The Tenet Experience and the Design of Protocols for Integrated-Services Internetworks

Domenico Ferrari

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

ABSTRACT

The research done by the Tenet Group in multimedia networking has reached a point where it may be useful to reflect on the significance of its results for the current debate on how integrated-services internetworks should be designed. Such reflections constitute the main subject of this paper. The principles of the work and the conclusions reached so far by the Tenet researchers are discussed in the light of the conflict between the two major technologies being proposed to build future information infrastructures: namely, the Internet and the ATM technologies. The Tenet approach suggests one feasible way for resolving the conflict to the advantage of all the users of those infrastructures. This paper discusses various fundamental aspects of integrated-services network design: the choice of the service model, the type of charging policy to be adopted, and the selection of a suitable architecture.

1. Introduction

In late 1987, while investigating the properties the operating systems of the future would have to exhibit, my attention was captured by an assumption my co-workers and I were making: that the networks of the future would provide real-time (e.g., continuous-media) users with the performance guarantees they need. Even then, I was deeply convinced that those networks would be based on packet-switching principles; thus, the question that obviously presented itself to my mind was whether ours was a realistic assumption. Is it possible to obtain performance guarantees from packet-switching networks which are prone to congestion? Out of my curiosity about the feasibility of such a seemingly impossible feat came (a) my fascination with what was later called "multimedia networking", (b) a scheme that showed how mathematically provable bounds on throughput, delay, delay jitter, and packet loss could indeed be obtained (later called "Tenet Scheme 0"), and (c) the Tenet Group, which I founded in 1989.

The Tenet Group has now started distributing the source code of its first protocol suite, Suite 1, and demonstrated a still incomplete prototype of its much more advanced suite, Suite 2. In the meantime, the context of our work has changed greatly. What was, even in 1989, the effort of a few, isolated individuals or small groups, has become a major subject for research and development, with imminent and important opportunities for practical use of ideas and results. The design of a suitable architecture and of effective protocols for integrated-services networks has been, from the beginning, a crucial concern for the ATM community, since asynchronous transfer mode was proposed as a viable technology with which broadband integrated-services digital networks could be built. However, more recently, the explosive success of the Internet and of its virtual multimedia incarnation, the Multicast Backbone or MBone, has forcefully put this concern also on the agenda of the Internet Engineering Task Force. The solutions that are being proposed in the two communities reflect their respective cultures and the principles of their respective networking technologies; hence, they could not be more different from, and more incompatible with, each other. The field of multimedia networking is split into two camps, each favoring its own approach. This would not, in itself, be a deplorable situation if the winning approach in the competition were to be adopted by everybody. Unfortunately, at least in the foreseeable future, it is unlikely that such a happy event will occur; a more likely result of the current distance between the two approaches is that the two types of networks will uneasily coexist, and will only communicate in difficult, limited, *ad-hoc* ways, to the detriment of the clients of both. For the sake of these clients, it would be much better if all possible efforts were made *now* to design the architectures and the protocols of both network types so as to maximize the compatibility and interoperability of the services they will offer.

The Tenet approach has characteristics that place it somewhere between those of the Internet and the ATM communities. It may therefore be useful to reflect, at this critical juncture, on the Tenet experience. I shall attempt to do so in this paper.

I assume that the reader is acquainted with the principles of multimedia networking and of the main approaches that have been proposed. The service models pertaining to these approaches are reviewed and compared in [5], where an essential bibliography regarding the approaches being pursued within the Internet and ATM communities can also be found. For the Tenet approach, the principles are discussed in [6], and Scheme 0 in [7]; the extensions to Scheme 0 that resulted in Scheme 1 are described in [4], the design and implementation of Suite 1 in [1], and some of the experiments performed on it in [2]. Scheme 2 and Suite 2 are the subjects of [9], [3], and [8]. Those referenced in this paragraph are the essential Tenet papers. They can be found, along with many others concerning various details of the Tenet schemes and suites, on the World Wide Web at URL http://tenet.berkeley.edu.

