P. Venkat Rangan (Ed.)



Network and Operating System Support for Digital Audio and Video

Third International Workshop La Jolla, California, USA, November 12-13, 1992 Proceedings

Springer-Verlag

Berlin Heidelberg New York London Paris Tokyo Hong Kong Barcelona Budapest Series Editors

Gerhard Goos Universität Karlsruhe Postfach 6980 Vincenz-Priessnitz-Straße 1 D-76131 Karlsruhe, Germany Juris Hartmanis Cornell University Department of Computer Science 4130 Upson Hali Ithaca, NY 14853, USA

Volume Editor

P. Venkat Rangan
Depart, of Computer Science & Engineering, University of California at San Diego
9500 Gilman Drive, La Jolla, California 92093-0114, USA

300

CR Subject Classification (1991): H.5.1, C.2, D.4, H.4.3

ISBN 3-540-57183-3 Springer-Verlag Berlin Heidelberg New York ISBN 0-387-57183-3 Springer-Verlag New York Berlin Heidelberg

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer-Verlag. Violations are liable for prosecution under the German Copyright Law.

© Springer-Verlag Berlin Heidelberg 1993 Printed in Germany

Typesetting: Camera-ready by authors Printing and binding: Druckhaus Beltz, Hemsbach/Bergstr. 45/3140-543210 - Printed on acid-free paper

Preface

Technological advances are revolutionizing computers and networks to support digital video and audio, leading to new design spaces in computer systems and applications. Under the surface of exciting multimedia technologies lies a mine of research problems. The goal of this workshop, which was sponsored by IEEE Computer and Communication Societies, in cooperation with ACM SIGCOMM, SIGOIS and SIGOPS, was to bring together the leading researchers in all aspects of multimedia computing, communication, storage, and applications.

The field of multimedia has witnessed an explosive growth in the last few years. This workshop, the third in the series (the first was in November 1990 at Berkeley, California, and the second in November 1991 in Heidelberg, Germany) attracted a record number of 128 submissions from four continents (America, Asia, Australia and Europe). Each submission was reviewed by three program committee members. The selection of papers was extremely competitive, with only 26 full papers and 14 short papers being accepted. There were seven full paper sessions and five short paper sessions. In addition, there was an illuminating invited keynote banquet address by Professor Edward A. Fox of Virginia Polytechnic and State University on the topic "Progressing Towards a Hypermedia National Library". In order to keep the workshop environment conducive to in-depth interactions among active researchers, attendance at the workshop was by invitation only. The number of attendees at the workshop was about 95 from 15 countries.

This book constitutes the formal proceedings of the workshop. Authors of the best papers of the workshop have been encouraged to submit full papers for early publication in the ACM/Springer-Verlag journal *Multimedia Systems* (the first journal to focus exclusively on issues relating to multimedia), and also to ACM Multimedia 93, the first ACM international conference on multimedia, August 2-6, 1993, Anaheim, California.

I would like to thank the program committee members for their timely and in-depth reviews, which contributed greatly to constituting what was a very solid workshop program. My thanks also go to my students, Harrick Vin and Srinivas Ramanathan, and my secretary, Ida O'Neil, for their help and involvement in making arrangements for this workshop.

In closing, I am pleased to present this book reporting on the state-of-the-art research in multimedia systems.

June 1993

P. Venkat Rangan Program Chair

Program Committee

P. Venkat Rangan (Chair)

University of California at San Diego, USA

Sudhir R. Ahuja Gordon Blair Rita Brennan Stephen Casner Stavros Christodoulakis Flaviu Christian Domenico Ferrari Riccardo Gusella Ralf Guido Herrtwich Jim Kurose Desai Narasimhalu Duane Northcutt Craig Partridge Jonathan Rosenberg **David Sincoskie** Jean-Bernard Stefani **Daniel Swinehart** David Tennenhouse Radu Popescu-Zeletin

