

*ISI Reprint Series*  
*ISI/RS-92-294*  
*August 1992*

Eve M. Schooler  
Stephen L. Casner



An Architecture for  
Multimedia Connection Management

Reprinted from the  
*Proceedings IEEE 4th Comsoc International*  
*Workshop on Multimedia Communications,*  
MM '92, pp.271-274, Monterey, CA (April 1992).

*University of Southern California*  
*Information Sciences Institute*  
4676 Admiralty Way, Marina del Rey, CA 90292-6695  
310-822-1511

This research was sponsored by the Defense Advanced Research Projects Agency under contract number DABT63-91-C-0001. Views and conclusions contained in this report are the authors' and should not be interpreted as representing the official opinion or policy of DARPA, the U.S. Government, or any person or agency connected with them.

<b>REPORT DOCUMENTATION PAGE</b>					
1a. REPORT SECURITY CLASSIFICATION Unclassified		1b. RESTRICTIVE MARKINGS			
2a. SECURITY CLASSIFICATION AUTHORITY		3. DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release; distribution is unlimited.			
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE					
4. PERFORMING ORGANIZATION REPORT NUMBER(S) ISI/RR-92-294		5. MONITORING ORGANIZATION REPORT NUMBER(S)			
6a. NAME OF PERFORMING ORGANIZATION USC/Information Sciences Institute	6b. OFFICE SYMBOL	7a. NAME OF MONITORING ORGANIZATION Department of the Army			
6c. ADDRESS (City, State, and ZIP Code) 4676 Admiralty Way Marina del Rey, CA 90292-6695		7b. ADDRESS (City, State, and ZIP Code) Directorate of Contracting, ATTN: ATZS-DKM Post Office Box 748 Fort Huachuca, AZ 85613-0748			
8a. NAME OF FUNDING/SPONSORING ORGANIZATION DARPA	8b. OFFICE SYMBOL	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER DARPA (Fort Huachuca) DABT63-91-C-0001			
8c. ADDRESS (City, State, and ZIP Code) 3701 N. Fairfax Drive Arlington, VA 22203-1714		10. SOURCE OF FUNDING NUMBERS			
		PROGRAM ELEMENT NO. ---	PROJECT NO. ---	TASK NO. ---	
		WORK UNIT ACCESSION NO. ---			
11. TITLE (Include Security Classification) An Architecture for Multimedia Connection Management (Unclassified)					
12. PERSONAL AUTHOR(S) Schooler, E.M., Casner, S.L.					
13a. TYPE OF REPORT Research Report	13b. TIME COVERED FROM _____ TO _____	14. DATE OF REPORT (y/m/d) 1992, August	15. PAGE COUNT 6		
16. SUPPLEMENTARY NOTATION					
17. COSATI CODES		18. SUBJECT TERMS (Use reverse if needed; identify by block no.)  Multimedia, teleconferencing, connection architecture, connection management, configuration management, connection control protocol			
FIELD	GROUP				SUB-GROUP
09	02				
19. ABSTRACT (Continue on reverse if necessary and identify by block number)  What ingredients are needed to enable widespread personal teleconferencing over the Internet and NREN? The major focus of our work in the Multimedia Conferencing Project at ISI is on the design and implementation of protocols at a number of levels in the protocol stack; at the lower levels, real-time data communication services for the Internet in general; and at higher levels, a connection management architecture to facilitate connections among heterogeneous systems. In this paper, we confine our discussion to higher level issues: an overall connection architecture, a connection control protocol, and configuration management.					
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT [X]Unclassified/Unlimited [ ]Same as Rpt [ ]DTIC Users		21. ABSTRACT SECURITY CLASSIFICATION Unclassified			
22a. NAME OF RESPONSIBLE INDIVIDUAL(S) Kathleen McLaughlin/Celeste Anderson		22b. TELEPHONE (310) 822-1511	22c. OFFICE SYMBOL		

# An Architecture for Multimedia Connection Management

Eve M. Schooler  
schooler@isi.edu

Stephen L. Casner  
casner@isi.edu

USC/Information Sciences Institute  
4676 Admiralty Way  
Marina del Rey, CA 90292

## 1 Introduction

What ingredients are needed to enable widespread personal teleconferencing over the Internet and NREN? Certainly the integration of audio and video compression hardware into workstations is essential. However, network protocols are as critical to this goal. This is the major focus of our work in the Multimedia Conferencing Project at ISI. Toward this end, ISI and BBN have developed an experimental packet teleconferencing system that currently is operating at several sites on the Terrestrial Wideband Network, TWBnet [1], and has more recently been ported to the DARPA Research Testbed, DARTnet, for further experimentation. The system allows geographically separated individuals to collaborate by combining real-time packet audio and video with shared computer workspaces, sometimes called *groupware*.


