Streaming Media Performance Across the Internet

Ben Odom
CSE 581
1/28/2001

Paper Overview

- Real Media Client/Server
  – An Empirical Study of Real Video Performance Across the Internet
    - Video
  – An Empirical Study of Real Audio Traffic
    - Audio
- Motion Picture Experts Group MPEG-4
  – Measurement Study of Low-bitrate Internet Video Streaming
    - Video

Streaming Video Characteristics

- Streaming Video
  - Sensitive to delay and jitter
  - Not sensitive to small amount of data loss
  - Prefers steady data rate as opposed to bursty type data stream
  - Often uses UDP as a transport rather than TCP
    - Suggest that may not be TCP friendly
    - Unresponsive to network congestion

Measuring Performance

- Video traffic is typically a long lived stream
  - Allows for in stream adjustment
- Dropped packets is not an accurate measure of performance
  - Recovery techniques can hide this
- Quality
  - Frame loss is a good measure of performance
  - smoothness and clarity of an image are what is important in video

Real Video Server Overview

- Content is created using various codecs
- Converted to Real's proprietary format for streaming
- Connections & Protocols
  - Primarily RTSP
  - Can wrap in HTTP for firewall clients
- Uses two connections for communication
  - Control ideally TCP
    - Stop, play, FF, REW
  - Data typically UDP
    - The actual clip

Playback Schemes

- Buffering
  - First 13 seconds
  - If buffer empties stream can stop for up to 20 seconds
- SureStream
  - Video is encoded at multiple bit rates and can switch mid-stream
- Special recovery packets
  - Try to make up for inter-frame packet loss

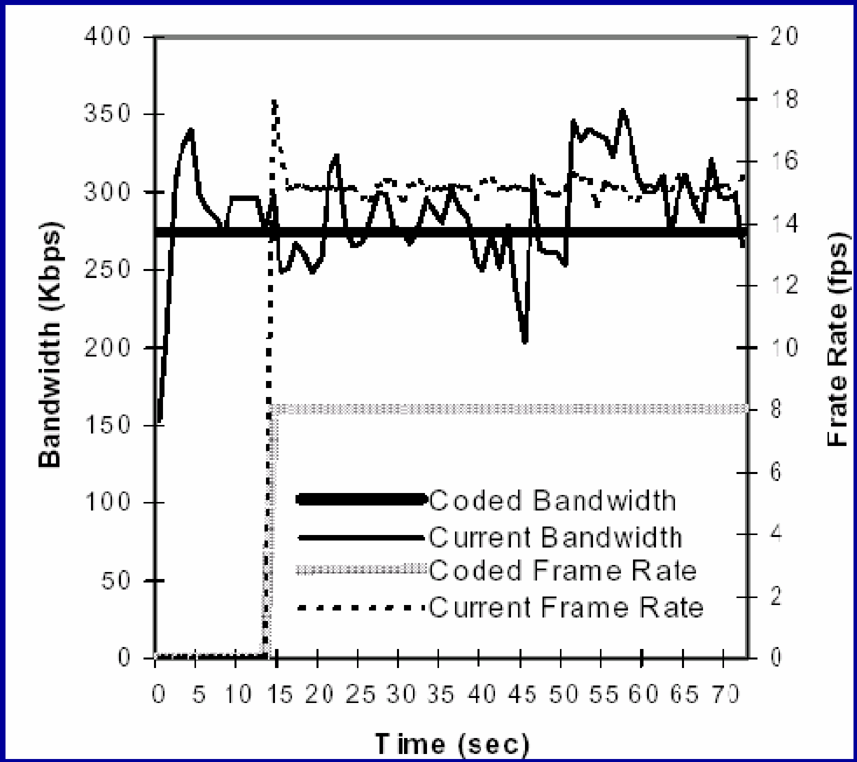

Test Setup

- Used a custom player called RealTracer
  - Used the Real Media SDK
  - Plays the media and records performance statistics
- Chose geographically diverse websites for both servers and clients
- Gathered volunteers from all over

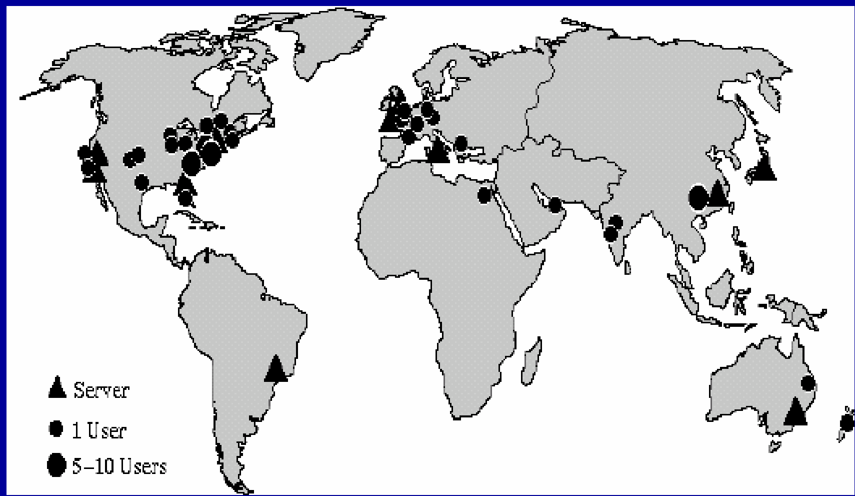

Quality Analysis Frame Rate

- Frame rate = smoothness
- Focused on 3 key rates
  - 3 fps, 15 fps, and 25 fps
- Quality can still be poor if frame rate is high but jitter is also high
- Average frame rate for modems get a proportionally higher frame rate for their available bandwidth
- High speed links operate at capacity less than 10% of the time
- Bottleneck not the end user