AT&T Bell Labs, USA University of Lancaster, UK Apple Computer, USA ISI, University of Southern California, USA Technical University of Crete, Greece University of California at San Diego, USA University of California at Berkeley, USA Hewlett Packard Labs, USA IBM European Networking Center, Germany University of Massachusetts at Amherst, USA National University of Singapore, Singapore Sun Labs, USA BBN, USA Belkore, USA Bellcore, USA CNET, France Xerox PARC, USA MIT, USA GMD FOKUS, Germany

Table of Contents

Session 1: Network and Operating System Support for Multimedia Chair: Domenico Ferrari	1
Adaptive, Best-Effort Delivery of Digital Audio and Video Across Packet-Switched Networks K. Jeffay, D.L. Stone, T. Talley, F.D. Smith	3
The Architecture of Rattlesnake: A Real-Time Multimedia Network Gerard J.M. Smit, Paul J.M. Havinga	15
Beyond ST-II: Fulfilling the Requirements of Multimedia Communication Ralf Guido Herrtwich, Luca Delgrossi	25
An Approach to Real-Time Scheduling - But is It Really a Problem for Multimedia? Roger Needham, Akira Nakamura	32
Session 2: Multimedia On-Demand Services Chair: Jonathan Rosenberg	40
Design and Analysis of a Grouped Sweeping Scheme for Multimedia Storage Management Philip S. Yu, Mon-Song Chen, Dilip D. Kandlur	44
Admission Control Algorithms for Multimedia On-Demand Servers Harrick M. Vin, P. Venkat Rangan	56
The Design and Implementation of a Continuous Media Storage Server Phillip Lougher, Doug Shepherd	69
Performance Studies of Digital Video in a Client/Server Environment Kathleen M. Nichols	81
Session 3: Media Synchronization Chair: Riccardo Gusella	92
Basic Synchronization Concepts in Muldimedia Systems Luiz F. Rust da Costa Carmo, Pierre de Saqui-Sannes, Jean-Pierre Courtiat	94
Synchronization in Joint-Viewing Environments Kurt Rothermel, Gabriel Dermler	106

Synchronization of Multi-Sourced Multimedia Data for Heterogeneous Target Systems Dick C.A. Bulterman	119
An Intermedia Skew Control System for Multimedia Data Presentation T.D.C. Little, F Kao	130
Session 4: Distributed Multimedia Systems Chair: J.J. Garcia-Luna Aceves	142
System Support for Efficient Dynamically-Configurable Multi-Party Interactive Multimedia Applications Mark Moran, Riccardo Gusella	143
Requirements for Network Delivery of Stored Interactive Multimedia Darren New, Jonathan Rosenberg, Gil Cruz, Tom Judd	157
Multimedia Processing Model for a Distributed Multimedia I/O System Rusti Baker, Alan Downing, Kate Finn, Earl Rennison, Doollyun David Kim, Young Hwan Lim	164
Enhancing the Touring Machine API to Support Integrated Digital Transport Mauricio Arango, Michael Kramer, Steven L. Rohall, Lillian Ruston, Abel Weinrib	176
Session 5: Network and Operating System Support for Multimedia Chair: Ralf Guido Herrtwich	183
Preliminary Measurement of the RMTP/RTIP Hui Zhang, Tom Fisher	185
The Multimedia Multicast Channel Joseph C. Pasquale, George C. Plyzos, Eric W. Anderson, Vachaspathi P. Kompella	197
Analysis of a Resequencer Model for Multicast over ATM Networks Liming Wei, FongChing Liaw, Deborah Estrin, Allyn Romanow, Tom Lyon	209
Session 6: Multimedia Models, Frameworks, and Document Architectures Chair: Rita Brennan	221
An Integrated Platform and Computational Model for Open Distributed Multimedia Applications G. Blair, G. Coulson, P. Auzimour, L. Hazard, F. Horn, J.B. Stefani	223

