PROCEEDINGS

ACM Multimedia '95

San Francisco, California, November 5-9, 1995

acm PRESS

Sponsored by the ACM SIG MULTIMEDIA, SIGCHI, SIGGRAPH, SIGMIS, SIGBIO, SIGCOMM, SIGIR AND SIGOIS, in cooperation with ACM SIGAPP, SIGCAPH, SIGMOD and SIGOPS

The Association for Computing Machinery, Inc.
1515 Broadway
New York, NY 10036

...In proceedings of the ACM Multimedia, 1995 (San Francisco, CA, USA, November 5-9, 1995) ACM, New York, 1995, pp. 23-35.

Ordering Information

Nonmembers
Nonmember orders placed within the U.S. should be directed to:

Addison-Wesley Publishing Company
Order Department
Jacob Way
Reading, MA 01867
Tel: +1 800 447 2226

Addison-Wesley will pay postage and handling on orders accompanied by check. Credit card orders may be placed by mail or by calling the Addison-Wesley Order Department at the number above. Follow-up inquiries should be directed to the Customer Service Department at the same number. Please include the Addison-Wesley ISBN number with your order:

A-W ISBN 0-201-87774-0

Nonmember orders from outside the U.S. should be addressed as noted below:

Europe/Middle East

Addison-Wesley Publishing Group
Concertgebouwplein 25
1071 LM Amsterdam
The Netherlands
Tel: +31 20 6717296
Fax: +31 20 6645334

Germany/Austria/Switzerland

Addison-Wesley Verlag Deutschland GmbH
Hildachstrasse 15d
Wachsbleiche 7-12
53111 Bonn
Germany
Tel: +49 228 98 515 0
Fax: +49 228 98 515 99

United Kingdom/Africa

Addison-Wesley Publishers Ltd.
Finchampstead Road
Wokingham, Berkshire RG11 2NZ
United Kingdom
Tel: +44 734 794000
Fax: +44 734 794035

After January 1, 1996:
Addison-Wesley Longman Publishers, Ltd.
Longman House
Burnt Mill
Harlow, Essex CM202JE
United Kingdom
Tel: +44 1279 623 623
Fax: +44 1279 431 059

Asia

Addison-Wesley Singapore Pte. Ltd.
15 Beach Road
#05-02/09/10 Beach Centre
Singapore 0718
Tel: +65 339 7503
Fax: +65 339 9709

Japan

Addison-Wesley Publishers Japan Ltd.
Nichibo Building
1-2-2 Sarugakucho
Chiyoda-ku, Tokyo 101
Japan
Tel: +81 33 291 4581
Fax: +81 33 291 4592

Australia/New Zealand

Addison-Wesley Publishers Pty. Ltd.
6 Byfield Street
North Ryde, N.S.W. 2113
Australia
Tel: +61 2 878 5411
Fax: +61 2 878 5830

Latin America

Addison-Wesley Iberoamericana S.A.
Boulevard de las Cataratas #3
Colonia Jardines del Pedregal
Delegacion Alvaro Obregon
01900 Mexico D. F.
Tel: +52 5 660 2695
Fax: +52 5 660 4930

Canada

Addison-Wesley Publishing (Canada) Ltd.
26 Prince Andrew Place
Don Mills, Ontario M3C 2T8
Canada
Tel: +416 447 5101
Fax: +416 443 0948

ACM Members
A limited number of copies are available at the ACM member discount. Send order with payment in US dollars to:

ACM Order Department
P.O. Box 12114
Church Street Station
New York, N.Y. 10257

OR

ACM European Service Center
Avenue Marcel Thiry 204
1200 Brussels
Belgium

Credit card orders from U.S.A. and Canada:
+1 800 342 6626

New York Metropolitan Area and outside of the U.S.:
+1 212 626 0500
Fax +1 212 944-1318
Email:acmpubs@acm.org

Please include your ACM member number and the ACM Order number with your order.

ACM Order Number: 433952
ACM ISBN: 0-89791-751-0

An Application Level Video Gateway

Elan Amir
University of California, Berkeley
elan@cs.berkeley.edu

Steven McCanne
University of California, Berkeley
mccanne@cs.berkeley.edu

Hui Zhang
Carnegie Mellon University
hzhang@cs.cmu.edu

ABSTRACT

The current model for multicast transmission of video over the Internet assumes that a fixed average bandwidth is uniformly present throughout the network. Consequently, sources limit their transmission rates to accommodate the lowest bandwidth links, even though high-bandwidth connectivity might be available to many of the participants. We propose an architecture where a video transmission can be decomposed into multiple sessions with different bandwidth requirements using an application-level gateway. Our *video gateway* transparently connects pairs of sessions into a single logical conference by manipulating the data and control information of the video streams. In particular, the gateway performs bandwidth adaptation through transcoding and rate-control. We describe an efficient algorithm for transcoding Motion-JPEG to H.261 that runs in real-time on standard workstations. By making the Real-time Transport Protocol (RTP) an integral component of our architecture, the video gateway interoperates with the current Internet video tools in a transparent fashion. We have built a prototype of the video gateway and used it to redistribute multi-megabit JPEG video seminars from the Bay Area Gigabit Network as 128 kb/s H.261 MBone sessions.