2. The service model

The context for which the Tenet real-time protocols have been designed is that of a large, heterogeneous, wide-area internetwork. The **wide-area** orientation was not meant to exclude, but to encompass, regional-, metropolitan-, and local-area networks. Indeed, all or most of these types of networks are usually traversed by a wide-area connection in an internetworking environment. It is conceivable that protocols implemented for wide-area networks may not be optimal for local-area ones, and vice versa; however, we believe that the structure and most of the design of such protocols can be done without having to tailor them to a specific geographical extent, while only a few aspects of the design and perhaps many more of the implementation can and should be influenced by it.

Heterogeneity is not yet a major concern of the ATM community. The pan-ATM vision of the world that many members of that community seem to hold is at best very futuristic, hence not very realistic. Even if ATM technology were to be wildly successful, we would still have to live with major non-ATM parts of the global infrastructure for a long time, and we cannot postpone the advent of universal integrated-services until the last non-ATM network will have been decommissioned. Thus, we must design architectures and protocols for heterogeneous internetworks even if we believe in the eventual elimination of all non-ATM networks (a belief, incidentally, that I find very hard to subscribe to). Another important aspect of heterogeneity is found among switches, routers, and gateways. All major approaches in the Internet and ATM camps allow for interoperable network nodes, as does the Tenet approach, but some of the other schemes being proposed in the literature impose restrictions that are likely to be excessive, usually without discussing these restrictions, their impact, and the possibility of removing them. Yet another dimension of design for heterogeneity is the tolerance for parts of an internetwork in which certain services (e.g., those offering performance guarantees) may be unavailable. This tolerance is essential, at least during the period when new services are being added to an internetwork, an operation that must be done gradually. In the Tenet approach, no guarantees can be offered on a path that includes nodes that neither run the real-time channel setup protocol (RCAP) and the network-layer protocol (RTIP) nor belong to a subnetwork that offers performance guarantees. However, communication can still take place on such a path, with generally better performance than if non-real-time protocols were used.

The internetworks we consider may be very **large.** This is, I think, a non-controversial issue. Although a careful study of multimedia protocols might reveal that protocols designed for large networks are not optimum in smaller environments, scalable designs and implementations would still be necessary. Scalability is not only very important but also very hard to achieve; it has several dimensions (geographic extent, number of hosts, number of hosts or users participating in an application, number of application types and applications, bandwidth, topological complexity), all of which are important. The Internet community is very interested in scalability; we in the Tenet Group have considered it one of the main design objectives in Suite 2, but only experimentation will determine how well we have succeeded in each dimension.

The service model of an integrated-services network incorporates some of the most fundamental decisions to be made by the network's designers. These decisions determine which of the many different services that might be offered should indeed be offered by the network. In this area, I believe we must proceed cautiously and conservatively, given that so little is known about the requirements of real multimedia applications, and about the behavior of the multimedia architectures and protocols with which we have just started to experiment. It makes sense to begin by offering a few essential services, and to adopt and supplement those that succeed with the more sophisticated services real applications will undoubtedly suggest. To facilitate adjustments, additions, and replacements, the first service models ought to be as flexible and free of constraints as possible. For example, they should permit experimental evaluation and minimize dependencies among parameters; instead of forcing real-time clients (of the transport layer) to choose one of several *quality of service (QoS) classes,* they should let them freely specify the values of the QoS parameters, i.e., of the bounds they need the network to guarantee. Leaving this freedom to transport-layer clients does not preclude offering QoS classes to application-layer clients; such classes can be tailored to fit the needs of real applications, after the experiments involving these applications have been performed, and can be more easily extended and modified when they are built on top of a general parameterized transport-level service.

