Report on DIMACS Workshop on
Multimedia Streaming on the Internet

Workshop Organizers:

Lixin Gao	Dept. of Computer Science, Smith College	gao@cs.smith.edu
Jennifer Rexford	AT&T Labs – Research	jrex@research.att.com

ReportAuthors:

Gabriela Alexa	RUTCOR, Rutgers University
Michael Bradshaw	University of Massachusetts
Lixin Gao	Dept. of Computer Science, Smith College
Jennifer Rexford	AT&T Labs – Research
Nicolas Schabanel	DIMACS, Rutgers University
Shubho Sen	University of Massachusetts
Bing Wang	University of Massachusetts

1 Workshop Focus

The last few years have witnessed a dramatic growth of multimedia applications such as video and audio streaming over the Internet. Server and network resources have proven to be a major limiting factor in the widespread usage of multimedia streaming applications. Supporting high-quality streams and a large population of clients requires innovative techniques that make efficient use of server and network resources. The purpose of this workshop was to bring together leading researchers in the field to discuss issues relating to server media-content delivery scheduling, the disk I/O scheduling, and network and proxy support for delivering multimedia content.

2 Keynote

"Internet Media Delivery Systems: Current Design, Open Issues, Research Problems"
Yuriy A. Reznik, Sr. Research Scientist, RealNetworks, Inc.

Summary: Incorporated in 1995, Real Networks is a pivotal player in the Internet media delivery space, with industry figures suggesting that as much as 85% of existing Internet streaming media content is available in the proprietary RealMedia format. Reznik gave an overview of the components in the Real architecture responsible for the various stages in streaming audio and video - creation and encoding (RealProducer), Serving (Real Server and Real Proxy), and playback (Real Player). RealSystem 8 uses a new proprietary video encoder which can produce VHS-quality video at $220 - 500$ kbps, and DVD-quality at 1 Mbps, and can provide the same quality video as an MPEG clip encoded at twice the bit rate. Reznik demonstrated the superior quality of this new encoder by simultaneously playing (i) RealPlayer - a clip encoded using the new encoder, and (ii) Windows Media Player - an MPEG-encoding of the same video clip at twice the bit rate as the proprietary encoding.
The talk focussed on a number of research problems in the media coding space including

- Supporting bounded delays, bounded bitrates, random-access decoding, transmission through multiple channels with joint decoding at the reception side.

- Source modeling issues - tradeoffs between transient and asymptotic efficiency of the codes, dealing with lossy compressions, the need to compute more global statistics to get better compression, and determining what distortion measure to use.
- IP & Channel coding issues - should we use joint source/channel coding.
- Tradeoffs between server storage usage and coding complexity.

Author biography : Yuriy Reznik received his MS in electrical engineering from Kiev Polytechnic Institute and his PhD in from the National Academy of Sciences in Ukraine. Research interests include theory of computing, combinatorics, analysis of algorithms, information theory, digital signal processing, and their applications. Since 1997, he has been a been a member of the research staff at RealNetworks. At RealNetworks, he was one of the key designers of the RealSystem G2, contributing most extensively to the RealVideo and RealAudio compression algorithms. Most recently, he was in charge of the development of statistical source and channel modelling techniques for the core components of RealSystem 8.0.

3 Proxy Servers

"The Role of Network Proxies in Supporting Multimedia Streaming"
Antonio Ortega, University of Southern California

Author abstract: In a network which shows increasing heterogeneity in terms of bandwidth, delay and reliability, proxies can serve as the preferred tool to enable adaptation, without requiring an end-to-end involvement of the application. Moreover, application- and media-specific functions should be preferred to generic ones.
In this talk, we presented an overview of recent research in video streaming. The focus is to explore the potential benefits of introducing video specific functions at proxies. Three examples of such functions include caching, local retransmission and introduction/removal of error protection.

- Video caching, where the cache has the ability to store only a selected set of frames from a video sequence. We show how this increases the robustness of delivery to the receiver, in cases where cache storage is limited.
- Local error control via ARQ, where the final link to the client is assumed to be lossy. We show how ARQ combined with rate control at the proxy can minimize the loss of video information at the receiver.
- Multiple description coding (MDC) for error robustness. We described an MDC approach where redundancy can be easily controlled. This allows a proxy to increase or decrease the level of redundancy, depending on whether the link to the client is more or less reliable than the rest of the network.

Summary: Antonio Ortega argued that proxy servers placed at the edge of the network can adapt a generic stream to the needs of the user in heterogeneous, asynchronous, lossy environments, and presented the gains of considering video specific functions as local retransmission, caching, and introduction/removal of error protection. Three examples illustrated the benefits of using proxy servers: a combination of ARQ-based local error control and rate control for loss- recover on a lossy client-proxy path - the technique can take advantage of a short roundtrip time (RTT) on the proxy-client path, and the ability to predict the channel behavior on a short time scale, and requires only local retransmissions; video caching (where the goal is to use selective frame caching to increase robustness); using multiple description coding (MDC) for adapting the redundancy in a stream to the loss characteristics of the proxy-client link, to ensure minimal distortion reception.

