

A Comment on Packet Video Remote Conferencing and the
Transport/Network Layers

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard. Distribution of this memo is unlimited.

Abstract

The new generation of multimedia applications demands new features and new mechanisms for proper performance. ATM technology has moved from concept to reality, delivering very high bandwidths and new capabilities to the data link layer user. In an effort to anticipate the high bandwidth-delay data link layer, Delta-t [Delta-t], NETBLT [RFC 988], and VMTP [RFC 1045] were developed. The excellent insights and mechanisms pioneered by the creators of these experimental Internet protocols were used in the design of Xpress Transfer Protocol (XTP) [XTP92] with the goal of eventually delivering ATM bandwidths to a user process. This RFC is a vehicle to inform the Internet community about XTP as it benefits from past Internet activity and targets general-purpose applications and multimedia applications with the emerging ATM networks in mind.

1. Introduction

Networking is no longer synonymous with analog telephony. High-performance lower-layer networks have made possible exciting new applications: collaborative environments, distributed client/server computing, remote conferencing, teleclassrooms, and distributed life-sciences imaging. These applications normally demand a great deal of bandwidth and often create operating system bottlenecks. Enabling these new multimedia applications entails delivering bandwidth to the applications, not just having bandwidth available on the network. This statement may appear obvious, but often solutions at the transport layer are satisfied by having bandwidth at that layer without sufficient sensitivity to higher-layer access to the bandwidth. The unavailability of bandwidth at upper layers is becoming the real issue as the networks are becoming a high-performance virtual backplane without concomitant high-performance control schemes. It appears that new services are needed that require communication with all layers. The ATM architecture calls

for such an integrated control scheme.

2. Remote Conferencing

The challenges of remote conferencing is an application whose challenges may be met at the data link layer by the emerging broadband networks. If so, important medical applications such as medical imaging for diagnosis and treatment planning would be possible [CHIM92]. Remote conferencing would permit imaging applications for life sciences through the use of national resource centers. Collaboratory conferences in molecular modeling, design efforts, and visualization of data in numerous disciplines could become possible.

At the Second Packet Video Workshop, held December, 1992, at MCNC in the Research Triangle Park, North Carolina, a recurrent theme was the use of multimedia in remote conferencing. Its applications included the use of interactive, synchronized voice and video transmission, multicast transmission, data transfer, graphics transmission, noninteractive video and audio transmission, and data base query within a virtually shared workspace. A few participants doubted the ability of current computer networks to handle these multimedia applications and preferred only connection-oriented, circuit-switched services. Most participants, however, looked forward to using an integrated network approach.

2.1. Remote Conferencing Functions and Requirements

Remote conferencing as seen at the workshop requires a set of functions. It must provide session scheduling that deals with initiating a session, joining in-progress sessions, leaving a session without tearing it down if there are multiple participants, and terminating a session.

The remote-conferencing session needs a control subsystem that is either tightly controlled for an n-to-n connection for two to 15 participants, or loosely controlled for a 1-to-n connection for any number of participants. The Multipeer-Multicast Consortium is working on defining the control requirements and the mechanisms for control. At the Packet Video Workshop, one participant presented a conference control protocol (CCP) shown in Figure 1 [CCP92]. In this architecture the CCP controls the Network Voice Protocol (NVP) [RFC741] and the Packet Video Protocol (PVP) [PVP81] over the experimental Internet Stream Protocol, Version 2 (ST-II) [RFC1190] rather than IP.

Latency and intramedia synchronization and intermedia synchronization (lip-sync) are critical for the interactive voice and video streams

of remote conferencing. Client/server applications including data base operations are equally important. The ability to display noninteractive video and high-resolution graphics is necessary.

As with most applications, security will eventually be an issue. At the very least, there must be a mechanism to determine who can find out what about conference and who can join a conference. This determination will be part of an upper-layer protocol.