What services are really *essential*? The most fundamental classification of the known traffic types is that which distinguishes *real-time traffic* from *non-real-time traffic,* where the former class can be seen as encompassing all types of traffic that require some form of performance bound, and the latter class includes traffic that may require guarantees on reliability of transmission but not on performance. The simplest types of service corresponding to these fundamental classes are those of *deterministic guaranteed services* and *best-effort services,* respectively. The former provides deterministic, i.e., absolute, bounds on each of the chosen QoS parameters; the latter try to transfer packets expeditiously and fairly, and to control congestion, but cannot enforce any bounds on performance. All the major approaches proposed so far include these two types of service; however, they also add some types of service that can be regarded as an intermediate one between the two extremes. This is the case of *predictive service,* which has been devised for an Internet-like environment, and, while not providing mathematically provable guarantees, is expected to offer equivalent performance bounds most of the time; it is also the case of Tenet's *statistical service,* which provides mathematically provable statistical bounds, less expensive for the network to support than the corresponding deterministic bounds. In the current ATM/AAL specifications, available bit-rate (ABR) service can be regarded as an intermediate service between guaranteed and best-effort, which is represented by unspecified bit-rate (UBR) service. ABR service does not provide guarantees, except that for a small amount of bandwidth, but is carefully flow-controlled so as to minimize cell losses. The distinction between the two classes of guaranteed traffic, constant bit-rate (CBR) and variable bit-rate (VBR), is of a different nature, and does not seem as fundamental; perhaps, it would be more reasonable to let this distinction be expressed by the values of the traffic parameters than to represent it as two different types of service.

One might wonder why we need to offer guaranteed-performance services to real-time traffic, since guarantees require admission control, which may cause access to the network to be denied to real-time clients when the network's real-time capacity is saturated, and access denials are clearly undesirable. The reason is that the alternative solutions that are usually proposed are either not effective against overloads (priority scheduling without admission control, adaptive applications, better congestion control algorithms) or likely to be too expensive and wasteful (overprovisioning). If quality of service is protected by admission control, there are techniques that can be used to increase the real-time capacity of the network when necessary: for example, *media scaling,* which consists of offering discounts to clients who will accept to reduce their input rates within a prespecified range when requested by the network; or *dynamic bandwidth renegotiation* during the transmission of multimedia streams. Incentives for the use of intermediate services should also result in substantial increases of the available capacity for real-time traffic. One might also argue that deterministic services are not really needed, if what I have called intermediate services are being offered by the network; unfortunately, however, there are applications (e.g., telesurgery, telecontrol, robotics, studio-quality video, and so on) that are intolerant of disruptions or even deteriorations of quality; unless we exclude these applications from our networks, a decision in clear contradiction with the notion of integrated services. Hence, we must offer deterministic guarantees.

Is it really necessary to offer guaranteed-performance services at the network layer? Could we build a transport-level real-time service on top of a best-effort network protocol, or an application-level service on top of a best-effort protocol stack? Regrettably, the answer to both parts of the latter question is negative. Unlike the case of transmission reliability, which can be obtained by layering a reliable protocol on top of an unreliable one, performance guarantees can be offered to users only if all layers support, in their own way, those guarantees.[1]

One might then wonder whether an integrated-services network could avoid offering non-real-time services. In principle one could, but it would result in a substantial reduction of the total capacity of the network. Thus, the two types of service I have called essential are indeed essential. But integrated-services networks will eventually offer a much richer menu of services. For instance, the traffic we consider non-real-time, because it has traditionally been using the Internet, actually contains subtypes with different performance requirements or preferences: interactive applications such as telnet are very close to real-time traffic (in fact, I believe that telnet clients are likely to use real-time services in congested situations when these services will be available), where non-time-critical applications such as electronic mail will not (except in the case of urgent messages). Other applications, such as file transfer, may benefit from the enforcement of a lower bound on throughput. Some of these traffic subtypes will probably justify the introduction of several flavors of non-real-time services and of intermediate services.