"Video Caching and Delivery Using Proxy Server for Reducing Bandwidth Requirement over Wide Area Networks"
Weihsiu Ma and David H.C. Du, University of Minnesota

Author abstract: Due to the high bandwidth requirement and rate variability of compressed video, delivering video across wide area networks (WANs) is a challenging issue. Proxy servers have been used to cache web objects to alleviate the load of the web servers and to improve client access time on the Internet.
In the model, a central server is connected to a proxy server via WAN and a proxy server can reach many clients via Local Area Networks (LANs). The proxy server allows partial video caching and thus a certain number of video frames are stored in its storage such that the network bandwidth requirement over WAN can be reduced. Since there are two video sources, video data has to be synchronized before playback at a client. Two delivery models, client-synchronization model and proxy-synchronization model in terms of where to synchronize data are identified. The two models have different complexity and resource consumption in proxy server and the client.
The main focus of the talk is on the fundamental understanding on delivering single video by using the two models and investigate the effectiveness of the partial video caching in bandwidth reduction. We study the tradeoffs between client buffer, start-up delay, resource requirement in the proxy server, and bandwidth requirement over WAN for video transport. Given a video delivery rate for the WAN, we propose several frame caching selection algorithms to determine which frames of the video are to be stored in the proxy server. In client-synchronization model, a scheme, which partitions a video into different segments (chunks of frames) with alternating chunks stored in the proxy server, is shown to offer the best tradeoff. In proxy-synchronization model, caching the initial portion (prefix) of a video is adopted and this approach is proved to minimize the space requirement in proxy storage.

Summary: David H.C. Du argued that partial video caching at video proxy servers can be used to reduce the WAN transmission bandwidth requirements, and to decrease client startup delays. Two video delivery models (client-synchronization a the proxy-synchronization) for synchronizing the video data from the origin server and the proxy were presented. A number of *frame selection* algorithms (Just in Time, Minimum Difference, and Chunk) were presented for determine what portions of the video should be stored at the proxy. A comparative performance study of these algorithms in the context of the two synchronization models was also reported.

"Issues in Design and Evaluations of Multimedia Proxy Caching Mechanisms for the Internet"
Reza Rejaie, AT&T Labs - Research, Haobo Yu, USC ISI

Author abstract: Most of the existing Internet multimedia streaming applications are based on the client-server architecture; a media server pipelines a stream to a client through the network while the client plays back the available portion of the stream. This client-server architecture has two major limitations:

1. The quality of the delivered streams is limited to the bottleneck bandwidth between the server and the client. Thus a client with a high bandwidth local connectivity may receive low quality streams due to a remote bottleneck.

2. Its scalability is limited, in that it is difficult to support a large number of concurrent high quality sessions due to network and server load.

Multimedia proxy caching(MCaching) is a natural extension of the client-server architecture that addresses all the above challenges simultaneously. The idea is to cache popular streams with maximum deliverable quality

at a proxy close to interested clients. Thus the proxy can effectively maximize delivered quality and improve scalability by significantly reducing the load on the network and the server.
In this talk, we provide insight in the design and performance evaluations of MCaching mechanisms. In addition to cache efficiency (measured by hit ratio), MCaching introduces the notion of "quality" for delivered streams as a new dimension to Web caching.
We first outline the design of an efficient MCaching mechanism that consists of

1. an internal structure (e.g. layered) for cached streams,
2. a fine-grain replacement mechanism based on fine-grain popularity and encoding-specific information (e.g. utility), and
3. an online fine-grain prefetching mechanism to further improve the delivered quality on-the-fly.

We argue that performance of MCaching mechanisms should be collectively evaluated along two dimensions: caching efficiency and delivered quality. Specifically, a fundamental tradeoff is identified between these two dimensions. The simulation results show that the proposed MCaching design can adaptively exploit this tradeoff to improve overall performance whereas simple extension of current Web caching schemes can only enhance the performance along one dimension.

Summary: Reza Rejaie presented the multimedia proxy caching mechanism (MCaching) as an improved extension of the existing client-server architecture, that allows replications and proxy caching, and which introduces the delivered streams quality as a new dimension to the proxy caching. The speaker analyzed the performance of MCaching mechanism of maximizing the byte-hit ratio (BHR) and the delivered quality. Finally, he presented some of the future directions of work (to design a prototype for MCaching, to use the utility function for prefetching and replacements, and to analyze the multimedia traffic).

"Traffic Smoothing for Network Wide Streaming of Multimedia Traffic"
JunBiao Zhang and Joseph Hui, Arizona State University

Author abstract: Traffic smoothing for video streams are considered for network configurations such as single nodes, nodes in tandem and routes in parallel. We place constraints on the smoothed traffic, such as rate constraints or on-off constraints. Optimal algorithms are derived for some of these cases minimizing delay or buffer requirements or maximum rates. Experimental results are also presented for traffic smoothing. Some exploration on how these results could have an impact on practical implementations of video streaming applications and signaling requirements is also presented.

Summary: Joseph Hui presented some concerns for traffic smoothing for single node, multi-node in tandem, and routes in parallel network configurations. The speaker presented the sliding-stretching algorithm as an optimal linear algorithm for minimizing the peak rate, the low rate, the number of rate changes and the buffer requirements. Also, some results obtained by applying the compromise of partial reverse smoothing and forward smoothing were discussed.

