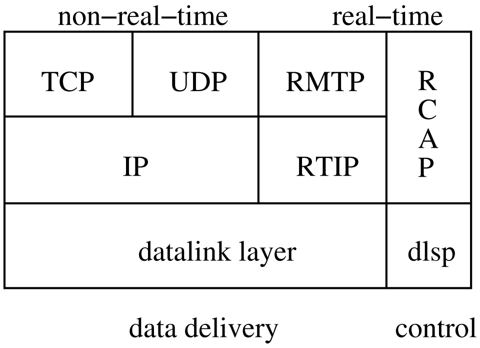

Figure 2. The architecture of an integrated-services network according to the Tenet approach[13]; **dlsp** is the datalink signalling protocol; some control protocols have been omitted.

RCAP provides real-time channel establishment and teardown services. The traffic and QoS parameters that are to be specified by a real-time client at the upper interface to RCAP in Suite 1 are listed in Table 1. $D sub max$ is the end-to-end delay bound; $Z sub min$ is the lower bound of the probability that the delay of a message is greater than $D sub max$; since $J sub max$ is the end-to-end delay jitter bound, the delays of all messages in the deterministic case will be between $D sub max - J sub max$ and $J sub max$; and $W sub min$ is the lower bound of the probability that a message is lost in the network because of a buffer overflow. Among the traffic parameters, $X sub min$ is the inverse of the peak rate, and $X sub roman ave$ the inverse of the worst-case average rate[++]; *I* is the interval that has been used to compute $X sub roman ave$; and $S sub max$ is the maximum length of a message to be presented to the transport layer. In Suite 2, the parameters are the same, but the values of the first two traffic parameters and of all the QoS parameters may be replaced by ranges of values to decrease the probability of a "busy signal". The presence of probabilities among the QoS parameters means that the Tenet suites can provide statistical services as well as deterministic ones. Statistical services represent an "intermediate" type of service between the non-real-time and deterministic (or hard) real-time extremes. Such a service is expected to be less expensive than its deterministic counterpart, hence of interest to tolerant applications; it may therefore be regarded as the Tenet equivalent of the predictive service mentioned in Section 3.1, which, however, does not offer any guarantees. In the ATM model, ABR service may be considered as being of intermediate type; it is also possible that some of the QoS classes will include probabilistically specified bounds, which could permit some controlled overbooking of resources for VBR applications.

Suite 1 has been running for almost two years on several testbeds; the source code of our implementation is being distributed free of charge for research and educational purposes. Suite 2, a prototype version of which has been demonstrated in March 1995, will be able to provide a much richer set of services, including:

- multicast connections with dynamic joining and leaving;

++ For correct admission control, traffic models must be bounding, and on the pessimistic side.

Table 1. QoS and traffic parameters in RCAP (Tenet Suite 1)

QoS parameters

- $D sub max$: the maximum delay
- $Z sub min$: the minimum probability of timely delivery
- $J sub max$: the maximum delay jitter
- $W sub min$: the minimum probability of no loss due to buffer overflow

Traffic parameters

- *$X sub min$* : the minimum intermessage time
- *$X sub roman ave$* : the minimum average intermessage time
- *I* : the averaging interval
- *$S sub max$* : the maximum message size

- sender-initiated, receiver-initiated, and third-party connection setup (Suite 1, in which establishment is sender-initiated, could be enabled to set up receiver-initiated, as well as duplex, connections with a very simple code modification);
- resource partitioning for real-time traffic, which permits link sharing and virtual private networks[16];
- advance reservations of real-time connections[17];
- media scaling[18].

The guarantees promised by the Tenet suites are valid if the client obeys at all times the worst-case traffic description it has given the network and in the absence of failures. The environment in which such end-to-end guarantees can be obtained consists of an internetwork composed of networks that either guarantee performance (e.g., ATM networks with Q.2931 signalling) or can be put under the Tenet protocols' control; the latter is the case of a network into whose switches or routers one can load and run RTIP and a datalink layer version of RCAP[19].

