Welcome to the on-line proceedings of NOSSDAV'95!

The 5th International Workshop on Network and Operating Systems Support for Digital Audio and Video

Sponsored by IEEE Communications Society

in cooperation with: ACM SIGCOMM, SIGOPS, SIGMM, SIGGRAPH, SIGIR,
Durham, New Hampshire, April 18 - 22, 1995.

Revised, and now available in print as Lecture Notes in Computer Science 1018 (Springer-Verlag)!

Introduction by the Chairs

It is with great pleasure and an even greater sense of anticipation that we welcome you to the 5th International Workshop on Networking and Operating System Support for Digital Audio and Video (NOSSDAV '95).

The Program Committee selected 23 of the 101 submissions to form the basis of eight technical sessions. We have asked each Session Chair to dedicate a generous amount of time to discussion. Although there will be some differences from session to session to accommodate the nature of the topics addressed and the preferences of the Chairs, in general the discussion will be centered on a round-table consisting of adjunct papers and various invited speakers.

While the technical program will look at the state of the art in networking and operating system support for multimedia, a ninth session will take a broader look at the NOSSDAV workshop charter. Starting from what the workshop has achieved over the last five years, this session will explore possible directions for future NOSSDAV workshops.

We have decided--a first for NOSSDAV--to make the proceedings available electronically, and hope that this will contribute to disseminating the workshop contributions to the largest possible audience.

We are indebted to our Program Committee, and in particular to the Session Chairs for the work that went into the planning and realization of the workshop. Special thanks to Dinesh Venkatesh and Mary Hendrix for their invaluable support in workshop registration and Proceedings publication. We would also like to acknowledge the support of the IEEE Communications Society, the workshop sponsor, and ACM SIGCOMM, SIGGRAPH, SIGOPS, SIGMM, and SIGIR. Finally, our appreciation goes to Hewlett-Packard Laboratories and Boston University for the resources made available to us during the last several months.

Because of the limited attendance and the intimate setting of Durham, New Hampshire, we hope that the workshop will prove to be a fruitful learning and exciting experience for all participants.

Riccardo Gusella
T.D.C. Little
April 1, 1995

Accessing the Proceedings

The proceedings are accessed by through the sequence
author --> abstract --> paper.ps, **or,** title --> abstract --> paper.ps

- Index by author
- Index by title
- Program Committee
- Worshop Particpants (pictures)

Comments on the electronic proceedings are welcome. Send to *mcl@spiderman.bu.edu*.

Lecture Notes in Computer Science 1018

The proceedings are published by Springer-Verlag.

The book information is as follows:

Volume editor: T.D.C. Little, R. Gusella
Title: Network and Operating System Support for
Digital Audio and Video
Subtitle: Fifth International Workshop, NOSSDAV'95
Durham, New Hampshire, USA, April 1995
Proceedings
Lecture Notes in Computer Science series, Vol. 1018
ISBN: 3-540-60647-5

Welcome to the on-line proceedings of NOSSDAV'95!

Revised, and now available in print as Lecture Notes in Computer Science 1018 (Springer-Verlag)!

Advance Reservation Systems, Chair: Kevin Jeffay

- **M. Degermark, T. Kohler, S. Pink and O. Schelen**, Advance Reservations for Predicted Service
- **D. Ferrari, A. Gupta and G. Ventre**, Distributed Advance Reservation of Real-Time Connections
- **L. Wolf, L. Delgrossi, R. Steinmetz, S. Schaller and H. Wittig**, Issues of Reserving Resources in Advance

Operating System Support, Chair: Duane Northcutt

- **P. Leydekkers, V. Gay and L. Franken**,A Computational and Engineering View on Open Distributed Real-Time Multimedia Exchange
- **M.B. Jones, P. Leach, R. Draves and J. Barrera, III**, Support for User-Centric Modular Real-Time Resource Management in the Rialto Operating System
- **K. Jeffay and D. Bennett**, A Rate-Based Execution Abstraction for Multimedia Computing
- **C.J. Lindblad**, VuSystem Performance Measurements
- **C.D. Cranor and G.M. Parulkar**, Design of Universal Continuous Media I/O[*]
- **H. Kitamura, K. Taniguchi, H. Sakamoto and T. Nishida**, A New OS Architecture for High Performance Communication Over ATM Networks -- Zero-copy Architecture --

Resource Management and Quality of Service, Chair: Aurel Lazar

- **A. Eleftheriadis and D. Anastassiou**, Meeting Arbitrary QoS Constraints Using Dynamic Rate Shaping of Coded Digital Video
- **A. Campbell, D. Hutchinson and C. Aurrecoechea**, Dynamic QoS Management for Scalable Video Flows
- **D. Venkatesh and T.D.C. Little**, Dynamic Service Aggregation for Efficient Use of Resources in Interactive Video Delivery
- **K. Kawachiya, M. Ogata, N. Nishio and H. Tokuda**, Evaluation of QoS-Control Servers on Real-Time Mach
- **L. Schreier and M. Davis**, System-Level Resource Management for Network-Based Multimedia Applications

NOSSDAV Retrospective, Chairs: Domenico Ferrari and Dan Swinehart

Audio and Video Systems, Chair: Raj Yavatkar

- **N. Chaddha, G. Wall and B. Schmidt**, An End-to-End Software Only Scalable Video Delivery System
- **S. Cen, C. Pu, R. Staehli, C. Cowan and J. Walpole**, A Distributed Real-Time MPEG Audio Player
- **J.C. Bolot, H. Crepin and A. Garcia**, Analysis of Audio Packet Loss over Packet-Switched Networks
- **H.C. Lauer, C. Shen, R. Osborne, J. Howard and Q. Zheng, M, Takegaki, H. Shimakawa and I. Mizunuma**, Digital Audio and Video in Industrial Systems
- **K. Fall, J. Pasquale and S. McCanne**, Workstation Video Playback Performance with Competitive Process Load