4 Content Distribution

"Decentralized Movie Distribution in the Internet"
Carsten Griwodz, Darmstadt University of Technology

Author abstract: After the failures of initial video-on-demand developments, the rapid growth of the Internet is rising a new attraction to VoD-like applications. Intranet VoD is successfully following the track of centrally managed video distribution system. For public VoD, we consider it likely that such an approach would fail as it did some years ago. Rather, decentralized distribution systems that are organized in a similar way as the existing video rental store infrastructures seem to be more adequate for adaptable growth. In this talk, we present one consistent distribution infrastructure based on caching and the results that we have produced so far in investigating the unsolved issues of the infrastructure. Specifically, this concerns protocol support, support for copyright violator tracing, and the efficiency of the distribution system. A protocol suite that allows low overhead reliable transfer of data into cache servers is introduced. In the absence of multicast-capable watermarking schemes, an alternative, straight-forward, personalized movie marking scheme is offered for discussion. Finally, the efficiency gains that can be achieved by an interaction of caching strategies and new distribution mechanisms are presented.

Summary: Carsten Griwodz argued that the traditional video-on-demand model of centrally managed video distribution would not scale for Internet-wide distribution, and a distributed scheme would be much better at decreasing network and server loads. He proposed a hierarchical technique called called gleaning, that combines the individual techniques of patching and caching. The stream is sent out through the distribution tree and each node stores a small part of the ongoing stream to serve as a patching channel for new clients. Experimental results for the gleaning shows that the technique is less expensive than simple caching or patching. Griwodz also introduced some new protocol suites (as LC-RTP, LC-RTCP etc.) that allow low reliable transfer of data into cache servers, and discussed techniques being developed for increased security.

"On Content Delivery Networks for Streaming Media"

David Shur, AT&T Labs

Author abstract: We describe our ongoing project on a content distribution network (CDN) for streaming media. CDNs support the delivery of streaming media in an efficient and scalable manner. Multimedia sources send a single stream to the CDN.
Based on the location and number of end-users, the CDN replicates the stream (typically via a tree topology) as needed. Our architecture consists of an overlay network of servers which exploit the functionality of IP Multicast where available. Where IP Multicast is not available, the overlay servers utilize unicast forwarding to reach end-systems. The servers also provide packet recovery protocols which enable high perceived quality of service through real-time recovery of lost packets. We compare our system with emerging commercial CDNs, and discuss the challenges in this area.

Summary: David Shur argued that a content distribution system needs to handle a number of important functions including access control and accounting, scalable delivery with consistent QoS, server selection policies (which server to connect to), and content peering. Shur argued that for widespread streaming content delivery, unicast methods do not scale well. Native IP multicast has a number of limitations including difficulty of deployment, lack of access control and detailed accounting, and lack of interoperability with non-multicast clients. To support the delivery of streaming multimedia in an efficient, scalable manner, Shur proposed a content distribution architecture consisting of an overlay network of servers performing various application-level functions, and communicating via IP multicasting where possible. In this architecture, clients access the various application services (splitters, gateways, packet recovery, repair, mirror, transcoding etc.) via edge servers. In addition to handling QoS problems, an edge server would host packet recovery and retransmission modules, with dedicated bandwidth to the corresponding modules on the other edge servers.

“Selecting among Replicated Multicast Video Servers”
Mostafa H. Ammar, Georgia Institute of Technology

Author abstract: Server replication is often used to improve the scalability of a service. One of the important factors in the efficient utilization of replicated servers is the ability to direct clients to a server according to some optimality criteria. Most server replication and selection work to date has focused on traditional unicast services. We will first motivate the need to replicate multicast video servers and argue that this is somewhat different rationale from the one used to motivate unicast server replication. We will then show some results from on-going work aimed at developing a framework and a set of algorithms and protocols for multicast server selection.

Summary: Mostafa H. Ammar proposed using multicast server replication as a solution for distributing the load for on-demand services, for balancing the feedback and for optimizing the client or the global performance. The speaker presented on-going work on the multicast server selection problem - allocating each client to an appropriate server such that a metric of interest is globally optimized (minimize the total hop cost, minimize the total resource consumption etc). The static multicast server selection problem can be modeled as a Steiner Tree Problem. For the dynamic variant of the problem, three heuristic policies (nearest server, nearest tree, nearest node) were presented. Other problems such as selection among rate-adaptive multicast servers, bandwidth bottlenecks, shared links, and heterogeneous access points for the client will be addressed in future research.

5 Panel: Multimedia Streaming and Network Interaction

Four panelists gave their viewpoints on the hot topics in Multimedia Streaming and answered questions from the floor. Below is a brief description of the panelists’ viewpoints and issues that were brought up.

Barry G. Haskell, AT&T : There are many dot-com companies starting up and a lot of good research is going into those areas. In particular, Haskell wanted us to focus on personal multimedia systems such as TIVO and the copyright holder’s responses to them. A second issue was how do we realize quality of service for streaming media. We need to create content that can be adapted to meet the current network capabilities, and whether we use a layered approach or MDC is hard to tell at this point.

Katherine Guo, Lucent: Guo presented some aspects of the multimedia content distribution over the Internet - research and practice. The speaker pointed out the fact that the Internet media is a billion industry and that there is a race to build content delivery systems. In order to achieve this, proxies are becoming very popular in solving the various delivery problems. Other problems encompass scalability, storage, concurrent clients, open API standards for communications and encoding, and how companies can make money from multimedia content on their web sites.