KEYWORDS

Conferencing protocols; digital video; image and video compression and processing; multicasting; networking and communication; efficient transcoding.

1 INTRODUCTION

The initial deployment of a number of multimedia applications such as nevot [13], vat [7], ivs [16], nv [5], and vic [9] over the Internet and the overlayed multicast backbone or MBone [4] has proven successful. However, a number of issues need to be resolved before the full potential of these new applications can be realized. One such issue is how to deal with heterogeneity in a conference with a large number of receivers.

Heterogeneity is intrinsic to the Internet technology. It exists both inside the network and at the end systems. Within the network, due to the mismatch between link speeds and unbalanced load distribution inside the network, the bandwidth available between different sender-receiver pairs is different. While one portion of the network may be a high speed ATM network, other portions may be ethernet, and the backbone may be composed of T1 and T3 links. In addition, different hosts have different processing power and video/audio hardware configurations. One host may be a 100 MIPS workstation without special video hardware, another host may be a PC with a JPEG board. As the Internet evolves to higher speed and larger size, the problems caused by heterogeneity will only get worse.

In existing applications, each sender transmits video at the same rate to all receivers. For a heterogeneous set of receivers, the source would have to run at a rate that matches the most constrained receiver. This is not satisfactory. Instead, receivers along high bandwidth paths should receive correspondingly higher quality.

One approach to this problem is the use of layered coding algorithms [10, 15, 17] where hosts along lower bandwidth paths receive fewer layers. While this approach is elegant and efficient, it is not yet deployed nor compatible with the installed base of JPEG hardware and existing video applications. Our goal is to create an architecture that is compatible with prevalent hardware and networking technology.

In this paper, we present the design and implementation of a video gateway that addresses the problem of heterogeneity. Our approach provides a mechanism for matching the transmission quality to the heterogeneous bandwidth constraints of distinct regions of a single logical multicast session. The gateway does so by intelligently managing incoming and outgoing video streams using transcoding and rate-control. By making the Real-time Transport Protocol (RTP) [14] an integral component of our architecture, the video gateway interoperates seamlessly with the current Internet video tools. We have demonstrated our implementation in the Bay Area Gigabit Network (BAGNet) and successfully used it to transcode high-rate JPEG [6] seminar video into

Multimedia '95, San Francisco, CA USA
© 1995 ACM 0-89791-751-0/95/11..$3.50

128 kb/s H.261 [18] streams for the MBone.

2 BACKGROUND AND MOTIVATION

In this section, we motivate the need for our video gateway architecture by giving several example scenarios.

2.1 BAGNet and MBone

One of our goals is to integrate the MBone applications into the BAGNet environment. We briefly describe BAGNet and MBone environments, then motivate the need for the video gateway by discussing issues of running existing MBone applications across both BAGNet and the general Internet.

BAGNet is an OC-3C (155 Mb/s) ATM network that connects 15 San Francisco Bay Area sites including universities, government and industrial research labs. One goal in the BAGNet project is to develop and deploy teleseminar applications with high quality video and audio. Currently, each site has only several workstations connected directly to the ATM network, and one of these machines is used as a gateway to connect BAGNet with other machines in the organization. Typically, ethernets are used within each organization. In a teleseminar scenario, one of the machines directly attached to BAGNet will multicast high quality thus high rate video across BAGNet. A stream of full motion JPEG compressed NTSC quality video will consume about 6 Mb/s bandwidth. While multiple 6 Mb/s video streams can be easily supported within BAGNet, transmitting even one of these streams onto an ethernet is not practical as the ethernets are shared. In addition, it may be desirable to transmit the video across the entire MBone where the aggregate bandwidth is only about 500 kb/s for all sessions. Therefore, we would like to accommodate three types of receivers in such an environment: hosts connecting directly to BAGNet, hosts connecting to BAGNet via ethernets and routers, and hosts on the global MBone.

