

671.352

Multimedia SYSTEMS

LIBRARY
NOV 17 1994
IOWA STATE UNIVERSITY
of Science & Technology

N **U** **M**
B **R** **I**

Volume 1 · 1993

RPX Exhibit 1019
RPX v. DAE

Multimedia SYSTEMS

Aims and Scope

Multimedia systems publishes original research articles and serves as a forum for stimulating and disseminating innovative research ideas, emerging technologies, state-of-the-art methods and tools in all aspects of multimedia computing, communication, storage, and applications among researchers, engineers, and practitioners. Theoretical, experimental, and survey articles are all appropriate to the journal.

Specific areas of interest include:

- 1 Integration of digital video and audio capabilities in computer systems
- 2 Multimedia information encoding and data interchange formats
- 3 Operating system mechanisms for digital multimedia
- 4 Digital video and audio networking and communication
- 5 Storage models and structures
- 6 Methodologies, paradigms, tools, and software architectures for supporting multimedia applications
- 7 Multimedia applications and application program interfaces, and multimedia endsystem architectures

Copyright

Submission of a manuscript implies: that the work described has not been published before (except in the form of an abstract or as part of a published lecture, review, or thesis); that it is not under consideration for publication elsewhere; that its publication has been approved by all coauthors, if any, as well as by the responsible authorities at the institute where the work has been carried out; that, if and when the manuscript is accepted for publication, the authors agree to automatic transfer of the copyright to the publisher; and that the manuscript will not be published elsewhere in any language without the consent of the copyright holders.

All articles published in this journal are protected by copyright, which covers the exclusive rights to reproduce and distribute the article (e.g., as offprints), as well as all translation rights. No material published in this journal may be reproduced photographically or stored on microfilm, in electronic data bases, video disks, etc., without first obtaining written permission from the publisher.

The use of general descriptive names, trade names, trademarks, etc., in this publication, even if not specifically identified, does not imply that these names are not protected by the relevant laws and regulations.

While the advice and information in this journal is believed to be true and accurate at the date of its going to press, neither the authors, the editors, nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Special regulations for photocopies in the USA: Photocopies may be made for personal or in-house use beyond the limitations stipulated under Section 107 or 108 of U.S. Copyright Law, provided a fee is paid. This fee is US \$ 0.20 per page per copy, plus a basic fee of US \$ 2.00 per article. All fees should be paid to the Copyright Clearance Center, Inc., 21 Congress Street, Salem, MA 01970, USA, stating the ISSN 0942-4962, the volume, and the first and last page numbers of each article copied. The copyright owner's consent does not include copying for general distribution, promotion, new works, or resale. In these cases, specific written permission must first

Subscription information

Volume 1 (6 issues) will appear in 1993.

North America. Recommended annual subscription rate: approx. US \$ 238.50 (single issue price: approx. US \$ 46.00) including carriage charges. Subscriptions are entered prepayment only. Orders should be addressed to:

Springer-Verlag New York Inc.
Service Center Secaucus
44 Hartz Way
Secaucus, NJ 07094, USA
Tel. (201) 348-4033, Telex 0023-125994
Fax (201) 348-4505

All other countries: Recommended annual subscription rate: DM 320.00 plus carriage charges; [Federal Republic of Germany: DM 11.24 incl. value added tax; all other countries: DM 21.00 except for the following countries to which SAL delivery (Surface Air-mail Lifted) is mandatory: Japan, India, Australia/New Zealand. SAL charges and airmail delivery to all other countries are available upon request].

Volume price: DM 320.00, single issue price: DM 64.00 plus carriage charges. Reduced rate for personal subscribers is available upon request.

Orders can either be placed via a bookdealer or sent directly to:

Springer-Verlag, Heidelberger Platz 3
D-14197 Berlin, Germany
Tel. (0) 30/8207-0, Telex 183319
FAX (0) 30/8214091

Changes of address: Allow six weeks for all changes to become effective. All communications should include both old and new addresses (with postal codes) and should be accompanied by a mailing label from a recent issue.

According to § 4 section 3 of the German Postal Services Data Protection Regulations, the German Federal Post Office can inform the publisher of a subscriber's new address even if the subscriber has not submitted a formal application for mail to be forwarded. Subscribers not in agreement with this procedure may send a written complaint to Springer-Verlag's Berlin office within 14 days of publication of this issue.

Microform

Microform editions are available from:
University Microfilms International 300 N.
Zeeb Road, Ann Arbor, MI 48106, USA

Production

Springer-Verlag, May Doehl
Journal Production Department II
Postfach 105280, D-69042 Heidelberg,
Germany
Tel. (0) 6221/487439, FAX (0) 6221/487625

Responsible for advertisements

Springer-Verlag, E. Lückermann
Heidelberger Platz 3, D-14197 Berlin, Germany
Tel. (0) 30/8207-0, Telex 185411
FAX (0) 30/8207300

Printers

Zechnersche Buchdruckerei
D-67346 Speyer, Germany
© Springer-Verlag Berlin Heidelberg 1993
Springer-Verlag GmbH & Co. KG
D-14197 Berlin, Germany
Printed in Germany



Springer

**DOCKET
ALARM**

Find authenticated court documents without watermarks at docketalarm.com.