Another way to make offerings richer would be to increase (within reason) the number of QoS parameters. What are the best parameters? What is their optimum number? No definitive answers can be given to these questions. One can only look at the proposals presented to date, and determine the most popular parameters contained therein and their numbers. With the exception of ST-II, whose larger number of parameters was due to the desire to experiment with several subsets of them, the typical number of QoS parameters in these first attempts is between 3 and 6, and the main performance indices they represent are throughput, delay, delay variation, and packet or cell loss. There are remarkable similarities in the types of QoS parameters used by the different approaches, though there are also non-negligible differences in their definitions.

Integrated-services networks will have to provide their clients with other types of services besides those I have discussed above: for example, there will be video conferencing tools, audio conferencing tools, and other kinds of cooperative work support services, such as conference control protocols; advance booking of real-time services, made necessary for multi-party application planning by the possibility of access denials; security services of various kinds; and virtual private network creation and modification services, which will be usable by network managers for the direct convenience of some corporate users. Some of these services must or may be implemented at the network or transport layers, on which this paper is focusing; the others are to be provided by higher layers, but need to be supported by network and/or transport layer services nevertheless.

3. The charging problem

Thus far, very little work has been reported in the open literature on the relative **costs** of the various services that may be offered by an integrated-services network. These costs, if known, would provide a technical basis for the charging policies to be adopted for such a network.

The Tenet Group cannot claim to be an exception: some work was done on the cost problem, but it was not conclusive, and the topic was therefore (perhaps only temporarily) abandoned. I do not have any results to point to, but I feel the need to make some observations on the charging problem anyway. These observations have been suggested by the thoughts about the subject prompted by the Tenet experience, and are therefore, in a sense, a result of the Tenet project.

It is impossible for me at this time to assign relative costs to the various services the Tenet protocols (or any other real-time protocols) offer. Nor can I make plausible arguments in favor of, or against, any specific charging policies one may wish to propose for an integrated-services network. However, I feel strongly that charges must be *service-dependent,* and that there must be some positive correlation between the price of a service and the cost of providing it. I do not regard these statements as breakthroughs, and I certainly do not claim to be the first to make them. On the contrary, they are obvious observations, which

[1] We can trade reliability for delay, but not for a delay bound.

need to be emphasized only because they do not seem to be universally held beliefs yet.

The point is sometimes made that network usage should be free; others argue that charging for access should not obey policies more complicated than flat-rate ones (e.g., a monthly fee); the latter position is often justified by stating that any more sophisticated policy would be too difficult and/or expensive to implement.

The partisans of free access and free usage must surely expect that the cost of building and running the networks will be supported by charitable institutions, or more plausibly by governments, i.e., by all taxpayers. This model is the one most countries have adopted for their highway systems, though tolls are often exacted for the use of *better* highways (i.e., those with no crossings), of bridges, and so on. It is not, however, the model used by most utilities in most countries, which apply in their charging policies the arguably fairer principles of usage sensitivity and service dependency. In any case, it is important to recognize that, *even if there were no other reasons to do so, service-dependent charging would have to be used in integrated-services networks.* Without service dependency, since different services usually have different costs, the network's clients would have no major incentives to be frugal, i.e., to choose the minimum-cost service that will satisfy their requirements; such behavior would appreciably reduce the effective capacity (and the revenue) of the network. If real-time and non-real-time services had the same prices, who would ever use the non-real-time ones, unless a real-time request were rejected?[2]

Note that I have used the term "service", not "service type", in the previous paragraphs. I am not only advocating that, for example, deterministic services have higher prices than their statistical counterparts, but also that a deterministic channel with an end-to-end delay bound of 250ms have a higher price than one with a delay bound of 400ms, as long as supporting the former costs more to the network than supporting the latter. The addition of usage sensitivity will make the charging policy more fair; this is true not only for the case of non-real-time traffic, in which no or very little resources are tied up for the duration of the transmission of a message, but also for real-time-traffic, where lower usage by a real-time stream allows more room for non-real-time packets to be transmitted. However, *usage-sensitive charging is only preferable,* while *service-dependent charging is essential.* Service dependencies in accounting should not be difficult to handle; but, even if they were, a major effort should be made to resolve these difficulties: if the objective of service-dependent charging is not achieved, there is no point in building integrated-services networks.