Scheduling Multimedia Documents Using Temporal Constraints M. Cecelia Buchanan, Polle T. Zellweger	237
The Stratification System: A Design Environment for Random Access Video Thomas G. Aguierre Smith. Glorianna Davenport	250
On the Design of Multimedia Interchange Formats John F. Koegel	262
Session 7: Multimedia Workstations and Platforms Chair: Daniel C. Swinehart	272
Bus Bandwidth Management in a High Resolution Video Workstation Gerard A. Wall, James G. Hanko, J. Duane Northcutt	274
Analysis of I/O Subsystem Design for Multimedia Workstation Peter Druschel, Mark B. Abbott, Michael Pagels, Larry L. Peterson	289s
Tactus: Toolkit-Level Support for Synchronized Interactive Multimedia Roger B. Dannenberg, Tom Neuendorffer, Joseph M. Newcomer, Dean Rubine	302
Short Paper Session I: Scheduling and Synchronization Chair: Jim Kurose	314
An Analytical Model for Real-Time Multimedia Disk Scheduling James Yee, Pravin Varaiya	315
Real-Time Scheduling Support in Ultrix-4.2 for Multimedia Communication Tom Fisher	321
Continuous Media Synchronization in Distributed Multimedia Systems Srinivas Ramanathan, P. Venkat Rangan	328
Short Paper Session II: Architectures and Environments Chair: Duane Northcutt	336
High Speed Networks and the Digital TV Studio Guy Cherry, Jim Nussbaum, Mayer Schwartz	337
The Impact of Scaling on a Multimedia Connection Architecture Eve M. Schooler	341

IX

Short Paper Session III: Networking and Protocol Support Chair: Stephen Casner	347
An Admission Control Algorithm for Predictive Real-Time Service (Extended Abstract) Sugih Jamin, Scott Shenker, Lixia Zhang, David D. Clark	349
MEGAPHONE: A Multimedia Application Based on Object- Oriented Communication Daniel P. Ingold	357
System Support for Dynamic QOS Control of Continuous Media Communication Stephen TC. Chou, Hideyuki Tokuda	363
Short Paper Session IV: Multimedia Retrieval Chair: David Sincoskie	369
NMFS: Network Multimedia File System Protocol Sameer Patel, Ghaleb Abdulla, Marc Abrams, Edward A. Fox	370
A Continuous Media Player Lawrence A. Rowe, Brian C. Smith	376
Architecture of a Multimedia Information System for Content-Based Retrieval Deborah Swanberg, Chiao Fe Shu, Ramesh Jain	387
Short Paper Session V: Object-Oriented Systems and Toolkits Chair: Jean-Bernard Stefani	393
Application Construction and Component Design in an Object- Oriented Multimedia Framework Simon Gibbs	394
Audio and Video Extensions to Graphical User Interface Toolkits Rei Hamakawa, Hidekazu Sakagami, Jun Rekimoto	399
An Introduction to HeiMAT: The Heidelberg Multimedia Application Toolkit Thomas Käppner, Dietmar Hehmann, Ralf Steinmetz	405

Session 1: Network and Operating System Support for Multimedia

Chair: Domenico Ferrari, University of California at Berkeley

The first session of the workshop consisted of four presentations that covered a wide variety of topics: from a new type of local-area network for multimedia applications to LAN support for video conferencing, from an implementation of an internetwork real-time protocol to priority inversion countermeasures in the context of multimedia applications.

Kevin Jeffay of the University of North Carolina described one of the many current projects intended to explore what can be done with present-day networks and protocols in the area of multimedia applications. Unlike other efforts, this one relied on a real-time operating system, YARTOS, developed by Jeffay and his coauthors. As long as their results are not interpreted in a much more general and long-term prospective than warranted by their here-and-now, ad-hoc scope (for example, if they are not used to "demonstrate" that admission control is not needed because "the system works even without it"), these studies are useful, as they often enhance our understanding of the requirements of continuous-media applications and our experience with them. The case described in the presentation was the one of desktop conferencing between two users across an Ethernet or another LAN. Particular attention in the talk was paid to the novel transport protocol designed to manage adaptively the scarce and time-variant available bandwidth, deliver audio packets reliably, reduce end-to-end delay and jitter, and deal satisfactorily with packet losses. The protocol runs on top of IP, and uses some more bandwidth than UDP, but is obviously better than UDP for conferencing applications.