Volume 1
Number 1 1993

Multimedia SYSTEMS

Rangan PV, Ferrari D, Herrtwich RG	Editorial 1
Schooler EM	The impact of scaling on a multimedia connection architecture 2
Zhang HJ, Kankanhalli A, Smoliar SW	Automatic partitioning of full-motion video 10
Borenstein NS	MIME: a portable and robust multimedia format for internet mail 29
Amenyo JT, Lazar AA, Pacifigi G	Proactive cooperative scheduling and buffer management for multimedia networks 37
	Events 50



Springer International

This material may be protected by Copyright law (Title 17 U.S. Code)

The impact of scaling on a multimedia connection architecture

Eve M. Schooler

University of Southern California/Information Sciences Institute, 4676 Admiralty Way, Marina del Rey, CA 90292, USA
e-mail schooler@isi.edu

Received January 10, 1993/Accepted April 10, 1993

Abstract. As the last two meetings of the Internet Engineering Task Force have shown, the demand for Internet teleconferencing has arrived. Packet audio and video have now been multicast to approximately 170 different hosts in ten countries, and for upcoming meetings the number of remote participants is likely to be substantially larger. Yet the network infrastructure to support wide-scale packet teleconferencing is not in place. These experiments represent a departure from the two- to ten-site telemeetings that are the norm today. They represent an increase in scale of multiple orders of magnitude in several interrelated dimensions.

This paper discusses the impact of scaling on our efforts to define a multimedia teleconferencing architecture. Three scaling dimensions of particular interest are: (1) very large numbers of participants per conference, (2) many simultaneous teleconferences, and (3) a widely dispersed user population. Here we present a strawman architecture and describe how conference-specific information is captured, then conveyed among end systems. We provide a comparison of connection models and outline the tradeoffs and requirements that change as we travel along each dimension of scale. In conclusion, we identify five critical needs for a scalable teleconferencing architecture.

Key words: Packet videoconferencing – Connection architecture – Scalability – Multimedia

1 Overview of a connection management architecture

We have proposed a multimedia connection architecture that has served as the basis for discussion on remote conferencing architectures within the Internet Engineering Task Force (IETF) (Schooler 1992b). At the core of the modular architecture is the notion of a *connection manager*, which resides at each end system to coordinate the orchestration, maintenance, and interaction of multiuser sessions. Per-site connection managers communicate with peers by using a distributed connection control protocol (Schooler 1992a). Conceptually, the connection manager is separate from user interfaces to

the connection manager from the user interface, conference-oriented tools avoid duplication of effort. This encompasses the management of participation, authentication, and presentation of coordinated user interfaces. The connection manager is also separate from the underlying components, so that it is shielded from the decisions specific to each type of shared media (audio, video, groupware).

The connection manager acts as a conduit for control information not only remotely among peer connection managers, but also among other local conference-related components as depicted in Fig. 1. Connection managers are loosely coupled with *media agents* that implement the media processing and data communication functions. With media-specific details relegated to underlying media agents, functional commonality is distilled in the connection manager. The connection manager provides general mechanisms for session-related tasks (connect, invite, etc.) and acts as a broker to share information across media agents (participant lists, admission policies, etc.).

Modularity allows dependencies on particular hardware or communications facilities to be encapsulated within individual components of the system for easier deployment into new environments and offers the connection manager a selection of choices in media agent capabilities. Thus, the connection manager's other principal responsibility is *configuration management* of end-system heterogeneity. End-system differences include asymmetries in available media, codec mismatches, variations in bandwidth capabilities, transport incompatibilities, etc. Accordingly, the connection manager's control protocol negotiates a workable set of capabilities among group members (e.g., quality vs cost, MPEG vs. H.261).