Scheduling and Synchronization, Chair: Chuck Kalmanek

- **K. Rothermel and T. Helbig**, An Adaptive Stream Synchronisation Protocol
- **B. Schmidt, J. Northcutt, and M. Lam**,A Method and Apparatus for Measuring Media Synchronization
- **J. Nieh and M.S. Lam**,Integrated Processor Scheduling for Multimedia
- **A. Mauthe and G. Coulson**, Scheduling and Admission Testing for Jitter Constrained Periodic Threads
- **R. Yavatkar and K. Lakshman**, A CPU Scheduling Algorithm for Continuous Media Applications

Multicasting, Chair: Jim Kurose

- **H. Schulzrinne**, Dynamic Configuration of Conferencing Applications Using Pattern-Matching Multicast
- **E. Biersack and J. Nonnenmacher**, Wave: A New Multicast Routing Algorithm for Static and Dynamic Multicast Groups
- **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
- **K.C. Almeroth and M.H. Ammar**, The Role of Multicast Communication in the Provision of Scalable and Interactive Video-on-Demand Service

Network Scheduling and Real-Time Networking, Chair: Steve Pink

- **H. Zhang and E.W. Knightly**, A New Approach to Support VBR Video in Packet-Switched Networks
- **P. Goyal, S. Lam and H. Vin**, Determining End-to-End Delay Bounds in Heterogeneous Networks
- **V. Bansal, R.J. Siracusa, J.P. Hearn, G. Ramamurthy and D. Raychaudhuri**, Adaptive QoS-Based API for ATM Networking
- **S.S. Lam and G.G. Xie**, Burst Scheduling Networks: Flow Specification and Performance Guarantees
- **M. Grossglauser, S. Keshav and D. Tse** , The Case Against Variable Bit Rate Service

Storage Architectures, Chair: Jon Walpole

- **S.W. Lau and J.C.S. Lui**, A Novel Video-on-demand Storage Architecture for Supporting Constant Frame Rate with Variable Bit Rate Retrieval
- **A.J. Chaney, I.D. Wilson and A. Hopper**, The Design and Implementation of a Raid-3 Multimedia File Server
- **M.M. Buddhikot and G. Parulkar**, Efficient Data Layout, Scheduling and Playout Control in a Multimedia Storage Server
- **S. Stoller and J. DeTreville**, Storage Replication and Layout in Video-on-Demand Servers
- **S. Paek, P. Bocheck and S. Chang**, Scalable MPEG2 Video Servers with Heterogeneous QoS on Parallel Disk Arrays
- **G. Neufeld, D. Makaroff and N. Hutchinson** , Design of a Variable Bit Rate Continuous Media Server

Welcome to the on-line proceedings of NOSSDAV'95!

Revised, and now available in print as Lecture Notes in Computer Science 1018 (Springer-Verlag)!

Advance Reservation Systems, Chair: Kevin Jeffay

- **Advance Reservations for Predicted Service**, M. Degemark , T. Kohler, S. Pink and O. Schelen
- **Distributed Advance Reservation of Real-Time Connections**, D. Ferrari, A. Gupta and G. Ventre
- **Issues of Reserving Resources in Advance**, L. Wolf, L. Delgrossi, R. Steinmetz, S. Schaller and H. Wittig

Operating System Support, Chair: Duane Northcutt

- **A Computational and Engineering View on Open Distributed Real-Time Multimedia Exchange**, P. Leydekkers, V. Gay and L. Franken
- **Support for User-Centric Modular Real-Time Resource Management in the Rialto Operating System**, M.B. Jones, P. Leach, R. Draves and J. Barrera, III
- **A Rate-Based Execution Abstraction for Multimedia Computing**, K. Jeffay and D. Bennett
- **VuSystem Performance Measurements**,C.J. Lindblad
- **Design of Universal Continuous Media I/O[*]**, C.D. Cranor and G.M. Parulkar
- **A New OS Architecture for High Performance Communication Over ATM Networks--Zero-Copy Architecture**, H. Kitamura, K. Taniguchi, H. Sakamoto and T. Nishida

Resource Management and Quality of Service, Chair: Aurel Lazar

- **Meeting Arbitrary QoS Constraints Using Dynamic Rate Shaping of Coded Digital Video**, A. Eleftheriadis and D. Anastassiou
- **Dynamic QoS Management for Scalable Video Flows**, A. Campbell, D. Hutchinson and C. Aurrecoechea
- **Dynamic Service Aggregation for Efficient Use of Resources in Interactive Video Delivery**,S. D. Venkatesh and T.D.C. Little
- **Evaluation of QoS-Control Servers on Real-Time Mach**, K. Kawachiya, M. Ogata, N. Nishio and H. Tokuda
- **System-Level Resource Management for Network-Based Multimedia Applications**, L. Schreier and M. Davis

NOSSDAV Retrospective, Chairs: Domenico Ferrari and Dan Swinehart

Audio and Video Systems, Chair: Raj Yavatkar

- **An End-to-End Software Only Scalable Video Delivery System**, N. Chaddha, G. Wall and B. Schmidt
- **A Distributed Real-Time MPEG Audio Player**, S. Cen, C. Pu, R. Staehli, C. Cowan and J. Walpole
- **Analysis of Audio Packet Loss over Packet-Switched Networks**, C. Bolot, H. Crepin and A. Garcia

- **Digital Audio and Video in Industrial Systems**, H. Lauer, C. Shen, R. Osborne, J. Howard and Q. Zheng, M, Takegaki, H. Shimakawa and I. Mizunuma
- **Workstation Video Playback Performance with Competitive Process Load**, K. Fall, J. Pasquale and S. McCanne

Scheduling and Synchronization, Chair: Chuck Kalmanek

- **An Adaptive Stream Synchronization Protocol**, K. Rothermel and T. Helbig
- **A Method and Apparatus for Measuring Media Synchronization**, B. Schmidt, J. Northcutt, and M. Lam
- **Integrated Processor Scheduling for Multimedia**, J. Nieh and M.S. Lam
- **Scheduling and Admission Testing for Jitter Constrained Periodic Threads**, A. Mauthe and G. Coulson
- **A CPU Scheduling Algorithm for Continuous Media Applications**, R. Yavatkar and K. Lakshman