4. Summary

We have briefly reviewed the service models being proposed for integrated-services networks within the Internet community, within the ATM community, and by the Tenet Group. The models differ from one another in many ways, but share some of the fundamental characteristics, and seem to allow for further convergence even in their more specific features.

First of all, there is agreement about the necessity that an integrated-services network provide at least the two extreme services: non-real-time or "best-effort" on one side, and deterministically guaranteed-performance service on the other. Also, because of the low utilization of resources expected to characterize the latter, the need is felt for at least one intermediate, less expensive service type that caters to tolerant real-time applications. The proposals in this area are different: predictive, non-guaranteed service in the Internet approach, statistically guaranteed service in the Tenet solution, and VBR and/or ABR service in the ATM proposals. Only thorough experimentation with real applications on real networks can help us decide which of these approaches is (or are) the best and should be included in the design of an integrated-services network. A crucial question to be answered is the one about tolerance: how tolerant will tolerant applications really be? It is easy to become spoiled by good quality, and most consumers are accustomed to the quality of broadcast TV and CD audio; will they accept anything less? What will have to be the difference in price between deterministic and intermediate services to convince clients to prefer the latter, if there are noticeable differences in quality? And will there be noticeable differences in quality often enough for this last question to be meaningful?

Another set of questions are concerned with the distinction between rigid and adaptive applications, which is important for the success of a media scaling service when it refers to bandwidth adjustments, and of predictive service when it refers to playback-delay modifications. Adapting the rate or the delay to an increase in real-time traffic may cause a reduction in quality that certain applications may be willing to tolerate, provided the reduction is bounded; also, in mobile wireless contexts, adaptivity might be inevitable or highly desirable, at least with current and projected technology; however, the existence of some delay-adaptive applications on the Internet cannot be taken as a proof of the usefulness of such adaptivity in future networks, in which those kinds of applications will not have to be adaptive to protect themselves from the effects of congestion, as they will be shielded by some of the services being discussed here.

The parametric approach to QoS requirements and traffic description seems to be prevailing over other possibilities such as media-specific protocols. The solution that had been adopted by the ATM Forum until very recently, based on *QoS classes*, differed from the others, which are more general and flexible. It is in fact easy to build a QoS class menu on top of the parametric approach as implemented in ST-II or RCAP, present it to the application, and modify and extend it whenever needed; after all, the QoS and traffic parameters in Table 1, for example, are those to be given to RCAP at the transport layer, not at the application layer. We believe that defining QoS classes at the datalink or network or transport layer unduly constrains the extension of real-time services to new applications, and forces all applications into (sometimes uneasy) conformance with the designer's choices.

As for the parameters themselves, there seems to be agreement that delay and throughput (peak and average) are the most important QoS parameters, though throughput is usually considered as a traffic parameter since it is under the sender's control (but this traffic parameter specification implicitly asks the network to allocate a bandwidth commensurate with the desired throughput). Another important QoS parameter is the loss rate, while delay jitter is slightly more controversial; however, the controversy is about how and where jitter should be bounded, not about the usefulness of bounding it.

Multicast and resource partitioning are proposed in various forms by all approaches. Other services, such as media scaling and advance reservations, can be found only in some of the models, but probably not because they are deemed useless by the others; rather, we believe this is due to the different stages of development the different service models have reached so far.

The agreements among, and the possible convergences of, the various service models surveyed are meaningful, as they point to services and features that are likely to be found in the first deployments of integrated-services networks. However, those services and features have not been experimented with extensively enough for us to be able to say that they are part of "the right model". Extensive experimentation with multimedia protocol implementations and applications is the next crucial step.

5. Acknowledgements

The author wishes to thank all the past and present members of the Tenet Group for the many lively exchanges of ideas that over the years have contributed to his forming the opinions expressed in this paper. Both the opinions and the errors in the paper are the sole responsibility of the author. He is also indebted to Bruce Mah, Steve McCanne, and Clemens Szyperski for their comments on a draft of the paper, and AT&T Bell Laboratories, Hitachi, and Mitsubishi Electric Research Laboratories for their generous support.