Leana Golubchik, University of Maryland: Leana focused on the issues involved in creating efficient storage for the multimedia servers that are being developed. She made the important point that though storage costs are decreasing, disk bandwidth costs are not, and hence there is a need to develop scalable multimedia storage systems. One of the unique aspects of disk storage is that the data layout can significantly impact the load imposed on the system, implying that the layout must be done depending on how the data will be used later.

Mostafa Hamed Ammar, Georgia Tech: Mostafa talked about “democratic” multimedia streaming. It should be made for the people (scalable) and by the people. The speaker emphasized that up to now the “by the people” aspect was ignored, and that in such a vision, everyone can be a source of multimedia, and broadcast their own TV /Radio or web/video server. Multicast was meant to address both of these issues but people have shied away from giving people enough control to achieve "democratic" streaming. In order to achieve this there needs to be a removal of last hop bottle necks, efficient small group multicast, QoS on a per-flow basis, security, and pricing for data rates.

General Questions

With the rate of bandwidth increases will there ever come a point where there is no reason to continue this research?

The panel and audience agreed that there will always be a lack of bandwidth. This was attributed to two factors. One is that people are not charged for the bit rate that they use and they have no incentive not to use more if the network can support it. The second point is that as more bandwidth becomes available new services are introduced that require that bandwidth. There were several mentions of virtual reality services and guesses at where the future of computing might lead.

Given that LAN environments are overprovisioned and the actual Internet Backbone is overprovisioned, where does the bottleneck lie on the end-end path?

This could occur at different places. It was widely suggested that one location of bandwidth bottlenecks would exist at peering points between different ISPs.

What are the thoughts on regulating UDP to be TCP friendly?

There was a question as to what TCP friendly is. TCP is particularly bad for streaming since it is very bursty. We need to look at solutions which will avoid congestion and allow us to stream data at constant rate. As of yet there are many propositions but few are widely accepted.

How do we cope with closed standards?

There was once again great division on this problem. Besides the obvious solution to open the standards, many people proposed that the community work only with open standards. While this doesn't allow for inter-operability with other closed standards it does allow useful research and patents to proceed. The research can then be applied to commercial products.

Panelist biographies:

- Mostafa Ammar is a Professor with the College of Computing at Georgia Tech. He received the S.B. and S.M. degrees from the Massachusetts Institute of Technology in 1978 and 1980, respectively, and the Ph.D. in Electrical Engineering from the University of Waterloo, Ontario, Canada in 1985. For the years 1980-82 he worked at Bell-Northern Research (BNR), first as a Member of Technical Staff and then as Manager of Data Network Planning. In 1999, he was on a sabbatical leave with BellSouth Telecommunications' Science and Technology organization and returned to his Ivory Tower in January 2000. His research interests are in the areas of computer network architectures and protocols, distributed computing systems, and performance evaluation. He currently serves as the Editor-in-Chief of the IEEE/ACM Transactions on Networking. He is the co-author of the textbook "Fundamentals of Telecommunication Networks," published by John Wiley and Sons.

- Leana Golubchik received her Ph.D. degree from the Computer Science Department at the University of California at Los Angeles in 1995. She is currently an Assistant Professor in the Department of Computer Science at the University of Maryland at College Park. From Fall of 1995 until Summer of 1997, she was an

Assistant Professor in the Department of Computer Science at Columbia University. Her research interests include computer systems modeling and performance evaluation, Internet-based computing, and multimedia storage systems.

- Katherine Guo received her PhD in Computer Science from Cornell University in 1998. She has been a Member of Technical Staff at Lucent Bell Labs' Networking Software Research Department ever since. She is one of the leaders in IPWorX WebCaching Project which is a Lucent product originated from Bell Labs Research.

- Barry G. Haskell received the B.S., M.S. and Ph.D. degrees in Electrical Engineering from the University of California, Berkeley in 1964, 1965 and 1968, respectively. From 1968 to 1996 he was at AT&T Bell Laboratories, Holmdel, New Jersey. Since 1996 he has been at AT&T Labs, Middletown, New Jersey where he is presently an AT&T Fellow in the Image Processing and Software Technology Research Department. His research interests include digital transmission and coding of images, videotelephone, satellite television transmission, medical imaging as well as most other applications of digital image processing. He has published over 60 papers on these subjects and has over 40 patents either granted or pending. He is a coauthor of several books on digital video and is very active in the establishment of International Video Communication Standards.