4. The architecture

Since the outset of my research on real-time protocols, I have been more interested in integrating services within an internetwork than in solving the problem of providing performance guarantees to real-time traffic, though, for obvious reasons, I was completely absorbed by the latter problem during the first two years of the research. When the Tenet Suite 1 was designed in 1990-91, one of the design objectives was that the protocols should be able to coexist with those in the Internet suite, so that a user with non-real-time applications developed for an Internet-like environment could continue using TCP/IP or UDP/IP without making any changes to those applications. Since that time, the Tenet Group has presented this architecture, which is based on two coexisting protocol stacks, the Internet one and the Tenet one, as its proposal for a simple integrated-services network architecture. The classification of applications into real-time and non-real-time ones is considered by many as the most fundamental among the several that could be introduced. If the services offered by the Internet stack are acceptable to non-real-time users, these users will be able to satisfy their communications requirements by using TCP or UDP without being adversely impacted by the presence of the real-time stack, except perhaps in their performance, due to the higher priority of real-time traffic. For real-time applications, if the services offered by the Tenet protocols include those needed by the user or acceptable replacements for them, then a real-time user will request the creation of one or more *channels* by invoking the RCAP channel setup protocols; if RCAP responds affirmatively to the request, the user can start sending traffic at the appropriate rate on the channel(s) just created using the data delivery protocols RMTP and RTIP.

---

[2] Also, it is interesting to recall from economics that a market-based system, in which prices are set by consumer demands and supplier costs, most *efficiently* and *fairly* allocates resources for a commodity, which is exactly what bandwidth will be if it is going to be as plentiful as "free-access partisans" expect.

Does ATM fit into this picture? In my opinion, it can be made to fit if on both the Internet and the ATM sides there is the will to converge to the same general architecture for integrated-services internetworks. An example of such an architecture is shown in Figure 1. The example uses the Tenet protocols, but of course they could be replaced by others: all the protocols referenced in the figure are indeed functional placeholders, as they stand for the services they offer or for services of the same types. Thus, for instance, a future version of IP might also offer real-time services similar to those now obtainable from RTIP, and thus render the presence of RTIP in the architecture superfluous. IP in the figure represents the functions that are needed to transport non-real-time data across a (possibly heterogeneous) internetwork, for example those embodied in the current version of IP, IPv4. The protocol is now evolving, though still very slowly, towards one capable of supporting also the types of real-time services that will be needed in future networks. Until the evolution in this direction is completed, the solution shown in Figure 1, which has two protocols at the network layer, may be the only viable one.

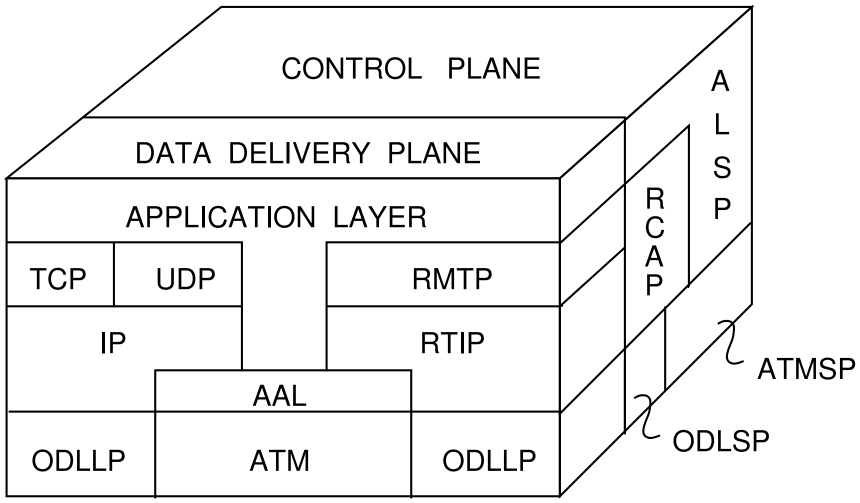

Figure 1. An architecture for an integrated-services internetwork. ODLLP: Other Data Link Layer Protocols. ALSP: Application Layer Signaling Protocols. ODLSP: Other Data Link Signaling Protocols. ATMSP: ATM Signaling Protocol.