Multicasting, Chair: Jim Kurose

- **Dynamic Configuration of Conferencing Applications Using Pattern-Matching Multicast**,H. Schulzrinne
- **Wave: A New Multicast Routing Algorithm for Static and Dynamic Multicast Groups**, E. Biersack and J. Nonnenmacher
- **Connection Establishment for Multi-Party Real-Time Communication**, R. Bettati,D. Ferrari, A. Gupta, W. Heffner, W. Howe, M. Moran, Q. Nguyen and R. Yavatkar
- **The Role of Multicast Communication in the Provision of Scalable and Interactive Video-on-Demand Service**, K.C. Almeroth and M.H. Ammar
- Invited Panelists

Network Scheduling and Real-Time Networking Chair: Steve Pink

- **A New Approach to Support VBR Video in Packet Switched Networks,**H. Zhang and E.W. Knightly
- **Determining End-to-End Delay Bounds in Heterogeneous Networks**, P. Goyal, S. Lam and H. Vin
- **Adaptive QoS-Based API for ATM Networking**, V. Bansal, R.J. Siracusa, J.P. Hearn, G. Ramamurthy and D. Raychaudhuri
- **Burst Scheduling Networks: Flow Specification and Performance Guarantees**, S.S. Lam and G.G. Xie
- **The Case Against Variable Bit Rate Service**, M. Grossglauser, S. Keshav and D. Tse

Storage Architectures, Chair: Jon Walpole

- **A Novel Video-on-demand Storage Architecture for Supporting Constant Frame Rate with Variable Bit Rate Retrieval**,S.W. Lau and J.C.S. Lui
- **The Design and Implementation of a Raid-3 Multimedia File Server**, A.J. Chaney and I.D. Wilson
- **Efficient Data Layout, Scheduling and Playout Control in a Multimedia Storage Server**, M.M. Buddhikot and G. Parulkar
- **Storage Replication and Layout in Video-on-Demand Servers**, S. Stoller and J. DeTreville
- **Scalable MPEG2 Video Servers with Heterogeneous QoS on Parallel Disk Arrays**, S. Paek, P. Bocheck and S. Chang
- **Design of a Variable Bit Rate Continuous Media Server**, G. Neufeld, D. Makaroff and N. Hutchinson

Program Co-Chairs:

Riccardo Gusella	HP Labs USA
T.D.C. Little	Boston University USA

Program Committee:

Stavros Christodoulakis	Technical University of Crete
Domenico Ferrari	University of California Berkeley
Ralf Herrtwich	IBM Creative Multimedia Studios
Andy Hopper	Olivetti & University of Cambridge
Kevin Jeffay	University of North Carolina at Chapel Hill
Chuck Kalmanek	AT&T Bell Laboratories
Jim Kurose	University of Massachusetts at Amherst
Aurel Lazar	Columbia University
Derek McAuley	Cambridge University
Duane Northcutt	Sun Microsystems Laboratories
Guru Parulkar	Washington University
Steve Pink	Swedish Institute of Computer Science
Radu Popescu-Zeletin	GMD-FOKUS
K.K. Ramakrishnan	AT&T Bell Laboratories
P. Venkat Rangan	University of California San Diego
Jon Rosenberg	C|NET
Eve Schooler	California Institute of Technology
Doug Shepherd	Lancaster University
Cormac Sreenan	AT&T Bell Laboratories
Jean-Bernard Stefani	France Telecom/CNET
Daniel Swinehart	Xerox PARC
Hideyuki Tokuda	Carnegie Mellon University/Keio University
Jon Walpole	Oregon Graduate Institute of Science & Technology
Raj Yavatkar	University of Kentucky

Additional Referees

Milind Buddikot	Washington University
Zubin Dittia	Washington University
Anand Krishnamurthy	Boston University
K. Lakshman	University of Kentucky
Christos Papadopoulos	Washington University
Gopal Raman	Washington University
Srinivas Ramanathan	HP Labs
Dinesh Venkatesh	Boston University

Return to the NOSSDAV page.
Last updated: November 13, 1995
Send Your Comments to Webmaster.

NOSSDAV'95 Paper Abstracts

- Kurt Rothermel and Tobias Helbig, **An Adaptive Stream Synchronization Protocol,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, pp. 189-202.

Abstract: Protocols for synchronizing data streams should be highly adaptive with regard to both changing network conditions as well as to individual user needs. The Adaptive Synchronization Protocol we are going to describe in this paper supports any type of distribution of the stream group to be synchronized. It incorporates buffer level control mechanisms allowing an immediate reaction on overflow or underflow situations. Moreover, the proposed mechanism is flexible enough to support a variety of synchronization policies and allows to switch them dynamically during presentation. Since control messages are only exchanged when the network conditions actually change, the message overhead of the protocol is very low.

Keywords: Communication protocols, stream synchronization, quality of service, buffer control algorithm.

- S.W. Lau and John C.S. Lui, **"A Novel Video-On-Demand Storage Architecture for Supporting Constant "Frame Rate with Variable Bit Rate Retrieval,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, 315-326.

Abstract: One of the quality of service (QOS) factors in video-on-demand (VOD) applications is to provide high resolution quality to end users. One way to achieve this is to provide a constant display frame rate (e.g., 30 frames/sec) at the display station. However, due to the nature of video files and compression technique applied, video frame sizes vary significantly from frame to frame. Therefore, although the display frame rate is fixed, data retrieval is a variable bit rate process. Conventional VOD storage servers assume a peak rate retrieval of video files. Therefore, the number of concurrent requests to the VOD server cannot be maximized. In this paper, we consider a VOD storage the number of concurrent requests to the VOD server cannot be maximized. In this paper, we consider a VOD storage server which can support a fixed frame rate at the display and at the same time, variable bit rates retrieval of compressed video files. We describe (1) video files layout strategy, (2) request scheduling algorithm, (3) buffering issues and, (4) various VCR features support such that the number of concurrent requests can be maximized.

Keywords: Multimedia storage architecture, scheduling, variable bit rate retrieval.

- G. Neufeld, D. Makaroff and N. Hutchinson, **"The Design of a Variable Bit Rate Continuous Media Server,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, pp. 375-378.