6 Reliability and Loss

"Adaptive Streaming of Stored Layered Video over Lossy Channels"

Srihari Nelakuditi, Zhi-Li Zhang, Sandeep C. Rao
University of Minnesota, Minneapolis

Author abstract: In recent years, most popular Internet application is web-based audio and video playback where stored video is streamed from the server to a client upon request. Rigid playback deadlines coupled with resource constraints make video delivery a challenging task. Video smoothing techniques reduce the bandwidth requirement by using client buffer for pre-fetching. However, when both network bandwidth and client buffer are limited it may not be possible to deliver full-quality video. In such a situation, it is desirable to minimize the degradation in the video quality while operating within the resource constraints. Layered encoding is proposed to provide finer control on video quality where the video signal is split into layers and a prefix of these layers is chosen such that the resource constraints are met. However it is not feasible to precompute the optimal number of layers since the network conditions are continually varying. So the main concern is how to choose the ideal number of layers adaptively based on the prevailing network conditions. This work addresses this layer selection problem for the case of stored layered video delivery over lossy networks such as wireless. Layer selection schemes is proposed that utilize the knowledge about the bandwidth and buffer requirements of the stored video in delivering smoother quality video.
One of the problems in assessing the performance of a video delivery scheme is the lack of a good metric that captures the user's perception of video quality. In general, the higher the amount of detail in the played video, the better is its quality. However, it is generally agreed that it is visually more pleasing to watch a video with consistent, albeit lower, quality than one with highly varying quality. Thus, a good metric should capture both, namely, the amount of detail per frame as well as its uniformity across frames. In a layered video delivery, ideally we would like to have a high mean and a low variance in the number of layers per frame played. We devised a metric called "mean layer run" which is defined as the number of layers a user is expected to see continuously when the video is watched for a given observation period starting at any arbitrary frame. This metric thus accounts for both mean and variance and hence measures both detail and uniformity of the video playback. This metric was used to measure the perceived quality of the delivered video for evaluating the performance of various video delivery schemes.

We proposed a MINimal Layer Discard algorithm (MINLD) that maximizes the quality, under given network bandwidth and client buffer constraints, by selectively discarding higher layers in order to minimize the likelihood of future lower layers arriving late, thereby increasing the overall quality of the video delivered. We develop an online variant of this offline algorithm, called Finite Horizon based Minimal Layer Discard (FHMLD) for adaptive layer selection. This scheme is applied to a wireless setting where channel bandwidth is fixed and known but of varying loss with retransmissions for error recovery. The video delivery based on FHMLD scheme can be summarized as follows: periodic estimation of loss rate based on observed channel conditions; estimation of effective bandwidth and effective delay based on measured loss rate; selective layer level discard based on effective bandwidth, effective delay; discarding of packets that would anyway arrive late.
The simulation shows that FHMLD is better in maintaining consistent quality playback than a simple greedy layer selection scheme. The future work is to extend the scheme to the case where both bandwidth and loss are unknown and varying.

Summary: Srihari Nelakuditi presented some solutions for the problem of smoothing video applications, under various assumptions. When the quality metrics and the available bandwidth are given, and all the frames in a layer have the same size (CBR), the speaker presented the generic and the offline algorithms and their performances. When there is no knowledge about the future available bandwidth and the problem is what layers to select, and what frames of a layer to prefetch, the speaker presented an adaptive algorithm for this case.

"Optimal Streaming of Layered Video"
Despina Saparilla, Institut Eurecom

Author abstract: This paper presents a model and theory for streaming layered video. The bandwidth available to the streaming application is modeled as a stochastic process whose statistical characteristics are unknown a priori. The random bandwidth models short term variations due to congestion control (such as TCP-friendly conformance). The assumptions are: the video has been encoded into a base and an enhancement layer, and that to decode the enhancement layer, the base layer has to be available to the client, and that the client has abundant local storage and attempts to prefetch as much of the video as possible during playback. At any instant of time, starvation or partial starvation can occur at the client in either of the two layers. During periods of starvation, the client applies video error concealment to hide the loss. We study the dynamic allocation of the available bandwidth to the two layers in order to minimize the impact of client starvation. For the case of an infinitely-long video, the optimal policy takes on a surprisingly simple and static form. For finite-length videos, the optimal policy is a simple static policy when the enhancement layer is deemed at least as important as the base layer. When the base layer is more important, we designed a threshold policy heuristic which switches between two static policies. We provide numerical results that compares the performance of no-prefetching, static and threshold policies.

Summary: Despina Separilla presented a model and a theory for streaming layered continuous media, and focused on streaming continuous media over fair-shared bandwidth. The layer streaming model allocates the available bandwidth between two layers. The speaker presented a detailed analysis of the model in the case of infinitely and finitely long video, and concluded that prefetching the continuous media during playback is critical.

7 Resource Tradeoff

"Frame-Based Periodic Broadcast and Fundamental Resource Tradeoffs"
Subhabrata Sen, University of Massachusetts, Amherst

Author abstract: The Internet is witnessing a rapidly increasing load of continuous media traffic in the form of streaming audio and video. Video streams typically have high transmission bandwidth requirements, and exhibit burstiness on multiple time scales, making it expensive to deliver it.
In this talk, we explore fundamental resource tradeoffs in periodic broadcast, a technique for reducing network transmission bandwidth requirements for streaming a popular video to multiple asynchronous clients. We consider a class of frame-based periodic broadcast schemes which assign a fixed transmission bandwidth to each frame in the video, and continuously transmit each frame at this rate. The model accommodates both CBR and VBR streams. We describe a novel broadcast scheme that minimizes the transmission bandwidth overhead under client buffer constraints. We also consider the problem of using a single transmission scheme to satisfy clients with heterogeneous resource constraints, and present a heuristic client reception scheme to jointly minimize the client playback startup delay and client buffer requirements. Finally, the talk concludes with results from extensive evaluations that explore the tradeoffs among the network transmission bandwidth requirement, and client resources (buffer, reception bandwidth, and playback startup delay).