The second presentation was given by Gerard Smit of the University of Twente in the Netherlands. It provided an overview of, and some of the details about, a new type of local-area network called Rattlesnake, which has been designed considering the needs of real-time applications as well as those of the more traditional ones. Like DEC's AN1, it is characterized by an aggregate bandwidth much greater than that of each link; however, Rattlesnake is based on a fixed topology, that of a Kautz graph, and is not selfreconfiguring; on the other hand, AN1 does not make any distinction between real-time and non-real-time traffic, whereas Rattlesnake relies on hybrid TDM to provide circuit switching to the former and packet switching to the latter. Also, cut-through is used for real-time packets, while the non-real-time ones are stored in the switches and then forwarded. The choice of hybrid TDM was, in the opinion of some of the attendees, not sufficiently motivated in the presentation, as the main argument seemed to ignore the existence of a number of schemes that can provide hard real-time guarantees even with AIM or, more generally, packet switching. The properties of Kautz topologies were considered interesting; in particular, the fixed degree, the fault tolerance, and the self-routing capabilities even in presence of link failures.

Ralf Herrtwich of the IBM European Networking Center in Heidelberg, Germany, talked about the implementation of the ST-II protocol that was done at ENC. He summarized the reasons for choosing ST-II as the internetwork protocol in the Heidelberg Transport System (HeiTS), which has been built to support a distributed multimedia platform. He also discussed the additions and enhancements made by him and his coauthor to ST-II to complete it and make it better suited to the needs of multimedia traffic; in particular, resource management mechanisms, characterized by optimistic (non-worstcase) reservation and graceful degradation, and a feedback scheme to allow rates to be controlled without having to keep all clocks in synchrony. This implementation of ST-II is probably the most complete and interesting to date. It is to be hoped that its authors will soon report on their experiments and experience with it.

The fourth and last paper of the session was presented by Akira Nakamura of Cambridge University in England. The connection between the two parts of the presentation was not easily grasped by the audience. In its first part, much longer than the second, the paper describes a new algorithm to deal with the priority inversion problem; the algorithm is less conservative than the priority ceiling protocol, which may cause waiting on a shared lock even when it is not necessary. In the second part, the question is raised whether the scheduling problems that may be caused by shared locks are really encountered in multimedia applications. The authors of the paper argue, unfortunately without providing any experimental evidence, that these applications may encounter lock conflicts, if any, between processing stages for the same stream rather than for access to data structures shared by several independent streams. Thus, in a sense, operating systems should not be extended, but rather reduced, and in any case, rethought and redesigned, for multimedia applications. The arguments presented in the second part of the talk sparked a heated debated between those who agreed and those who disagreed with the main thesis of the authors; the discussion was unfortunately cut short by the expiration of the session's deadline.

Adaptive, Best-Effort Delivery of Digital Audio and Video Across Packet-Switched Networks

K. Jeffay, D.L. Stone, T. Talley, F.D. Smith University of North Carolina at Chapel Hill Department of Computer Science Chapel Hill, NC 27599-3175 USA {jeffay,stone,talley,smithfd}@cs.unc.edu

Abstract: We present an overview of a "best-effort" transport protocol that supports conferencing with digital audio and video across interconnected packet switched networks. The protocol delivers the highest quality conference service possible given the current load in the network. Quality is defined in terms of synchronization between audio and video, the number of frames played out of order, and the end-to-end latency in the conference. High quality conferences are realized through four transport and display mechanisms and a real-time implementation of these mechanisms that integrates operating system services (e.g., scheduling and resource allocation, and device management) with network communication services (e.g., transport protocols). In concert these mechanisms dynamically adapt the conference frame rate to the bandwidth available in the network, minimize the latency in the displayed streams while avoiding discontinuities, and provide quasi-reliable delivery of audio frames.