Abstract: This paper describes the design and implementation of a file server for variable bit rate continuous media. Most continuous media file servers have been designed for constant bit rate streams. We address the problem of building a server where each stream may have a different bit rate and, more importantly, where the bit rate within a single stream may vary considerably. Such servers will be come increasingly more important because of Variable Bit Rate (VBR) compression standards such as MPEG-2.

Keywords:

- S. Lam and G. Xie, **"Burst Scheduling Networks: Flow Specification and Performance Guarantees,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, pp. 303-306.

Abstract: We present a class of packet switching networks, called Burst Scheduling Networks, designed to provide throughput, delay, and delay jitter guarantees. These performance guarantees are derived from the delay guarantee of a VC server, and a new traffic model called Flow Specification. The delay guarantee of a VC server has several desirable properties, including the following firewall property: The guarantee to a flow is unaffected by the behavior of other flows sharing the same server. There is no assumption that sources are flow-controlled or well-behaved. Each guaranteed flow is modeled as a sequence of bursts, each of which is a sequence of packets. Bursts are needed to specify two types of jitter bounds: over the delays of packets in a burst, and over the delays of bursts in a flow. For video flows, each encoded picture is naturally modeled by a burst. The model is also appropriate for audio and data flows that require delay and delay jitter guarantees. With the new traffic model, a flow can be partitioned into intervals (bursts) that have substantially different average rates; the first packet of a burst carries information on the size and average rate of the burst. Switches are designed to process flows efficiently in bursts.

Keywords: Packet switching, delay guarantee, delay jitter guarantee, throughput guarantee, firewall property, virtual clock, Burst Scheduling, video.

- K. Almeroth and M. Ammar, **"The Role of Multicast Communication in the Provision of Scalable and Interactive Video-On-Demand Service,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, pp. 267-270.

Abstract: Multicast delivery can improve the scalability of Video-On-Demand (VOD) systems by allowing multiple customers to share one set of video server and network resources. In such a system, providing customer interactivity can be difficult. In this paper we overview our work which addresses how interactivity may be accomplished in a multicast delivery environment and analyzes the performance of such systems.

Keywords: Video-On-Demand, video, multicast, interactive, scalable

- E. Biersack and J. Nonnenmacher, **"WAVE: A New Multicast Routing Algorithm for Static and Dynamic Multicast Groups,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, pp. 243-254.

Abstract: We present a new multicast algorithm called WAVE for establishing source-specific multicast trees. WAVE meets multiple quality of service requirements (constraints) such as delay, cost, and available bandwidth, simultaneously. Simulation results show that WAVE performs very good in terms of delay and cost for both, static and dynamic multi- cast groups, when compared to the best multicast algorithms known.

Keywords:

- L. Wolf, L. Delgrossi, R. Steinmetz, S. Schaller and H. Wittig, **"Issues of Reserving Resources in Advance,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, pp. 27-37.

Abstract: Resource management offers Quality-of-Service reliability for time-critical continuous-media applications. Currently existing resource management systems provide only means to reserve resources starting with the reservation

attempt and lasting for an unspecified duration. However, for several applications such as video conferencing the ability to reserve the required resources in advance is needed. This paper explains a model for resource reservation in advance. We identify and discuss issues which must be resolved in resource reservation in advance systems. Some of the possible scenarios to be considered are described and we show how the resource reservation in advance scheme can be embedded in a general architecture.

Keywords:

- A. Mauthe and G. Coulson, **"Scheduling and Admission Testing for Jitter Constrained Periodic Threads,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, pp. 219-222.

Abstract: We present new admission tests for periodic real-time threads with explicitly stated deadlines. In standard earliest-deadline-first scheduling, the deadline of a periodic thread is conventionally the end of the current period. In contrast, our tests supports periodic threads in which the deadline may be earlier than the end of the current period. In the extreme case, the deadline may be specified as identical to the execution time, which results in perfectly isochronous periodic threads. The provision of such threads, which we refer to as jitter constrained threads, helps end-systems to honour jitter- as well as throughput-related QoS parameters in distributed multimedia systems. In addition, such threads can reduce end-to-end delay and buffer memory requirements as less buffering is needed to smooth excessive delay jitter. Our admission tests are primarily designed around the earliest deadline first (EDF) scheduling algorithm but the tests can also be applied to priority based scheduling schemes, such as rate monotonic, with minimal modification.

Keywords: Admission Testing, CPU-Scheduling

- K. Fall, J. Pasquale and S. McCanne, **"Workstation Video Playback Performance with Competitive Process Load,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, pp. 179-182.

Abstract: While many researchers believe that multimedia applications are best managed with hard, real-time scheduling mechanisms, models based on application-level adaptation with relaxed scheduling constraints are gaining acceptance. We analyze an existing video conferencing application, which was designed without explicit support for CPU resource management, and propose modifications to its architecture to support CPU load adaptation. Inter-frame display times (IDT) for a 30 frames/sec video segment are analyzed for increasing multiprocessing loads. As expected, the IDT variance (and mean) increases markedly with load, especially beyond saturation. We show that this display jitter can be significantly reduced by gracefully adapting the application's load requirements to match the available CPU resources.

Keywords:

- V. Bansal, R.J. Siracusa, J.P. Hearn, G. Ramamurthy and D. Raycaudhuri, **"Adaptive QoS-Based API for ATM Networking,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, pp. 299-302.

Abstract: This paper describes work-in-progress on a new adaptive QoS-based application programming interface (API) for ATM networking. This ATM API, referred to as the "ATM Service Manager (ASM)" is motivated by the observation that current transport interfaces do not provide QoS features necessary for multimedia applications to achieve desired performance/cost objectives on multiservice ATM networks. In general, an application connected to ATM has a choice of transport protocols (e.g. TCP, UDP, alternative multimedia stream protocols), ATM service class

(e.g. ABR, VBR, CBR) and QoS parameters (e.g. peak/average bandwidth, delay jitter, cell loss) for each media type. It is our view that application software should be shielded from the complexity of QoS-based service management by an adaptive API which is responsible for mapping the application requirements to ATM (and vice versa).