Summary: Subhabrata Sen presented work on periodic broadcast, a technique for reducing network transmission bandwidth requirements for streaming a popular video to multiple asynchronous clients. In this technique, the server segments the video and repeatedly transmits each segment on a separate multicast channel. To ensure starvation free streaming playback, a client may need to perform workahead buffering and to receive from more than one channel at a time. The speaker considered a class of frame-based periodic broadcast schemes which assign a fixed transmission bandwidth to each frame in the video, and continuously transmit each frame at this rate. Two new broadcast schemes were presented, along with the tradeoffs among the network transmission bandwidth requirement, and client resources (buffer, reception bandwidth, and playback startup delay). The speaker also presented some topics of ongoing research (investigation of periodic broadcast and patching, development of AMPS streaming testbed etc).

"Algorithms for On-Demand Stream Merging"

Amotz Bar-Noy, AT&T Shannon Labs
Richard E. Ladner, University of Washington and AT&T Shannon Labs

Author abstract: As the Internet grows so does the desire for on-demand streams of many types: movies, songs, news stories, stock quotes, and others. The popularity of a specific stream may be so high that multicasting may be the only way to satisfy the demand. In addition, clients requesting a stream will want service as quickly as possible. This may require repeated multicasts of the same stream to satisfy the delay guarantees. Stream merging has the potential to help solve the bandwidth problems created by heavy, low delay, demand for the same stream.
This talk describes the stream merging technique and how it can be used to reduce bandwidth requirements at the stream server. We describe a new quadratic off-line algorithm for minimizing bandwidth and the optimal solution for the fully loaded case, where streams are requested at unit time intervals. We analyze the approximation ratio for oblivious off-line algorithms that only know the number of requests and not their arrival times. Finally, We briefly describe a new approach to on-line stream merging based on our optimal oblivious offline algorithm. It turns out that the Fibonacci numbers play an important role in the analysis.

Summary: Amotz Bar-Noy presented the Stream Merging technique for on-demand media streaming to reduce the resource needs of the server, reduce the bandwidth, and to satisfy the client by reducing the startup delay. The speaker argued that stream merging may save up to 1/4 of the bandwidth. The performances of the off-line merging, oblivious merging, and on-line merging algorithms were analyzed. In particular, the optimal off-line merging was improved to a quadratic complexity, and in the case of oblivious merging, their exact approximation ratio is defined as a function of the resulting merging tree.

8 Multimedia Delivery

"The Design and Implementation of a Media Independent Streaming Service for Stored Video Applications"

Wu-chi Feng, Ohio State University

Author abstract: In this talk, we describe our work focused on delivering high-quality video content over best-effort networks (such as the Internet) for stored video streams. The goal of this work is two-fold. First, we are interested in using algorithms that adapt stored video sources efficiently to the available network resources, avoiding high packet losses and congestion collapse. Our approach advocates the use of TCP as the transport mechanism for stored video delivery. In addition, our delivery technique uses a priority-based algorithm that helps smooth the frame rates delivered to the end user effectively. Second, we have designed a media independent streaming service to show the viability of our approach. By providing this media-independent interface, we hope to catalyze future video-based applications such as scientific visualization.

Summary: Wu-chi Feng argued that currently all multimedia applications are trying to achieve some form of congestion control using UDP. Feng presented a TCP transmission based technique for the delivery of advanced streaming video (in particular non-live video applications) over Internet, which avoids the packet losses, and provides flow/congestion control. Feng is currently working on a session layer protocol which lies on top of the network layer. This layer provides the application with information on the stream size and allows prioritization of packets so that important information is sent out ahead of time to ensure continuance in quality. By doing this, the stream can more easily fit into the TCP session. Feedback to upper layers allows the application to match its flows more carefully to what is currently available.

" TCP-like flow control for multimedia streaming using TCP emulation at receiver (TEAR)"

Injong Rhee, North Carolina State University

Author abstract: Congestion and flow control is an integral part of any Internet data transport protocol. It is widely accepted that the congestion avoidance mechanisms of TCP have been one of the key contributors to the success of the Internet. However, TCP is ill-suited to real-time multimedia streaming applications. Its bursty transmission, and abrupt and frequent wide rate fluctuations cause delay jitters and sudden quality degradation of multimedia applications. For asymmetric networks such as wireless networks, cable modems, ADSL, and satellite networks, transmitting feedback for (almost) every packet received as it is done in TCP causes congestion in the reverse path, causing feedback losses and delays. In this environment, TCP may severely under-utilize the forward path throughput. Use of multicast further complicates the problem; TCP-like frequent feedback from each receiver to the sender in a large scale multicast session cause well-known scalability limitations, such as acknowledgment implosion.