Figure 1(a) shows such a scenario. Assume that host H0 is transmitting high quality JPEG video at 6 Mb/s. The ATM switch S forwards the transmission to BAGNet and to a router R0 connected to an ethernet, which contains among its hosts an MBone router R1. In order to not flood the ethernet, we run a video gateway on R0 that transcodes the 6 Mb/s JPEG down to 1 Mb/s. Thus, ethernet hosts such as H1 still receive the video transmission at a reasonably high quality. We place an additional video gateway between R1 and the MBone. This gateway further reduces the bandwidth requirements by transcoding the JPEG stream to H.261 and limits its output rate to 128 kb/s.

2.2 Linking Remote MBone Sites

Another application of our video gateway architecture is to link several remotely located MBone sites together. Consider the following example: a seminar at Berkeley is multicast locally to the campus, and a group of researchers at Carnegie Mellon University wants to tune in and participate. Unfortunately, a scenario like this — where small groups communicate via long haul networks — is not effectively supported by the current IP Multicast infrastructure.

The major mechanism of limiting the scope of multicasting traffic over today's MBone is use of TTL. The mechanism works as follows. Each multicast packet is assigned a Time-To-Live number in the IP header and each link is assigned a static threshold value. If the packet's time-to-live value is less than the threshold, the packet is not forwarded. By placing larger thresholds on the links that connect remote sites and using smaller TTL numbers for each multicast packet, the multicast traffic can be contained in a local region.

In the example mentioned above, in order for researchers at CMU to receive the multicast video, the video has to be sent out from Berkeley using a TTL value that is higher than the threshold value of cross-country links, which means that the video will be distributed across the entire country. This is not satisfactory.

With our video gateway, this problem can be easily solved by placing one video gateway at each site, and linking the two video gateways by unicast connections. Figure 1(b) illustrates this scenario. Multicast video from CMU is forwarded by the local video gateway via unicast connection to the remote video gateway, which in turn multicasts the video to the remote site at UCB.

2.3 Multicast Video across ISDN Links

Most MBone transmissions use a target rate of 128 kb/s for video and 64 kb/s for audio, precluding simple bridging of MBone sessions across a 128 kb/s ISDN line. By transcoding a 128 kb/s H.261 coded video stream into a 64 kb/s H.261 stream (and by similarly adapting the audio bit rate), an MBone session can be bridged over an ISDN link, as depicted in Figure 1(c).

2.4 Multicast Video over Wireless Links

Wireless links are characterized by relatively low bandwidth and high transmission error rates. Figure 1(d) illustrates the role of the video gateway in a topology containing mobile wireless hosts. By placing a video gateway at the basestation router (BS), we can transcode the incoming video stream to a lower bandwidth stream and control the rate of output transmission over the wireless link.

3 THE VIDEO GATEWAY

We now address the design of a video gateway architecture that will flexibly support the configurations discussed in the previous section. In order to interoperate with the existing MBone video tools, the video gateway must be RTP compatible. In this section, we give a brief overview of RTP, discuss the ramifications of transcoding an RTP data stream, and propose an architecture that can perform this transcoding in a protocol-consistent fashion.

3.1 Real-time Transport Protocol

The Real-time Transport Protocol [14] is an application-level protocol that is designed to satisfy the needs of multi-party multimedia applications. In the IP protocol stack, RTP lies