We are designing and implementing protocols at a number of levels in the protocol stack: at the lower levels, real-time data communication services for the Internet in general [2–4]; and at higher levels, a connection management architecture to facilitate connections among heterogeneous systems. In this abstract, we confine our discussion to higher level issues: an overall connection architecture, a connection control protocol, and configuration management.

## 2 A Connection Management Architecture

We have designed a framework for applications such as teleconferencing that require the establishment of multiple participant, multiple media sessions. We propose that such applications be constructed as shown in the diagram below. The central component that orchestrates the multiparty, multimedia connections is the *connection manager*. Conceptually, it is separate from user interfaces (UIs) to the system, which sit above it offering services up to the user and relaying requests down from the user. By separating the connection manager from the user interface, conference-oriented tools avoid duplication of effort; this encompasses management of participation, authentication, and presentation of coordinated UIs. This is along the lines of Swinehart's switching kernel [5] and is related to Lantz' and Lauwers' models for UIs that endow programs with the ability to be multi-user [6]. The connection manager is also separate from underlying components, called *media agents*, which handle communication protocol decisions (transport and internetwork protocols) and the devices specific to each type of shared media (audio, video, groupware). This organization also allows the connection manager to convey timing information between media agents for inter-media synchronization.

We refer to this layered organization as the *connection management architecture*, since management of connections is the primary focus of the model and since the connection manager acts as the conduit through which control information flows between layers. Its modularity allows functionally-equivalent components to substitute for each other. Our current teleconferencing system is the forerunner to this architecture, but we have also drawn ideas from other connection management schemes [7–10]. While these schemes are diverse,

each has suggested the need for a connection management abstraction. The emphasis of our architecture is to provide reliability across WANs, to accommodate heterogeneity, and to facilitate interoperability with other teleconferencing implementations.



*Coordinated Management of Separate Services*

### 3 The Connection Control Protocol

The Connection Control Protocol (CCP) is an application layer protocol used by connection managers to communicate among themselves. It is based on the protocol used by the Multimedia Conference Control Program (MMCC), which has served as both user interface and connection manager for the existing system [11]. The CCP is essentially a multicast, transaction-based protocol. It aims to provide reliable, group communication and to accommodate variability in request-reply response time due to WAN operation and to heterogeneous end system configurations. The specification for CCP has been drafted [12] and implementation is underway.

Conference orchestration is achieved through use of a distributed, peer-to-peer model. Peer connection managers reside on machines scattered throughout the Internet. Each connection manager acts as both client and server, notifying users of requests from other connection managers, and placing requests to other connection managers on the local user's behalf. During conference setup, a connection manager first communicates remotely with peer connection managers via CCP, then communicates locally with the media agents via well-defined call interfaces to actually create the underlying voice, video and groupware data flows.

The CCP specification emphasizes the connection establishment and disconnection procedures. A main task is the four-phase conference setup process. During conference setup, the initiator is responsible for negotiating a common set of capabilities, requesting participation, initiating media connections and propagating information among peers; it acts as a leader until conference creation completes. With the conference in place, the initiator reverts to having no special status, and other sites may be invited, join or leave the conference at any time. When a participant disconnects, the rest of the conference is left intact, unless it was a two-party connection which would be torn down entirely.

Because functionality may vary greatly between teleconferencing implementations, the protocol must be extensible. Thus, the CCP has a built-in mechanism that acts as a transparent conduit for other types of operations between user interfaces and media agents. These operations might occur during or outside of a conference, without CCP having to know the particulars of each operation. An example would be a media-specific function (e.g., mute a particular audio stream).


Due to the distributed nature of the system, cooperating connection managers may get out of synchronization. Site A thinks it is conferencing with site B, but site B thinks it is not in conference at all. CCP counteracts these problems by exchanging state information with every message, triggering active state queries, and using a resynchronization algorithm to resolve state mismatches.

#### 4 Configuration Management

Gaging from our own experience with WAN conferencing, as the number of sites scales up, so does the likelihood for heterogeneity among them. When telecollaboration becomes popular over the Internet, users and their equipment will inevitably fall into different administrative domains. Decisions on codec types, bounds on network bandwidth, the number of cameras connected, etc., will be made by different individuals. One end system's *capabilities* or *configuration* is likely to vary from another. Therefore one of the connection manager's functions is *configuration management* — the need to communicate configuration information and to implement service selection (on demand) so that end systems supporting different combinations of services can still be interoperable in the subset of compatible services. Although we are in the beginning stages of the configuration management design, we propose several mechanisms to deal with heterogeneity: a configuration language, a distributed resource locator service, and a resource synthesizer.