I have developed a new flow control approach for multimedia streaming, called TCP emulation at receivers (TEAR). TEAR shifts most of flow control mechanisms to receivers. In TEAR, a receiver does not send to the sender the congestion signals detected in its forward path but rather processes them immediately to determine its own appropriate receiving rate. TEAR can determine this rate using congestion signals observed at the receiver, such as packet arrivals, packet losses, and timeouts. These signals are used to emulate the TCP sender's flow control functions at receivers including slow start, fast recovery, and congestion avoidance. The emulation allows receivers to

estimate a TCP-friendly rate for the congestion conditions observed in their forward paths. TEAR also smoothes estimated values of steady-state TCP throughput by filtering out noise. This smoothed rate estimate will be reflected into the rate adjustment of receiving rates. Therefore, TEAR-based flow control can adjust receiving rates to a TCP-friendly rate without actually modulating the rates to probe for spare bandwidth, or to react to packet losses directly. Thus, the perceived rate fluctuations at the application are much more smooth than in TCP.
A unicast version of TEAR is implemented. In this talk, I will describe the implementation of TEAR, examine the performance of this TEAR implementation from the NS simulation and Internet experiments, and compare it with that of other TCP-friendly flow control techniques. Our preliminary tests indicate that TEAR shows superior fairness to TCP with significantly lower rate fluctuations than TCP. TEAR's sensitivity to feedback interval is very low, so that even under high feedback latency, TEAR flows exhibit acceptable performance in terms of fairness, TCP-friendliness, and rate fluctuations.

Summary: Injong Rhee showed us that with the increasing popularity of multimedia, congestion control is becoming more and more important. TCP's bursty nature does not suit multimedia application. Instead Rhee offered a protocol called TEAR (TCP emulation at receivers). A TEAR receiver observes the reception stream (packet arrivals, packet losses, and timeouts) and emulates the TCP sender's flow control functions over a period of time. At the end of the time period, the TEAR receiver estimates a TCP-friendly rate and sends this back to the TEAR sender. In this manner it emulates TCP congestion control without massive shifts in bandwidth and data recovery. TEAR is suitable for Internet-based multimedia streaming because (1) its rate adjustment is very smooth, and (2) it is TCP-friendly. TEAR has potential application in congestion and flow control for multicast because most of the functionalities are performed at receivers and it is thus very scalable. The performance evaluation of a TEAR implementation was also illustrated.

9 Open Problem and Challenge Session

Jennifer Rexford, AT&T Labs, observed that the topics covered in this workshop are related to content distribution (proxy services, overlay networks, asynchronous clients), and transport (multiple encoding, retransmission, selective discard, workahead smoothing). Next, she referred to some open problems from research area not covered in the workshop.

Media Adaptation

- Appropriate adaptive coding scheme (layered, multiple description, or multiple coding)
- Characteristics of the adaptation scheme (e.g., coding efficiency, granularity of adaptation, effects on efficiency)
- Protocol support for fine-grained control of adaptation (time, layers)
- System overheads for fine-grained control (data structure, storage, processing)
- Interaction between server-proxy and proxy-client adaptation (competing control loops)?
- Disk layout for adaptive streaming - anything special?
- Caching in the presence of adaptation (layered coding)?

Measurements/Characterization

- Characterization (popularity of streams, arrival process, VCR interaction, join/leave patterns)
- Techniques for identifying bottleneck

- Simulation platform for multimedia streaming (do we need an ns-like system for streaming simulation)?
- Monitoring/benchmarking of overlay networks
- Testbeds - open software for clients, proxy and server

Applications

- Virtual reality, telmedicine, multi-party game, user-developed content
- Networking environments, wireless networks
- Policy issues

Lixin Gao, Smith College, continued the session, by enumerating some ”religious” open questions.

Religious Questions

- Shall we ever have enough bandwidth? -How many bits does one person really need?
- Why should multimedia traffic be required to be TCP-friendly? What does being fair mean?
- De facto standards vs. IETF standards? Will the DoJ split RealNetworks?
- Is research in multimedia streaming relevant? Should we all go to startup companies?
- Is research in networking relevant?
- Will we ever have widespread deployment of QoS, multicast? Back to ATM? Separate networks?
- Security? Various protection schemes in proxy?
- Long-term research?

10 Multimedia System Design

“Striping Doesn’t Scale: How to Achieve Scalability for Continuous Media Servers with Replication”
Leana Golubchik, University of Maryland

Author abstract: Multimedia applications place high demands for QoS, performance, and reliability on storage servers and communication networks. These, often stringent, requirements make design of cost-effective and scalable continuous media (CM) servers difficult. In particular, the choice of data placement techniques can have a significant effect on the scalability of the CM server and its ability to utilize resources efficiently.
In the recent past, a great deal of work has focused on “wide” data striping. Another approach to dealing with load imbalance problems is replication. The appropriate compromise between the degree of striping and the degree of replication is key to the design of scalable CM servers. Thus, the main focus of this work is a study of scalability characteristics of CM servers as a function of tradeoffs between striping and replication.

Summary: Leana Golubchik presented striping as a solution for the load imbalance problem for video file placement. High quality videos are typically very large. Unfortunately the bandwidth between the disks and the processor is limited. When requests for different videos come in, the disk would not be able to handle the multiple requests and would fail. There are two mainstream approaches to solve this problem. Striping involves dividing the file into smaller pieces and spreading the pieces on multiple disks. Another technique called replication involves copying a popular video on multiple disks so that there is enough bandwidth to retrieve them. The speaker illustrated the performances of the striping method (load balance, less cost for storage) vs. the performances of the replication (scalability, reliability), as well as the limits for striping (scalability, implementation and expansiveness) and for replication (the cost for storage). Golubchick showed that by themselves neither of these techniques scale well. A hybrid systems consisting of multiple nodes with striping across disks within a single node, and replication across nodes is shown to possess a greater ability to scale and incrementally grow. Golubchik argued that intra node striping and inter node replication is preferable, and presented simulation results for the comparison between wide data striping vs. the hybrid method.