1. Introduction

The focus of this work is on the real-time transmission of live digital audio and video across interconnected packet-switched networks. Our goal is to support high-fidelity audio/video conferences, *e.g.*, conferences with high quality audio and full-motion color video. The approach one adopts for the real-time communication of audio and video will depend on numerous factors including the architecture of the audio/video subsystem, the encoding and compression technologies employed in the audio/video subsystem, the available network bandwidth, and the degree of physical layer support in the network for real-time communication. As a starting point we consider existing local area networks (*i.e.*, ethernets, token rings, and FDDI rings) interconnected by bridges or routers, and an audio/video system that acquires and compresses individual frames of video at NTSC rates.

One technical challenge in this environment is to manage and accommodate jitter in the audio/video streams seen by the processes responsible for displaying the streams. In order to sustain a high-fidelity conference, frames of audio/video data must arrive at a displaying workstation so as to be played at a precise rate (e.g., one frame every 33 ms for NTSC video). Although these frames can be generated at the desired rate by the originating audio/video hardware, the exact rate at which frames arrive at a receiver can be grossly distorted by poor operating system scheduling and resource allocation on the transmitting and receiving workstations, and varying load in the network. Jitter is problematic because it can cause discontinuities in the playing of audio/video streams and can increase the latency of the conference. Jitter therefore directly impacts conference quality.

^{*} This work supported by the National Science Foundation (grant numbers CCR-9110938 and ICI-9015443), and by the Digital Equipment Corporation and the IBM Corporation.

It is instructive to view a video conferencing system as a distributed pipeline. Frames of audio/video data are acquired, digitized, and compressed at one machine, transmitted over a network to a second machine where they are decompressed and displayed. Ideally each stage of this pipeline should operate in real-time, that is, each stage should process a frame before the previous stage outputs the following frame. Through the use of a real-time operating system that provides computation and communication services specifically tailored to the needs of applications that process continuous-time media streams such as digital audio and video [4], we can closely control the processing of audio/video frames on a transmitting workstation to ensure that frames progress from acquisition and compression hardware to the network interface with minimal jitter. Correspondingly, on the receiving machine, we ensure the arriving frames are delivered to the audio/video subsystems for display with jitter no worse than in the arriving stream. However, in our environment we cannot control other uses of the network and hence we cannot provide the same level of control over the transmission of individual frames.

We have developed a set of techniques for the transmission and display of audio/video frames that yield the highest quality conference possible given the conditions present in the network. First, in the transport layer, audio is transmitted redundantly (*i.e.*, units of audio data are transmitted multiple times). This allows a conference receiver to function well in the face of packet loss. Second, real-time scheduling and queue management techniques are used in the transport and network interface layers to manage conference latency and automatically adapt the data rate of the conference to varying network bandwidth. At a conference receiver the display of audio and video data follows a protocol to (further) manage conference latency and ameliorate deviations in inter-arrival times of audio/video data. The third mechanism dynamically varies the synchronization between displayed audio and video. This allows a receiver to trade off synchronization between audio and video for low conference latency and the ability to accommodate large deviations in interarrival times of audio/video data. Lastly, a data aging mechanism is used to reduce the latency in a conference after periods of high network load. Together, these mechanisms form a transport and display protocol for conferencing. It is argued that these mechanisms, in concert, provide the best possible conference service given the available network bandwidth at any point in time.

These techniques have been implemented in a multimedia transport protocol (MTP) that has been used to transmit live digital audio and video across packet-switched networks using Internet (IP) protocols [7]. We have investigated the effects of these transport and display techniques on the quality of the audio and video streams displayed at the receiver under a variety of network configurations and loading conditions. In this paper we survey two of the techniques we have developed to ameliorate the effects of jitter. These are the dynamic variation of audio/video synchronization and the adaptation of the conference frame rate to that currently sustainable in the network. The following section presents our requirements for audio/video transmission and describes the technologies on which our work is based, and describes the salient communications problems this work addresses. We follow in Section 3 with a description of the transport and display mechanisms that address jitter. We conclude in Section 4 with a review of our results.