Keywords: ATM networks, QoS, API, Multimedia network interfaces

---

- M. Grossglauser, S. Keshav and D. Tse, **"The Case Against Variable Bit Rate Service,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, pp. 307-310.

Abstract:

Keywords:

- K. Kawachiya, M. Ogata, N. Nishio and Hideyuki Tokuda, **"Evaluation of QOS-Control Servers on Real-Time Mach,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, pp. 123-126.

Abstract: The Keio-MMP (MultiMedia Platform) project has as its aim the development of a common software platform for distributed multimedia computing. Its base software is Real-Time Mach 3.0 (RT-Mach). We are extending RT-Mach and constructing servers for multimedia processing on our extension. One goal of our project is to incorporate the concept of Quality of Service (QOS), and to manage system resources over multiple continuous-media sessions on the basis of this concept. However, it is difficult for multimedia applications to manage their QOS appropriately by themselves without system support. There should be some framework for QOS management at the system-software level. To provide such support, we adopted a QOS Manager-based scheme, and developed an experimental QOS-Control Server on RT-Mach. This paper gives the results of several experiments with the QOS-Control Server and describes the extensions of RT-Mach needed for efficient QOS control. The thread-model extension makes it possible to create a thread suitable for continuous-media processing, and the CPU-reservation mechanism gives more efficient QOS control. Based on these experiments, we are now designing a ``QOS-Ticket'' model as a new QOS-control scheme.

Keywords: QOS control, resource management, operating system, multimedia, continuous media

---

- H. Kitamura, K. Taniguchi, H. Sakamoto and T. Nishidam, **"A New OS Architecture for High Performance Communication over ATM Networks -- Zero-copy architecture -- ,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995. pp. 87-90.

Abstract: A new OS architecture, referred to as Zero-copy architecture, for high performance communication is proposed. This architecture dissolved memory copy bottleneck that is a major overhead in protocol processing. This can reduce CPU processing overhead, in addition to realizing high speed data communication. This architecture is shown to be suitable for large volume of data communication like video and image transfer.

Keywords: Communication Protocol, Performance Evaluation, ATM, Multimedia Communication, Workstation, UNIX

- C. Lindblad, **"VuSystem Performance Measurements,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, pp. 79-82.

Abstract: In this paper I discuss some performance measurements made of the VuSystem, a programming system for

the software-based processing of audio and video data. The VuSystem is designed to run on ordinary Unix workstations with no specific support for the manipulation of multimedia data. Measurements made of processing times of representative filter modules demonstrate the viability of the approach.

Keywords:

- S. Cen, C. Pu, R. Staehli, C. Cowan and J. Walpole, **" A Distributed Real-Time MPEG Video Audio Player,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, pp. 151-162.

Abstract: This paper presents the design, implementation and experimental analysis of a distributed, real-time MPEG video and audio player. The player is designed for use across the Internet, a shared environment with variable traffic and with great diversity in network bandwidth and host processing speed. We use a novel toolkit approach to build software feedback mechanisms for client/server synchronization, dynamic Quality-of-Service control, and system adaptiveness. Our experimental results show that the feedback mechanisms are effective, and that the player performs very well in the Internet environment.

Keywords: multimedia, real-time system, software feedback, synchronization, Quality-of-Service control, adaptiveness

---

- M. Degermark, T. Kohler, S. Pink and O. Schelen, **"Advance Reservations for Predictive Service,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, pp. 3-14.

Abstract: We extend a measurement-based admission control algorithm suggested for predictive service to provide advance reservations for guaranteed and predictive service while keeping the attractive features of predictive service. The admission decision for advance reservations is based on information about flows that overlap in time. For flows that have not yet started, the requested values are used, and for those that have already started measurements are used. This allows us to estimate the network load accurately for the near future. To provide advance reservations we ask users to include durations in their requests. We provide simulation results to show that predictive service with advance reservations provides utilization levels significantly higher than those for guaranteed service.

Keywords: Predictive Service, Admission Control, Advance Reservations, Quality-of-Service, Networking

---

- M. Jones, P. Leach, R. Draves and J. Barrera, III, **"Support for User-Centric Modular Real-Time Resource Management in the Rialto Operating System,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, pp. 55-66.

Abstract: This paper describes ongoing investigations into algorithms for user-centric modular distributed real-time resource management. These investigations are being conducted in the context of the Rialto operating system -- an object-based real-time kernel and programming environment currently being developed within Microsoft Research. A primary goal of this research is to develop appropriate real-time programming abstractions to allow multiple independent real-time programs to dynamically coexist and share resources on the same hardware platforms. Use of these abstractions is intended both to allow individual applications to reason about their own resource requirements and for per-machine system resource planner applications to reason about and control resource allocations between potentially competing applications. The set of resources being managed is dynamically extensible, and may include remote resources in distributed environments. The local planner conducts resource negotiations with individual applications on behalf of the user, with the goal of maximizing the user's perceived utility of the set of running applications with respect to resource allocations for those applications.

Keywords:

- H. Zhang and E. Knightly, **" A New Approach to Support Delay-Sensitive VBR Video in Packet-Switched Networks,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, pp. 275-286.

Abstract: Previous approaches to supporting video on packet-switched networks include deterministic service, statistical service, predicted service, and feedback-based schemes. These schemes represent different tradeoffs in quality of service (QOS), achievable network utilization, and method of dealing with overload. In this paper, we propose a new service that attempts to strike an efficient balance with the above tradeoffs. The approach is based on deterministic guarantees with client controlled renegotiation of QOS parameters and graceful adaptation during overload periods. We evaluate the scheme using two traces of MPEG-compressed video and show that, even with simple renegotiation polices and relatively low renegotiation frequencies, high network utilization in the range of 50% to 80% can be achieved. For traffic that is bursty over long intervals, this represents a 100% to 150% improvement in network utilization compared to deterministic service. Compared to statistical and predicted service, our approach allows more graceful and client-controlled QOS degradation during overload period.

Keywords: VBR video, guaranteed QOS, resource management, multimedia communication

---

- 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. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, pp. 255-266.