ATM, whenever endowed with a signaling protocol capable of exploiting its rich integrated-services potential, will offer a similar spectrum of services directly to applications running on each ATM network. When, however, an application will involve computers connected to non-ATM networks, some other protocols layered on top of ATM will have to be invoked, as shown in the figure. Such steps will also be necessary even within an ATM network if a client needs a service (e.g., in the field of security) that is not offered by the native ATM stack. Each use of the ATM protocol will obviously include that of an appropriate adaptation layer (AAL), as shown in the diagram.

In the control plane, the block labeled ALSP translates the requirements of a real-time application into the form needed by the ATM signaling protocol or by a transport-layer connection establishment protocol like RCAP, depending on whether the request is for a connection inside a single ATM network or for one involving additional networks. RCAP, in turn, is to be interfaced to the ATM signaling protocol and to other datalink-level signaling protocols collectively labeled ODLSP. For non-real-time services, the architecture in Figure 1 assumes that all or most of these users will prefer the connectionless datagram model when accessing the internetwork in non-native-ATM mode. Of course, any crossing of an ATM network by datagrams requires the use of a virtual connection, but the establishment and management of this connection will be invisible to the higher layers. The addition of other protocols is obviously possible; the architecture in Figure 1 is only an example of how existing protocols can be combined to offer a variety of services encompassing most of those currently believed to be useful in a heterogeneous information infrastructure. Unfortunately, not all the pieces are currently available, the most notable absence being that of an adequate, standardized ATMSP, but they could be there in a relatively short time if there were a clear agreement among the various constituencies to make this happen.

5. Concluding remarks

In this paper, I have tried to interpret the significance of the work done by the Tenet Group in the current debate about how an integrated-services information infrastructure should be architected. In particular, I have discussed the service model, the charging model, and the architectural aspects of the problem, and shown in what ways the Tenet experience could help attack each one of them. With all its limitations (just to quote a major omission as an example, it has not considered the needs of wireless mobile hosts and users), the Tenet Group's research in real-time networking has demonstrated that it is possible to obtain an integrated-services internetwork that preserves the successful Internet model for non-real-time traffic while providing all real-time applications with the performance guarantees they need even over routes not confined to be within a single ATM network. In light of the current debate, these results take on crucial importance. Even though it has not yet proved the feasibility of this convergence by actual construction, as the ATM part is still missing, I believe the Tenet approach has shown that such convergence is possible and highly desirable.

The main strength of the approach is what some fellow researchers might regard as a weakness: its uncompromising rigor in those aspects that we felt would require such a rigor. We insisted that the guarantees be mathematically provable, an objective that some mistakenly interpreted as equivalent to the exclusion of non-deterministic bounds. In order to make admission control practically feasible, we compromised on the traffic model, by choosing a particularly simple and coarse one; we could do so because we had concluded that, unlike provability, the choice of the traffic model is not a fundamental one, and we felt (correctly, it turned out) that such an initial model could be refined later. We insisted that the relationship between the client and the service must be regarded as a contractual one, and this decision has actually provided us with a powerful tool for reasoning about that relationship as well as about the design of the corresponding interface. We wanted connection establishment to be fast, and this goal called not only for efficient admission tests, but also for single-round-trip setup and for the determination of bounds local to each network node. These decisions entailed making compromises, but they were adopted because these are all areas where compromise does not affect the principles which we wanted to be rigorously obeyed. Localizing bounds also allowed us to decouple each node from the others traversed by the same connection, facilitating scalable implementations and the interoperability of routers not only with different admission control algorithms, but also with different service disciplines.