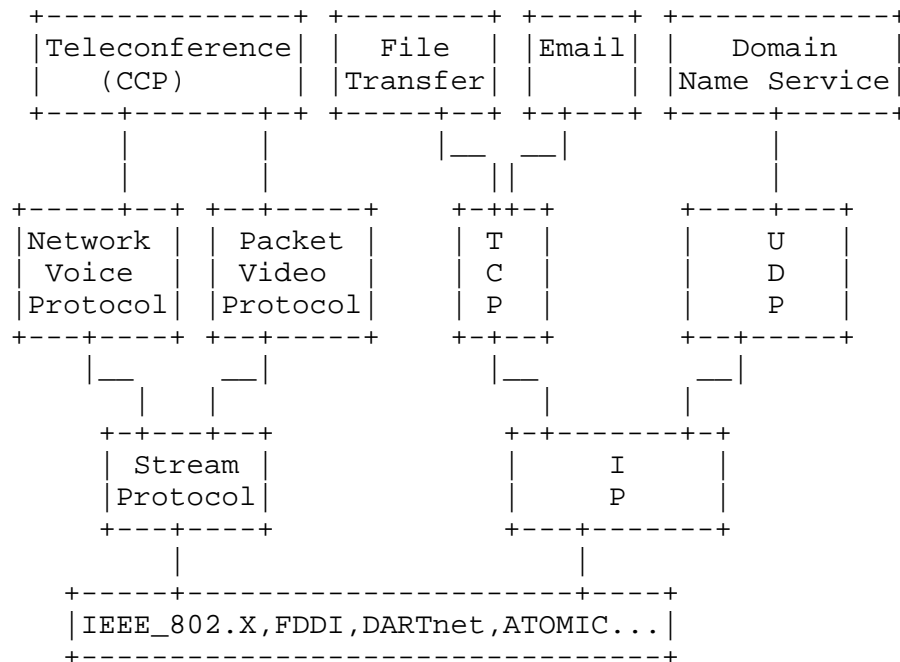


Figure 1: The Connection Control Protocol Architecture

The solutions must range in geography from single machines through LAN, CAN, MAN, WAN conferences, as well as in data content from PCs to high-end workstations. Implicit in the scaling is the notion of product and application interoperability.

Remote conferencing applications should take advantage of future network enhancements, as well as the functions provided by ATM, FDDI, and SMDS. Doing so should provide function versus resource trade-offs such as cost versus error control and cost versus rate control. As a result, the transport layer should be able to provide the services offered by the data link layer.

Most of the presentation on remote conferencing indicated a need for some degree of multicast functionality, ranging from the 1-to-n, with conference membership completely known, to conferences for which

existence of a group is known, but exact membership is not.

In remote conferencing, it is preferable to use one network for all information traffic. This network should have an offered quality of service (QOS) criteria to accommodate tradeoff metrics, which would include guaranteed throughput, connection reliability, high call-completion rate, few dropped calls, tolerable error rate, varying degrees of compression on the video and audio streams, and tolerable motion artifacts, flow control, and latency. The QOS should be an input to the transport layer from the application or transport service user.

The compression/coding function should provide time-stamping and packetizing information, as well as real-time echo cancellation. These functions are usually at the presentation and session layer of the Open System Interconnection (OSI) model or the at the application in some Internet models, but not the transport layer.

3. Potential Solutions

RFC 1193 deals with the requirements of real-time communications, which include some of the same capabilities [RFC1193]. But the requirements articulated create the necessity for new transport/network protocols. The new protocols under development by the Audio Visual Transport [SCHU92] (RTC, RTCP), and other protocols in this area (ST-II) suggest an architecture like that shown in Figure 2.

These approaches may work. However, they encourage a discipline that creates a protocol for each new class of application. Another approach might be to unify the protocols. It is felt that this is one of the main goals of XTP (see Figure 3).

Other design considerations of XTP include the following:

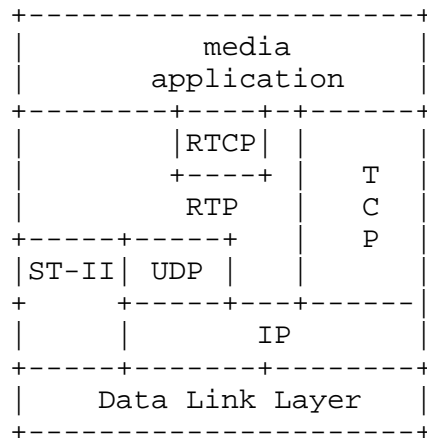


Figure 2: One emerging multimedia architecture

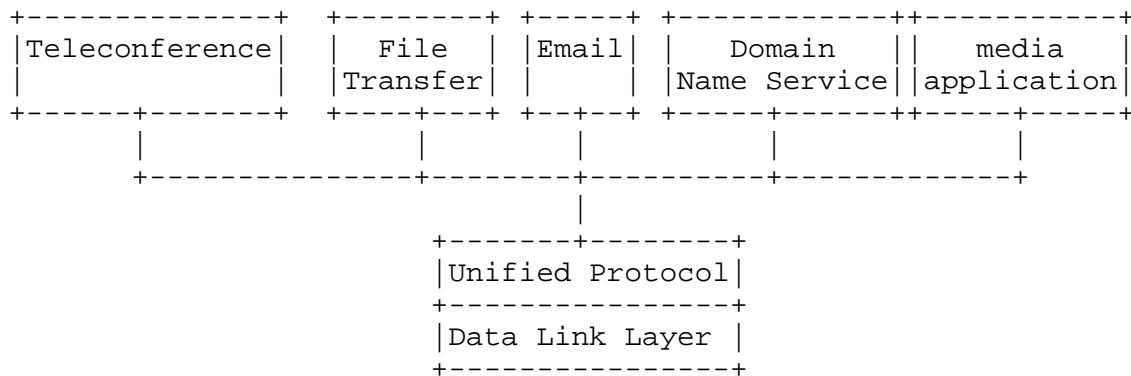


Figure 3: Another integrated multimedia architecture

- (1) Developing a protocol based on the work and experience of the protocol groups such as IETF, which produced NETBLT, VMTP, TCP, IP, and UDP.
- (2) Incorporating lessons from Delta-t.
- (3) Observing the design paradigm shift of ATM, SONET, and VMTP in the header, trailer, and information segment construction.
- (4) Targeting the real-time requirements articulated by the Department of Defense SAFENET committee and the French Ministry of Defense military real-time specification [GAM-T-103].

Mechanisms in XTP allow an application to approach the performance desired. A session-scheduling mechanism for joining and leaving a

multicast conference exists. The XTP mechanism for multicast satisfies the loosely controlled session requirements. The tightly controlled session would require the use of multiple XTP associations.

The priority mechanism that uses the 32-bit SORT field permits an application to prioritize data. Because XTP is a transport layer, this priority mechanism follows through every node tranversed. There is also an out-of-band delivery mechanism. However, XTP does not offer latency control by itself; it just provides a priority mechanism.

The selective acknowledgement, fast negative acknowledgement, and selective retransmission permit an application to choose an appropriate level of error control. The combination of the priority mechanism and these error-control mechanisms is likely to approach the latency and synchronization requirements of remote conferencing.

Noninteractive audio and video, as well as graphics presentation, can be accommodated in many ways by the application. It is important that the transport layer provides adequate mechanisms to deliver the appropriate data streams in a manner compatible with the applications. These applications can probably be accomplished by means of extant protocols, as well as XTP.

The scalability of the solution will be a function of the standards used. At the Packet Video Workshop, some of the applications sacrificed computer network standards in favor of desired performance. This approach usually impedes scalability. From the explanation of the applications taking this approach, it appeared that using XTP would have provided an adequate transport service for the applications.

XTP was designed to provide mechanisms to accommodate application requirements, that is, the ability to respond to QOS requests. For example, guaranteed throughput may be accomplished by using XTP's rate and burst control together with flow control or no flow control. Tolerable error rate can be accomplished with partially error controlled connections (PECC), a service which can be placed in the application or just above the transport layer [PECC92]. Motion artifacts and varying degrees of compression should be done above the transport layer in coordination with the transport layer or possibly in a network management function.

3.1. Synthesize the Hardware Fabrication Process into the Design

To produce an affordable solution, the hardware fabrication process should be a design consideration. Technologies are evolving too

rapidly to assume that a generic protocol design will anticipate all fabrication advances, but this fact should not impede use of the features of advanced hardware fabrication processes.

System interface problems and VLSI techniques should be considered in the specification of the protocol. An examination of the ATM and SONET standards appears to support this philosophy. Similarly, NETBLT and VMTP design efforts seem to support this approach. XTP does use it.

It is very helpful to break down the protocol into parallel-state machines for execution on more inexpensive hardware. This procedure reduces the context switching and interrupt handling requirements of the hardware, thereby decreasing production costs while producing a scalable protocol machine.

4. Multimedia Applications over XTP

In parallel with the IETF efforts to enable multimedia applications such as remote conferencing, the XTP forum members have been experimenting with major elements of these applications.

- (1) At the University of Virginia, more than 100 simulated voice channels were run on an FDDI network [UVAVOICE92]. The purpose of this experiment was to test the limits of FDDI and a software version of XTP in a simulated interactive voice environment. Multicast, noninteractive video has been supported there for several years.

UVA also has a video-mail demo over XTP/FDDI that uses Fluent multimedia interface and standard JPEG compression. This PC-based demo delivers full frame, full color, 30 frames/sec video from any network disk to a remote VGA screen. It is important that users could not discern any difference in playback between a local disk and a remote disk. An Xpress File System (XFS) is used on server and client systems.

- (2) The Technical University of Berlin, Germany, reports that the coordination, implementation, and operation of multimedia services (CIO) of the R&D in Advanced Communications Technologies in Europe (RACE) is using XTP as a starting point for design [XTPRACE].
- (3) At the Naval Command, Control, and Ocean Surveillance Center Research, Development, Test and Evaluation Division NRaD (formerly the Naval Ocean Systems Command (NOSC)), voice is

multicast over XTP/FDDI. A simple multicast is distributed to a group with a latency of around 25 ms, where the latency represents delay from the voice signal from the microphone to the audio signal to the speaker. This group is currently doing research on an n-party multicast of voice (telephone conference-call paradigm [n x n multicast]).