Abstract: There is considerable interest in the network community in supporting real-time multi-party applications, such as video conferencing. The Tenet Group at UC Berkeley and ICSI has designed and implemented protocols that provide quality of service (QoS) guarantees for real- time traffic on packet switching networks. Suite 2 of the Tenet protocols provides scalable, flexible and efficient network support for real-time multi-party connections. We outline our method of connection establishment and describe the design issues and alter- natives, and our decisions. Preliminary measurements confirm the viability of our approach for real-time multicast connection establishment.

Keywords: Multicast connection establishment; real-time communication; multimedia networking

---

- J. Bolot, H. Crepin and A. Garcia, **" Analysis of Audio Packet Loss in the Internet,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, pp. 163-174.

Abstract: We consider the problem of distributing audio data over networks such as the Internet that do not provide support for real-time applications. Experiments with such networks indicate that audio quality is mediocre in large part because of excessive audio packet losses. In this paper, we show using measurements over the Internet as well as analytic modeling that the number of consecutively lost audio packets is small unless the network load is very high. This indicates that open loop error control mechanisms based on forward error correction would be adequate to reconstruct most lost audio packets.

Keywords: Packet audio, loss process, Internet, forward error correction

---

- N. Chaddha, G. Wall and B. Schmidt, **"An End to End Software Only Scalable Video Delivery System,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire,

April 18-21, 1995, pp. 139-150.

Abstract: Precompressed video delivery systems commonly operate at fixed data rates. However, variations in the availability of network bandwidth and processor cycles are common in dynamic general purpose computing environments. Variability arises from the outright lack of resources (e.g. network bandwidth and cpu cycles), contention for available resources due to congestion, or a user's unwillingness to allocate needed resources to the task. Users of a scalable video delivery system have greater flexibility and therefore, the system can more effectively deliver video in the presence of system resource scarcity. This paper describes an end-to-end system combining a new scalable video compression algorithm, video delivery software, a software video decoder, and a market-based mechanism for the resolution of conflicts in providing video to the user.

Keywords: Scalable Video, Software Only, End-to-End Video Delivery System

---

- C. Cranor and G. Parulkar, **"Design of Universal Continuous Media I/O[*],"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, pp. 83-86.

Abstract: The problem this research addresses is how to modify an existing operating system's I/O subsystem to support new high-speed networks and high-bandwidth multimedia applications that will play an important role in future computing environments. The proposed I/O subsystem is called universal continuous media I/O (UCM I/O).

Keywords: I/O, networking, operating systems, multimedia

- D. Ferrari, A. Gupta and G. Ventre, **"Distributed advance reservation of real-time connections,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, pp. 15-26.

Abstract: The ability to reserve real-time connections in advance is essential in all distributed multi-party applications (i.e., applications involving multiple human beings) using a network that controls admissions to provide good quality of service. This paper discusses the requirements of the clients of an advance reservation service, and a distributed design for such a service. The design is described within the context of the Tenet Real-Time Protocol Suite 2, a suite being developed for multi-party communication, which will offer advance reservation capabilities to its clients based on the principles and the mechanisms proposed in the paper. Some simulation results about the performance of these mechanisms are also presented.

Keywords: Resource management/reservation, admission control, communication protocols, advance reservations

---

- P. Goyal, S. Lam and H. Vin, **"Determining End-to-End Delay Bounds In Heterogeneous Networks,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, pp. 287-298.

Abstract: We define a class of Guaranteed Rate (GR) scheduling algorithms. The GR class includes Virtual Clock, Packet-by-Packet Generalized Processor Sharing and Self Clocked Fair Queuing. For networks that employ scheduling algorithms belonging to GR, we present a method for determining an upper bound on end-to-end delay. The method facilitates determination of end-to-end delay bounds for a variety of sources. We illustrate the method by determining end-to-end delay bounds for sources conforming to Leaky Bucket and Exponentially Bounded Burstiness.

Keywords:

- K. Jeffay and D. Bennett, **" A Rate-Based Execution Abstraction For Multimedia Computing,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, pp. 67-78.

Abstract: Process models for multimedia computing must allow applications to adapt their pattern of execution as resources become scarce or abundant. As processes adapt, they should be able to express their desired performance in terms of application-defined data units or events. We propose a process model wherein processes execute according to a general rate specification of x process executions every y time units. In addition, a separate parameter is used to specify the desired response time for the completion of each execution. In all cases the real-time performance of a rate-based process is predictable. The model is general enough to encompass or extend many of the existing models proposed for multimedia systems. Our model of rate-based execution is described along with an implementation that detects when processes should adapt their execution rate and minimizes latency in interprocess communication.

Keywords:

- H. Lauer, C. Shen, R. Osborne, J. Howard and Q. Zheng, M, Takegaki, H. Shimakawa and I. Mizunuma, **"Digital Audio and Video in Industrial Systems,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, pp. 175-178.

Abstract: This is a position paper discussing the requirements of networks in industrial environments, especially with respect to digital audio and video. Topics include resource allocation, issues surrounding switching, scheduling, and priorities, end system interface requirements, and traffic characteristics and their implications on flow control.

Keywords:

- P. Leydekkers, V. Gay and L. Franken, **"A Computational and Engineering View on Open Distributed Real-time MultiMedia Exchange,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, pp. 43-54.

Abstract: An important requirement for distributed multimedia applications is the support of real-time communication and the means to specify real-time aspects. The aim of this paper is to extend RM-ODP and TINA-C computational and engineering views on distributed systems for the specification and support of real-time communication. It is expected that these bodies have a major impact in the area of distributed processing. However, concepts and mechanisms to support real-time communication are not yet fully included or detailed in these standards. In particular this paper addresses Quality of Service (QoS) specifications for continuous dataflows. These QoS specifications are described from the ODP computational and engineering viewpoint and the repercussions of these QoS specifications for functions located in both the computing and telecommunications environment are discussed.

Keywords: Distributed Multimedia Architectures, QoS, ODP, TINA-DPE

- J. Nieh and M. Lam, **"Integrated Processor Scheduling for Multimedia,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, 215-218.