The intent of the architecture is to facilitate interoperation among users' teleconferencing implementations across the Internet. Therefore, the connection manager is used to capture high-level configuration descriptions from users (e.g., the collection of media in which the user is interested, quality of service preferences, etc.), then conveys the requested configuration to peer connection managers. Each connection manager in turn provides more detailed descriptions to its media agents, which translate the configuration requests into

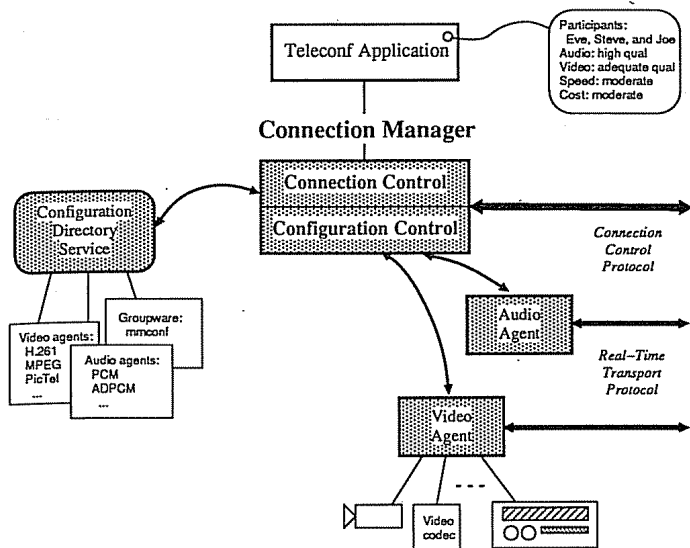


Fig. 1. Flow of control information

tridge 1992). Peer connection managers work to negotiate a suitable configuration by relying on interactions between each connection manager and its media agents, and between each media agent and the underlying network services.

For example, in the simplified scenario in Fig. 1, an application asks the connection manager for high-quality audio and adequate-quality video over moderate-speed links for moderate cost. The configuration directory service is consulted by the connection manager and identifies media agents that both meet the specifications and are available. In this case, the configuration directory service translates quality, speed and cost into media agents that provide specific encoding/data rate capabilities. Once notified, the initiator's local media agents may opt to reserve any devices (cameras, codecs, etc.) and network bandwidth upon which they will rely. The initiator's connection manager then communicates the request to the other participants' connection managers, negotiating particulars as needed. At this stage, the remote connection managers go through the same process of locating appropriate media agents and reserving the required resources. Finally, each connection manager instructs its local media agents to begin sending data, which means that the media agents establish real-time transport sessions (Schulzrinne 1992b). In a more optimistic scheme, the media agents would wait to reserve resources until all members have actually responded to the initial participation request; delayed reservation however may lead to service denial.

2 The problem of scale

Most experimentation with packet teleconferencing systems has been conducted within local area network (LAN) settings, with few users and with a modest degree of support for simultaneous conferences. In Fig. 2 we display a sampling of these systems. The x -axis denotes users per conference, the

which each system supports simultaneous teleconferencing sessions.

Although shared workspace applications, such as MM-Conf, function across WANs, they perform markedly better within LANs (Crowley 1990). This comes as no surprise since to maintain an actively changing global view of the workspace, these applications require reliable communication among all users. Typically the application is built on top of an N -by- N collection of TCP/IP streams, which can be problematic within the general Internet (Postel 1981; Postel 1980). A badly timed network outage or routing problem between one pair of conferees might lead to inconsistency in the shared view. To reconstitute the state, a WAN-sensitive session protocol might be layered above the transport to detect and correct peers that are out of synchronization.

Real-time teleconferencing systems, such as Etherphone and Phoenixphone, the CAR project and various CoDesk applications, support digital media over a LAN with centralized conference management (CM) (Eriksson 1992; Handley 1992; Swinehart 1990; Vin et al. 1991). In contrast, the Touring Machine and Rapport represent a class of systems that combine analog media with centralized computer-based session control (Ahuja et al. 1988; Arango 1992). In both cases, concurrency is supported, but only as much as the media crossbar switches or the LANs can physically support. To approximate WAN conferencing, analog systems use a proxy to link two distinct LAN communities through a commercial codec.

The second row of diagrams shows systems that are well-equipped for certain aspects of WAN operation by virtue of their decentralized architectures (Chen et al. 1992; Schooler, in press; Schulzrinne 1992a; Turletti 1993). In addition, the multimedia conference control program (MMCC) was designed to accommodate the likelihood in a WAN environment of heterogeneity at the end systems and the need to provide robust sessions across the network (Schooler, in press). Popular IETF tools, such as the visual audio tool (vat) from Lawrence Berkeley Laboratory, Xerox PARC's network video program (nv), the INRIA videoconferencing system (ivs), the network voice terminal (nevot) from University of Massachusetts, and Bolt, Beranek and Newman's desktop video conferencing program (dvc), specifically use a lightweight session model to support larger conferences of widely distributed participants (Schulzrinne 1992a; Turletti 1993). All of these systems, however, are bound in varying degrees by the number of users per conference. None provide explicit support for large numbers of concurrent conferences, due to the Internet's lack of infrastructure for real-time media and wide-scale multicast addressing. These last two classes of systems, formally differentiated by their style of session moderation, will be contrasted in a later section.

As can be seen in all five diagrams, even projects that scale in one dimension typically have architectural deficiencies in the other dimensions. To understand the prob-

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.