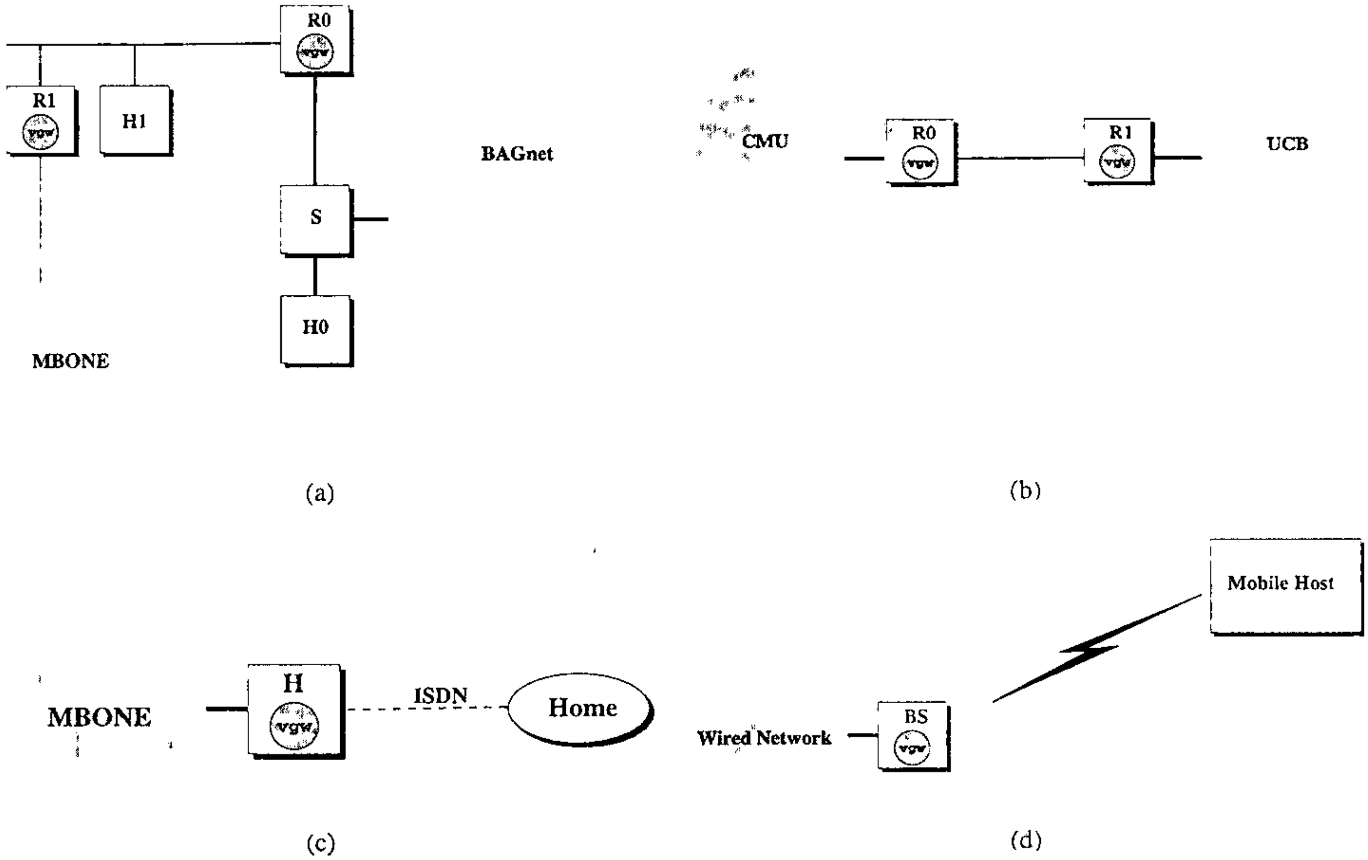

Figure 1: Four example scenarios for a video gateway.

just above UDP. The protocol consists of two parts: the data transfer protocol RTP and the control protocol RTCP.

Each RTP data packet consists of an RTP packet header and the RTP payload. The packet header includes a sequence number, a media-specific timestamp, and a synchronization source (SSRC) identifier. The SSRC provides a mechanism for identifying media sources in a fashion independent of the underlying transport or network layers. End-hosts allocate their SSRC randomly and because the SSRC's must be globally unique within an RTP session, a collision detection algorithm is employed to avoid conflicts. All packets from an SSRC form part of the same timing and sequence number space. Receivers group packets by SSRC identifiers for playback.

The RTP control protocol (RTCP) provides mechanisms for data distribution monitoring, cross-media synchronization, and sender identification. This control information is disseminated by periodically transmitting control packets to all participants in the session, using the same distribution mechanism as for data packets. The transmission interval is randomized and adjusted according to the session size to maintain the RTCP bandwidth below some configurable limit.

Distribution Monitoring. The primary function of RTCP is to provide feedback to the session on the quality of the data distribution. This information is critical in diagnosing failures and monitoring performance, and can be used by applications to dynamically adapt to network congestion [3]. Monitoring statistics are disseminated from active sources via RTCP sender reports (SR) and from receivers back to the entire session via receiver reports (RR). The SR statistics include, among other things, the sender's cumulative packet count and sender's cumulative byte count. Each receiver generates a separate reception report for each active source. The RR statistics include a cumulative count of lost packets, a short-term loss indicator, an estimate of the jitter in data packets, and timestamps for round-trip time estimation.

Synchronization. Since media streams are distributed on distinct RTP sessions (with distinct SSRCs), the protocol provides a mechanism for synchronizing media across sessions, which relies on canonical-name (CNAME) identifiers and SRs. Each SR contains that source's CNAME and the correspondence between local real-time and the media-specific time units. By matching sender reports across different media according to their CNAME, a receiver can align the time offsets of different media streams to reconstruct the original synchronization.

Sender Identification. Each session participant identifies itself by binding text descriptions to their SSRC using "source description" (SDES) messages.

In addition to the base protocol semantics, the RTP specification defines two types of "intermediate systems": *mix-*

DOCKET ALARM

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.