Abstract: The advent of multimedia ushers forth a growing class of applications that must manipulate digital audio and video within well-defined timeliness requirements. Existing processor schedulers are inadequate in supporting these requirements. They fail to allow the integration of these continuous media computations with conventional interactive and batch activities. By observing that all tasks do in fact have timeliness requirements, we have created a

new scheduler that provides integrated processor allocation for all classes of computational activities. By exploiting information on the relative importance of different activities to the user, the scheduler seeks to maximize the value the system delivers to the user, and gracefully degrades activities in order of importance when all timeliness requirements can not be met. The solution is unique in the degree to which it allows users control over the allocation of processing resources, and in that it provides an integrated scheduling solution in supporting multimedia applications. We have implemented the scheduler in the Solaris operating system to demonstrate its effectiveness for real applications on real systems.

Keywords: processor scheduling, operating systems, multimedia, time constraints, overload management

---

- L. Schreier and M. Davis, **" System-Level Resource Management for Network-Based Multimedia Applications,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, 127-130.

Abstract: SRI International (SRI) has developed a model for system-level resource management in distributed systems, and applied this model in the context of a multimedia conferencing application. The model considers user objectives, resource constraints, and adaptable execution techniques. User objectives are specified by means of benefit functions. We have implemented a prototype of a distributed multimedia display application that demonstrates key aspects of the model, including adaptation to a changing execution environment.

Keywords: resource management, adaptable systems, quality of service, benefit functions

---

- H. Schulzrinne, **"When can we unplug the radio and telephone?"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, pp. 183-184.

Abstract: Despite high-performance workstations, special-purpose communication devices like telephones and radio still dominate. Reasons include inadequate system and network support, both in terms of performance and the ability to write applications without specialized signal processing knowledge. The interaction of applications with the network will depend largely on the future tariff structures rather than protocol issues.

Keywords: audio, video, research issues, ATM, multimedia, programming abstractions

---

- H. Schulzrinne, **"Dynamic Configuration of Conferencing Applications using Pattern-Matching Multicast,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, pp. 231-242.

Abstract: Multimedia conferencing systems are usually large, complex software systems. We describe a local control architecture and communication protocols that allow to tie together media agents, controllers and auxiliary applications such as media recorders and management proxies into a single conference application. Unlike other systems, control of a single conference can be shared between several controllers. Each media can be handled by one or more independent media agents. Parts of the system have been implemented using an IP-multicast-based audio conferencing tool NeVoT. The communicating applications disseminate state and control information through a replicator. The replicator mainly limits distribution of messages based on expressed interest of other applications, thus implementing an application-level, receiver-driven local multicast. It also automatically starts applications as needed.

Keywords: conference control, multicast, RPC, CSCW

- S. Paek, P. Bocheck and S. Chang, **"Scalable MPEG2 Video Servers with Heterogeneous QoS on Parallel Disk Arrays,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, 363-374.

Abstract: In this paper we focus on the video storage unit of a video server. We present a new, flexible data placement strategy for independent parallel disk arrays. The trade-off between utilization efficiency and interactive delay is investigated for this data placement strategy. Based on this trade-off, we show the advantage of video servers supporting a range of interactivity QoS. For our data placement strategy, we show that using scalable video improves the utilization and interactivity performance of a video server. We use three-layer, scalable MPEG2 digital video to support resolution QoS at a video server. Finally, we show that the data placement strategy reduces the complexity of admission control at the video server to that of a single disk system.

Keywords:

- S. Stoller and J. DeTreville, **"Storage Replication and Layout in Video-on-Demand Servers,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, 351-363.

Abstract: We propose and analyze an architecture for storage servers in large Video on Demand (VoD) systems. We describe a method for distributing the collection of titles among the levels of the storage hierarchy, based on estimates of the mean demand for each title. The resulting distribution minimizes cost for a given level of performance. Since high availability is desirable in VoD systems, we consider the use of mirroring or parity-based redundancy (a la RAID) and estimate the effect on the system's cost and availability. In the very-large-scale storage systems needed for VoD, the placement of disk arrays on the pool of computers must be chosen carefully to provide high availability for the least cost. We propose a strategy for arranging disk arrays on a pool of computers; our strategy is inspired by Holland and Gibson's work on parity declustering for RAID. To estimate the availability benefit of redundancy, we developed an analytical availability model that takes into account the system's tolerance to all kinds of failures (e.g., software bugs as well as disk crashes).

Keywords: video-on-demand, RAID, availability, layout, storage

---

- A. Chaney, I. Wilson and A. Hopper, **"The Design and Implementation of a RAID-3 Multimedia File Server,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, pp. 327-338.

Abstract: The Olivetti Research Laboratory has developed an experimental system based on intelligent peripherals connected directly to an ATM network. As well as multimedia modules (e.g. audio and video) the system also includes a directly connected RAID-3 storage server called the "Disc Brick". This paper describes the architecture of the Disc Brick, and discusses some of the hardware and software issues raised by its design. It also presents measurements taken from a Disc Brick in operation, and discusses how the observations relate to the original design objectives. Finally, the paper attempts to evaluate the Disc Brick as part of ORL's family of directly connected peripherals.

Keywords:

- R. Yavatkar and K. Lakshman, **"A CPU Scheduling Algorithm for Continuous Media Applications,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995. pp. 223-226.

Abstract: We address the problem of designing a CPU scheduling algorithm that can guarantee a steady, average rate

of execution to multimedia (MM) applications. In particular, our goal is to provide dynamic QOS control in the presence of varying compute times, arrival and departure of processes, and CPU overloads. We have designed and evaluated a new CPU management algorithm called Rate-based Adjustable Priority Scheduling (RAP). that makes the following new contributions. First, we do not assume a priori knowledge of compute times for the MM applications. RAP includes an admission control algorithm that estimates and allocates available capacity to new processes. Second, we consider only average computing time needed by a process and admission control algorithm provides guarantees on average rate of execution. Third, our algorithm includes a rate adaptation mechanism and an application-level PLL (Phase Locked Loop) to adapt to CPU overload and variations in compute times.

Keywords:

- M.M. Buddhikot and G. Parulkar, **" Efficient Data Layout, Scheduling and Playout Control in MARS,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, pp. 339-350.