- (4) Commercially, Starlight Networks Inc., migrated a subset of XTP into the transport layer of its video application server. By using XTP rate control, full-motion, full-screen compressed video is delivered at a constant 1.2 Mbps, over switched-hub Ethernet to viewstations. This network delivers at least 10 simultaneous video streams.

Therefore, XTP has been used in applications that were previously placed over IP or even a data link layer.

5. Policy versus Mechanism

Separating protocol policies and mechanisms [XTPbk] permits adoption of new policies without altering offered mechanisms. An excellent example is UVA's Partially Error Controlled Connections (PECC). This control system maximizes error control in such a way that receiving FIFOs are never starved i.e., the application, driver, or operating system buffer control, and not the transport layer becomes the bottleneck.

Because XTP is mechanism-rich and policy-tolerant, this very dynamic error control policy mechanism is possible. Separating policy and mechanism permits an error-control or flow-control policy to adapt to the data link layer conditions without shutting down a connection and rebuilding (or multiplexing) a new one on a different protocol stack. This may also provide an easier way for a network management subsystem to maintain a desired QOS.

6. Summary

Remote conferencing presents new opportunities for research, business, and administration. Although some are proposing that only classical circuit-switched mechanisms be used, most network engineers are searching for ways to use the new features of FDDI, SMDS, and ATM in multimedia applications such as remote conferencing. Some new applications have been placed directly on a data link layer. New Transport/Network layer combinations have been proposed and are being tested. It is believed that consideration should be given to XTP as a possible solution because its forum membership includes commercial, government, and research institutions, some of which have implemented various applications that contribute to an overall remote-

conferencing application.

7. References

- [CCP92] Schooler, E., "An Architecture for Multimedia Connection Management", in Proceedings of the 4th IEEE ComSoc International Workshop on Multimedia Communications, Monterey, CA, April 1992.
- [CHIM92] Chimiak, W., "The Digital Radiology Environment", IEEE JSAC, Vol. 10, No. 7, pp. 1133-1144, September 1992.
- [Delta-t] Watson, R. W., "Delta-t Protocols Specification", Lawrence Livermore Laboratory, April 15, 1983.
- [GAM-T-103] French Ministry of Defense, "GAM-T-103 Military Real-Time Local Area Network Reference Model (Transfer Layer)", February 7, 1987.
- [PECC92] Dempsey, B., Strayer, T. and Weaver A., "Adaptive Error Control for Multimedia Data Transfer", in Proceedings of the IWACA 92, Munich, Germany, pp. 279-288, March 1992.
- [PVP81] Cole, R., "PVP - A Packet Video Protocol", W-Note 28, Information Sciences institute, University of California, Los Angeles, CA, August 1981.
- [RFC1045] Cheriton, D., "VMTP: Versatile Message Transaction Protocol Specification", RFC 1045, Stanford University, February 1988.
- [RFC998] Clark, D., Lambert, M., and L. Zhang, "NETBLT: A Bulk Data Transfer Protocol", RFC 998, MIT, March 1987.
- [RFC1193] Ferrari, D., "Client Requirements For Real-Time Communication Services", RFC 1193, UC Berkeley, November 1990.
- [RFC1190] Topolcic, C., Editor, "Experimental Internet Stream Protocol: Version 2 (ST-II)", RFC 1190, CIP Working Group, October 1990.
- [SCHU92] Schulzrinne, H., "A Transport Protocol for Audio and Video Conferences and other Multiparticipant Real-Time Applications", Internet Engineering Task Force, Internet-Draft, October 1992.

- [UVAVOICE92] Weaver, A. C. and McNabb, J.F., "Digitized Voice Distribution Using XTP and FDDI", Transfer, Vol. 5, No. 6, pp. 2-7 (November/December 1992).
- [XTP92] Xpress Transfer Protocol, version 3.6, XTP Forum, 1900 State Street, Suite D, Santa Barbara, California 93101 USA, January 11, 1992.
- [XTPbk] Strayer, W.T., Dempsey, B.J., and Weaver, A.C., "XTP: The Xpress Transfer Protocol", Addison-Wesley Publishing Company, Inc., 1992.
- [XTPRACE] Rebensburg, K. and Miloucheva, I., "The Use of XTP in a Large European Communication Project", XTP Forum Research Affiliate Annual Report, Document 92-183, pp. 105-112, 1992.

Security Considerations

Security issues are discussed in section 2.1.

Author's Address

William J. Chimiak
Department of Radiology
Bowman Gray School of Medicine
Medical Center Boulevard
Winston-Salem, NC 27157-1022

Phone: 919-716-2815
EMail: chim@relito.medeng.wfu.edu