6. References

1. C. Topolcic, Ed., "Experimental Internet Stream Protocol, Version 2 (ST-II)," Request for Comments 1190, Oct. 1990.
2. C. Partridge and S. Pink, "An implementation of the revised Internet Stream Protocol (ST-2)," *Internetworking: Research and Experience,* Vol. 3, pp. 27-54, March 1992.
3. L. Delgrossi, R.G. Herrtwich, and F.O. Hoffmann, "An implementation of ST-II for the Heidelberg Transport System," *Internetworking: Research and Experience,* Vol. 5, pp. 43-69, 1994.
4. S. Shenker, D.D. Clark, and L. Zhang, "A scheduling service model and a scheduling architecture for an integrated-services packet network," preprint, 1993.
5. B. Braden, D. Clark, and S. Shenker, "Integrated services in the Internet architecture: an overview," Internet Working Draft, Oct. 1993.
6. D. Clark, S. Shenker, and L. Zhang, "Supporting real-time applications in an integrated services packet network: architecture and mechanism," *Proc. ACM SIGCOMM '92,* pp. 14-26, 1992.
7. L. Zhang, S. Deering, D. Estrin, S. Shenker, and D. Zappala, "RSVP: a new resource ReSerVation Protocol," *IEEE Network,* Vol. 7, pp. 8-18, Sept. 1993.
8. ATM Forum, *ATM User-Network Interface Specification Version 3.0,* Prentice-Hall, 1993.
9. "ITU-TS Draft Recommendation Q.2931," ITU-TS, SG11/WG2 Temporary Document, Dec. 1993.
10. M. Graf and H.J. Stuettgen, "An internetworking architecture for multimedia communication over heterogeneous networks," in O. Spaniol, A. Danthine, and W. Effelsberg, Eds., *Architecture and Protocols for High-Speed Networks*, Kluwer, pp. 161-178, 1994.
11. Various Authors, Papers on flow control for ABR traffic, *IEEE Network*, Vol. 9, March/April 1995.

12. D. Ferrari, "Real-time communication in packet-switching wide-area networks," Tech. Rept. TR-89-022, International Computer Science Institute, Berkeley, CA, May 1989.

13. D. Ferrari, A. Banerjea, and H. Zhang, "Network support for multimedia: a discussion of the Tenet approach," *Computer Networks and ISDN Systems,* Vol. 26, pp. 1267-1280, 1994.

14. A. Banerjea, D. Ferrari, B. Mah, M. Moran, D. Verma, and H. Zhang, "The Tenet Real-Time Protocol Suite: design, implementation, and experiences," Tech. Rept. TR-94-059, International Computer Science Institute, Berkeley, CA, Nov. 1994.

15. A. Gupta, W. Heffner, M. Moran, and C. Szyperski, "Network support for real-time multi-party applications," *Abstracts of the 4th International Workshop on Network and Operating System Support for Digital Audio and Video,* Lancaster, England, pp. 37-39, Nov. 1993.

16. A. Gupta and D. Ferrari, "Resource partitioning for multi-party real-time communication," Tech. Rept. TR-94-061, International Computer Science Institute, Berkeley, CA, Nov. 1994; to be published in *IEEE/ACM Trans. on Networking,* 1995.

17. D. Ferrari, A. Gupta, and G. Ventre, "Distributed advance reservation of real-time connections," *Proc. 5th International Workshop on Network and Operating System Support for Digital Audio and Video,* Durham, NH, pp. 255-266, April 1995.

18. C. Parris, H. Zhang, and D. Ferrari, "Dynamic management of guaranteed performance multimedia connections," *Multimedia Systems,* Vol. 1, pp. 267-283, 1994.

19. D. Ferrari, "Real-time communication in an internetwork," *Journal of High Speed Networks,* Vol. 1, n. 1, pp. 79-103, 1992.