Abstract: Large scale on-demand multimedia servers, that can provide independent and interactive access to a vast amount of multimedia information to a large number of concurrent clients, will be required for a wide spread deployment of exciting multimedia applications. Our project, called Massively-parallel And Real-time Storage (MARS), is aimed at prototype development of such a large scale server. This paper primarily focuses on the distributed data layout and scheduling techniques developed in this project. These techniques support a high degree of parallelism and concurrency, and efficiently implement various playout control operations, such as fast forward, rewind, pause, resume, frame advance and random access.

Keywords: Large scale multimedia storage servers, Interactive playout control, Desk Area Networks, Storage systems

---

- A. Campbell, D. Hutchinson and C. Aurrecoechea, **"Dynamic QoS Management for Scalable Video Flows,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, pp. 107-118.

Abstract: We introduce the concept of Dynamic QoS Management (DQM) for control and management of hierarchically coded flows operating in heterogeneous multimedia networking environments. The motivation that underpins our scheme is to bridge the heterogeneity gap that exists between applications, end- systems and networks. QoS adaptors, QoS filters and QoS groups are key scalable objects used in resolving quality of service capability mismatch. QoS filters manipulate hierarchically coded flows as they progress through the communications system, QoS adaptors scale flows at the end-systems based on the flow's measured performance and user supplied QoS scaling policy, and QoS groups provide baseline quality of service for multicast communication.The focus of the work is driven by a) the special features of scalable video flows-in particular MPEG2, b) the needs of both scalable and single-layer video for transmission over multimedia networks such as ATM. A novel adaptive network service is proposed for the transmission of multi-layer coded flows that offers "hard" guarantees to the base layer, and "fairness" guarantees to the enhancement layers based on a new bandwidth allocation technique called Weighted Fair Sharing (WFS).

Keywords: MPEG-2, Adaptive Service, Scalable Flows, Dynamic QoS Management, End-to-end QoS

- D. Venkatesh and T.D.C. Little, **"Dynamic Service Aggregation for Efficient Use of Resources in Interactive Video Delivery,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, pp. 119-122.

Abstract: To support future interactive information delivery services there is a need to balance individual interactivity with the desire to maximize the number of supported sessions. Currently, few techniques have demonstrated the ability

to renegotiate and scale service parameters per session in progress as required to adapt to differing terminal equipment characteristics and network congestion. This paper addresses this problem through the definition of decomposable service groups that permit aggregation of interactivity, terminal characteristics, and levels of service scaling. The proposed approach applies the characteristics of end applications and data storage requirements to the design of a data scaling mechanism.

Keywords: interactive multimedia services, service scaling, multicast, dynamic service aggregation, video-on-demand

---

- Alexandros Eleftheriadis and Dimitris Anastassiou, **" Meeting Arbitrary QoS Constraints Using Dynamic Rate Shaping of Coded Digital Video,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, pp. 95-106.

Abstract: We introduce the concept of {\em Dynamic Rate Shaping}, a technique to adapt the rate of compressed video bitstreams (MPEG-1, MPEG-2, H.261, as well as JPEG) to dynamically varying rate (and delay) constraints. The approach provides an interface (or filter) between the encoder and the network, with which the encoder's output can be perfectly matched to the network's quality of service characteristics. Since the presented algorithms do not require interaction with the encoder, they are fully applicable to precoded, stored video (as in, for example, video-on-demand systems). By providing decoupling of the encoder and the network, universal interoperability can be achieved. A set of low-complexity algorithms for dynamic rate shaping is presented, and both optimal and extremely fast designs are discussed. The latter are simple enough to allow software-based implementation. Experimental results are provided using actual MPEG-2 bitstreams.

Keywords:

- B. Schmidt, J. Northcutt, and M. Lam **"A Method and Apparatus for Measuring Media Synchronization,"** *Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video*, Durham, New Hampshire, April 18-21, 1995, pp. 203-214.

Abstract: Media synchronization is widely regarded as a fundamental problem in the field of multimedia. While much work has been conducted in this area, and many different solutions have been proposed, no method for obtaining a repeatable, objective measure of synchronization performance exists. Thus, there has been no means for determining the effectiveness of potential media synchronization solutions. In this paper we present an experimental methodology for quantitatively measuring the performance of different media synchronization schemes. We describe a complete (hardware and software) test environment for measuring audio/video synchronization quality of various media players, and we also present empirical performance measurements of an example media player. The results show that external observation is necessary for accurate assessments of synchronization performance. This test and evaluation methodology is applicable to other media delivery systems and can serve as the first step in isolating and quantifying the effects of individual components of a media delivery system.

Keywords: Multimedia, Synchronization, Measurement, Clocks

NOSSDAV Participants

Page One

Almeroth	Ammar	Anirudhan	Bansal
Bettatti	Bettatti	Biesack	Bocheck
Bolot	Buddhikot	Campbell	Buddhikot
Campbell	Cen	Chang	Degermark

Return to the index.

Last updated: November 13, 1995
Send Your Comments to Webmaster.

NOSSDAV Participants

Page Two

Delgrossi	Eleftheriadis	Escobar	Fall
Ferrari	Gay	Goodall	Gopal
Goyal	Gupta	Heffner	Hendrix
Herrtwich	Hsieh	Jeffay	Jones
Kalmanek	Kawachiya	Keshav	Kitamura

Return to the index.

Last updated: November 13, 1995
Send Your Comments to Webmaster.

NOSSDAV Participants

Page Three

Knightly	Kurose	Lackshman	Lauer
Lazar	Leslie	Leydekkers	Lindblad
Little	Lougher	Lui	M. Lam
Mao	Mauthe	McAuley	Meyer
Moran	Nakajima	Chaddah	Neih

Return to the index.

Last updated: November 13, 1995
Send Your Comments to Webmaster.

NOSSDAV Participants

Page Four

Neufeld	Northcut	Ott	Paek
Pink	Rothermel	S. Lam	Sandstaa
Schmidt	Schulzrinne	Smith	Steinmetz
Stoller	Swinehart	Tokuda	Vin
Vogels	Walpole	Weilbach	Wilson
Wolf	Yavatkar	Zhang	

Return to the index.

Last updated: November 13, 1995
Send Your Comments to Webmaster.