"Semantic Transformation of Multimedia Streams inside an Active Network"
Maximilian Ott, C&C Research Laboratories, NEC USA, Inc.

Author abstract: Providing multimedia services to a large number of heterogeneous clients raises the fundamental problem of incompatibility between the native format of the media on the server, and the optimal format for each client. This brings up the need for transforming a multimedia stream, not only to reduce bit-rate, but also to enhance the acceptability of its play-out. We discuss the implications of performing such media stream transformations at different positions in the network, and propose a programmable network architecture for providing this service. We also describe an experimental testbed we built to determine the feasibility of our ideas.

Summary: Maximilian Ott examined the upcoming role of active networks. In particular he examined the concept of allowing the network to compute for the end users in a model called Journey. Journey is a network filled with data processing nodes. When packets pass by the nodes, some are randomly brought in and processed. Journey makes the same guarantees that UDP does. Packets might not reach the user and if they do they might not be processed. This allows for the network to be created with a stateless design. Ott highlighted other various points on his architecture.

"eSeminar Project Overview"
Martin G. Kienzle, IBM Research

Author abstract: The research objective of the eSeminar project is to provide a test bed for experiments with advanced multimedia technology in the operational setting of a complete end-to-end system, and to provide an integration point for many multimedia technologies. The operational goal of the system is to improve communication in the IBM Research Division by recording video of talks and meetings, as well as collateral information, and making it available to all Research employees worldwide. In order to make this system truly usable, we want to automate all aspects of recording and distribution to radically reduce the operational complexity cost, and to let the system completely "fade into the background". A second usability goal is to facilitate focused video access through indexing and other meta data, and to improve ease of access, to make the information available and usable to a very large number of people.
One of the greatest inhibitors of the deployment of digital video information systems has been that many technologies, such as capture, encoding, application design, hosting, and distribution have been developed as individual technologies, but have rarely been integrated in systems that have sufficient functional breadth and operational stability to evaluate the individual technologies in a broader context. eSeminar is an integration point to evaluate advanced media technologies in an operational setting that is to be used in the everyday lives of people. More specifically, the eSeminar objectives are:

- Create a video library of events for everybody at any IBM Research site to watch on-demand.

- Support a large number of videos for on-demand viewing .
- Support a large number of simultaneous on-demand streams.
- Transmit videos of events live as appropriate
- Make slides and other collateral data available where possible.
- Support sophisticated indexing tools and allow users focused access to the material: turn video from a sequential medium into a direct access medium.

eSeminar is operational at the Watson, Zurich, Austin, Almaden, and Tokyo sites, with work ongoing to bring the Beijing and Haifa labs into the system. Each site has a VideoCharger server to stream the videos locally on the LAN, and a web site to provide access to the videos and to collateral information. The eSeminar system is comprised of three workflow stages:
Production: The videos are being recorded on portable video servers in MPEG-1. In a specially equipped conference room, we use an automated camera management system for operator less video recording. When appropriate, speakers' slides are captured as images and are synchronized with the video. Our next research target is to automate the production and recording to the point that a user has only to fill some information into a web form, click a button, and the recording proceeds automatically without further human intervention.
Postproduction: The videos are transcoded into low bit rate videos (MPEG-1, CIF, 10 fps, 200 kbps video, 64 kbps audio). This low bit rate is required to avoid overloading the LANs at the various sites. Next, we are using scene change recognition to create thumbnails of video scenes for directly accessing scenes of the video. Then, we are integrating thumbnails of recorded slides into the video story board. Once an MPEG-1 file exists, the story board creation and the web page generation are well automated. We are now looking to integrate other indexing methods such as audio-based speech recognition and keyword indexing. After indexing off-line is complete we expect to move to on-line indexing. This will allow users instantaneous access to meta data. For instance, somebody who joins a talk late can use the capture data to catch up.
Distribution & hosting: The videos, the collateral material, and the related web pages are transmitted to the other sites using ftp. The content of the remote servers is managed from a central site. We have prototyped a content distribution management system that uses usage frequency as well as storage constraints to manage the media on distributed media server. We expect this system to be integrated into eSeminar to automate content management.
This is a brief overview of the current function, and of our near-term research objectives. In addition, we will present our experience in operating the system, and we will discuss problems we are currently addressing.

Summary: Martin G. Kienzle detailed the development of the eSeminar project. One of the many goals was to integrate and use existing technologies and network restrictions to create a highly automated indexing and distribution of recorded video. Kienzle's system uses a central server to distribute video data to slave servers that reside in a viewing LAN. In addition much work has gone in to tie in existing technologies to offer automated slide shots of the presentation medium, and indexing of keywords.

11 Acknowledgements

The workshop was held under the auspices of the Special Focus on Next Generation Networks Technologies and Applications and the Special Year on Networks. DIMACS thanks the National Security Agency, the New Jersey Commission on Science and Technology and the Office of Naval Research for support of this workshop.