The rigor of the approach is probably responsible for the ease with which, once a new problem was correctly formulated and constrained, a solution emerged almost naturally and effortlessly. This experience occurred in several instances: for example, extending the original Tenet scheme to an internetworking environment; devising the resource partitioning and advance reservation mechanisms; designing dynamic channel management algorithms; and developing a mechanism for setting up a connection starting from the receiver or for establishing a duplex connection. Of course, I cannot predict that the protocols that will implement real-time services in practical networks will necessarily be based on such a rigorous approach, though I believe there are applications that truly require mathematically provable guarantees. However, even if the actual protocols end up following a less rigorous approach, it is always going to be useful to know how far they are from their rigorous counterparts. If necessary, it is better to introduce approximations into a clean, well-understood, and easy-to analyze "exact" solution than to start with a heuristic, complex approach with unknown properties. For example, it is easier to understand, and reason about, an "intermediate" service derived by approximation from deterministic guaranteed service than an "intermediate" service created in isolation without a rigorous point of reference.

Perhaps the most useful result that the Tenet effort has produced can be summarized in the following message: *forcing the traditional Internet model to offer the wide spectrum of services an integrated-services internetwork will have to offer, which include deterministic guarantees and "intermediate" types, is neither wise nor necessary.* It is not wise because the Internet model, based on datagrams, variable routes, no admission control, and flat-rate charges, is totally inadequate for several important types of traffic the new networks will have to carry; and it is not necessary because the Internet protocols can easily be made to coexist within the same network with real-time protocols that can satisfy the quality requirements of the real-time portion of the traffic. Even macroscopic modifications of the Internet model are not likely to do the job as long as they preserve the fundamental characteristics of the Internet philosophy. This message is implicit in the ST-II specification, and is not diminished by the shortcomings of that early protocol; the fact that the Tenet Group has independently reached the same conclusions having started with the same problem, in spite of all the differences between the two approaches, may be seen as a stronger indication that the message deserves to be seriously considered.

Acknowledgements

I am indebted to all the researchers in multimedia networking, and particularly to the past and present members of the Tenet Group, for stimulating many of the thoughts reported in this paper, and sometimes even contributing directly to their formulation. The many improvements to the text of the first draft suggested by Mark Moran are gratefully acknowledged. Unfortunately, I cannot put the blame for any errors or misjudgements on anybody but myself.

The support, throughout the Tenet years, of the National Science Foundation, the Advanced Research Projects Agency, the Department of Energy, the Corporation for National Research Initiatives, the University of California, the International Computer Science Institute, AT&T Bell Laboratories, Digital Equipment Corporation, Hitachi, Hitachi America, Mitsubishi Electric Research Laboratories, and Pacific Bell is gratefully acknowledged. All opinions in this paper are those of the author, and do not represent opinions or policies, either expressed or implied, of the U.S. Government or of any of the other sponsoring organizations.

References

[1] 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.

[2] 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.

[3] 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," *Proc. Fifth International Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV),* Durham, NH, pp. 255-266, April 1995.

[4] D. Ferrari, "Real-Time Communication in an Internetwork," *Journal of High Speed Networks,* vol. 1, n. 1, pp. 79-103, 1992.

[5] D. Ferrari, "Multimedia Network Protocols: Where Are We?," *Proc. IST/SPIE's Symp. on Electronic Imaging: Science and Technology, Technical Conf. on Multimedia Computing and Networking 1995,* San Jose, CA, February 1995.

[6] 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, July 1994.

[7] D. Ferrari and D. Verma, "A Scheme for Real-Time Channel Establishment in Wide-Area Networks," *IEEE J. Selected Areas in Communications,* vol. 8, pp. 368-379, April 1990.

[8] 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.

[9] A. Gupta, W. Heffner, M. Moran, and C. Szyperski, "Network Support for Real-Time Multi-Party Applications," *Abstracts of the Fourth International Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV),* Lancaster, England, pp. 37-39, November 1993.