A *configuration language* must be developed to describe resources and devices located at each end system. This language will categorize the set of services offered by each system, and must be extensible to accommodate new services and devices as they are added. The system will maintain a mapping of configuration attributes to media agents, allowing the connection manager module to choose an appropriate agent or agents to meet configuration requirements made by applications. This is crucial for the negotiation phase of the CCP connection establishment procedure.

We also expect to support the dynamic configuration of shared resources so that multimedia devices could be shared among users or applications. For example, a crossbar switch or video codec might be shared among users on a LAN, or devices such as video bridges or coding translators might be remote services available to users throughout the Internet. The system will also provide a procedure for devices, resources and agents to register themselves as available to be shared and will define a *distributed locator service* to provide resource directory functions — the resource description, its location in the network, and the associated configuration parameters.



*Dynamic Configuration Management  
of Shared Resources*

In those instances where capabilities at end systems mismatch entirely (i.e., no compromise configuration can be found), a *resource synthesizer* will synthesize solutions to configuration requests [13]. The resource

synthesizer must find a sequence of communication mechanisms and intermediate translators to construct paths among all the end systems in the connection. The resource synthesizer would draw information from the locator service about shared capabilities that could be used throughout the network. We expect that experience from the network routing domain will be applicable to this problem.

## 5 Future Work

A strategy to bring widespread teleconferencing to the Internet community must include an open connection management architecture. That architecture must address issues in heterogeneous configurations and reliability across wide area networks. The blueprints for just such a framework have been outlined in this abstract and we are now turning our efforts toward its implementation.

## 6 Acknowledgment

This research is sponsored by the Defense Advanced Research Projects Agency under Contract Number DABT63-91-C-0001. The views and conclusions in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency of the U.S. Government.

## 7 References

- [1] Casner, S., Seo, K., Edmond, W., Topolcic, C., "N-Way Conferencing with Packet Video", *Proceedings 3rd International Workshop on Packet Video*, VISICOM '90, Morristown, NJ (Mar 1990).
- [2] Topolcic, C., "Experimental Internet Stream Protocol", RFC 1190, IETF CIP Working Group (Oct 1991).
- [3] Cohen, D., "Specifications for the Network Voice Protocol, Technical Report ISI/RR-75-39, USC/Information Sciences Institute, Marina del Rey, CA (Mar 1976).
- [4] Cole, R., "PVP - A Packet Video Protocol", Internal Document, USC/Information Sciences Institute, Marina del Rey, CA (July 1981).
- [5] Swinehart, D., "The Connection Architecture for the Etherphone System", Xerox PARC, Palo Alto, CA (Dec 1990).
- [6] Lauwers, J.C., Lantz, K.A., "Collaboration Awareness in Support of Collaboration Transparency: Requirements for the Next Generation of Shared Window Systems", *Proceedings CHI '90 Conference on Human Factors in Computer Systems* (1990).
- [7] Bates, P.C., Segal, M.E., "Touring Machine: A Video Telecommunications Software Testbed", *Proceedings First International Workshop on Network and Operating System Support for Digital Audio and Video*, Berkeley, CA (Nov 1990).
- [8] Leung, W.H., Baumgartner, T.J., Hwang, Y.H., Morgan, M.J., Tu, S.C., "A Software Architecture for Workstations Supporting Multimedia Conferencing in Packet Switching Networks", *IEEE Journal on Selected Areas in Communications*, Vol.8, No.3, pp.380-390 (Apr 1990).
- [9] Vin, H.M., Zellweger, P.T., Swinehart, D.C., Rangan, P.V., "Multimedia Conferencing in the Etherphone Environment", *IEEE Computer*, Vol.24, No.10, pp.69-79 (Oct 1991).
- [10] Garcia-Luna-Aceves, J.J., Craighill, E.J., Lang, R., "An Open-Systems Model for Computer-Supported Collaboration", *Proceedings 2nd IEEE Conference on Workstations* (Mar 1988).
- [11] Schooler, E.M., "A Distributed Architecture for Multimedia Conference Control", Technical Report ISI/RS-91-289, USC/Information Sciences Institute, Marina del Rey, CA (Nov 1991).
- [12] Schooler, E.M., "The Connection Control Protocol: Specification", Version 1.0, USC/Information Sciences Institute, Marina del Rey, CA (Jan 1992).
- [13] Nicolaou, C., "An Architecture for Real-Time Multimedia Communications Systems", *IEEE Journal on Selected Areas in Communications*, Vol.8.No.3, pp.391-400 (Apr 1990).