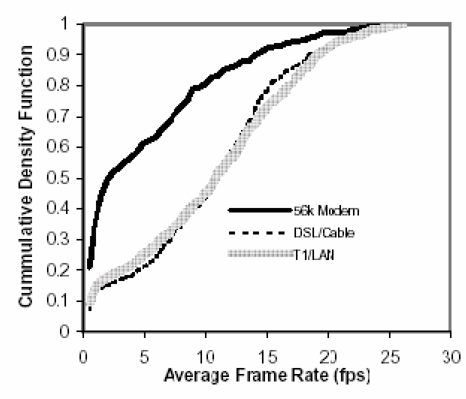

Figure 11. CDF of Frame Rate for Different End-Host Network Configurations

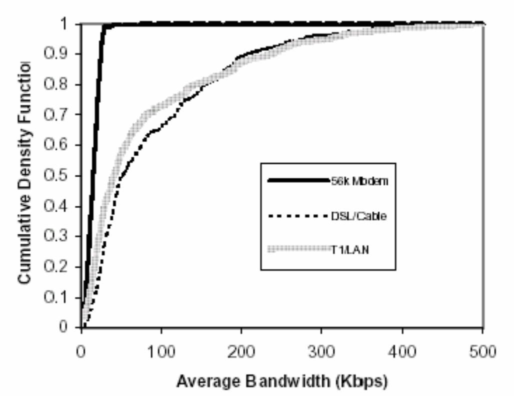

Figure 12. CDF of Bandwidth for Different End-Host Network Configurations

Quality Analysis by Geography

- Server side is almost identical in all geographies
- Client side in different geographic locations show variances
- Australia/New Zealand
  – 75% < 3 fps
- North America, Europe and Asia
  – All about the same at > 20 fps

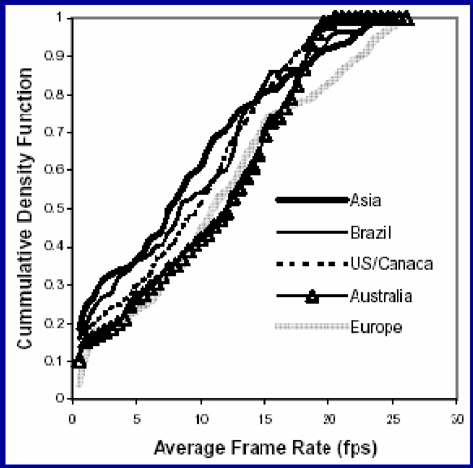

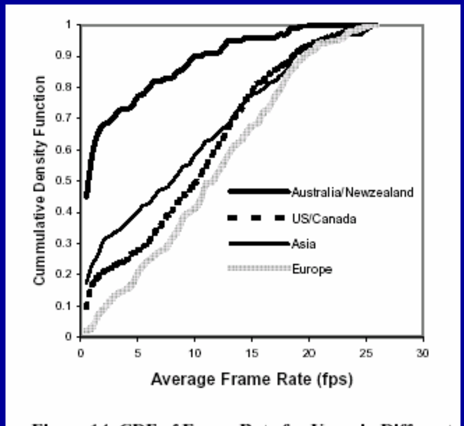

Figure 14. CDF of Frame Rate for Users in Different Geographic Regions

Analysis Jitter

- Jitter
  - When packets arrives late often too late to be used for scheduled decoding time
  - 50ms or less is smooth > 300ms is rough
  - Only about 15% play with > 300ms jitter due to buffering

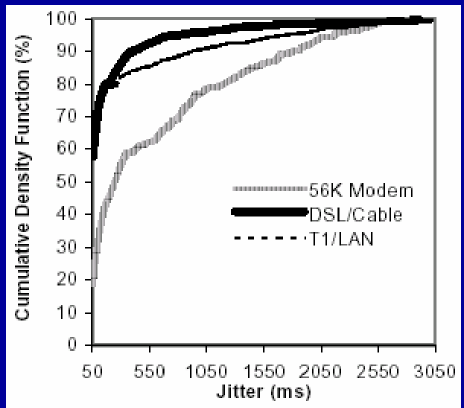

Figure 18. CDF of Jitter for Different Network Configurations

Figure 20. CDF of Jitter for Users in Different Geographic Regions

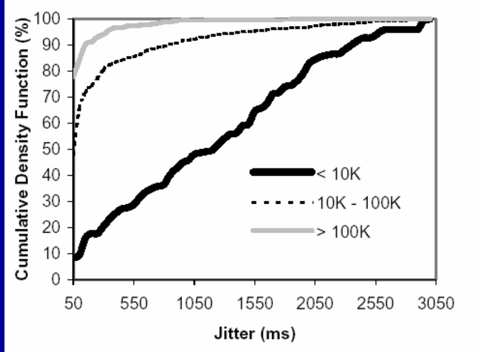

Figure 21. CDF of Jitter for Observed Bandwidth

Perceptual Quality

- Subjective measurement
- Hope to determine a relationship between system measurement of frame rate and jitter with perceptual quality
- Turns out DSL/Cable had a better perceptual quality probably due to less contention
- Measured
  - Geographic location, network configuration and PC type
- Not Measured
  - Playback length, time of day, video clip encoding bandwidth

Figure 22. CDF of Overall Quality

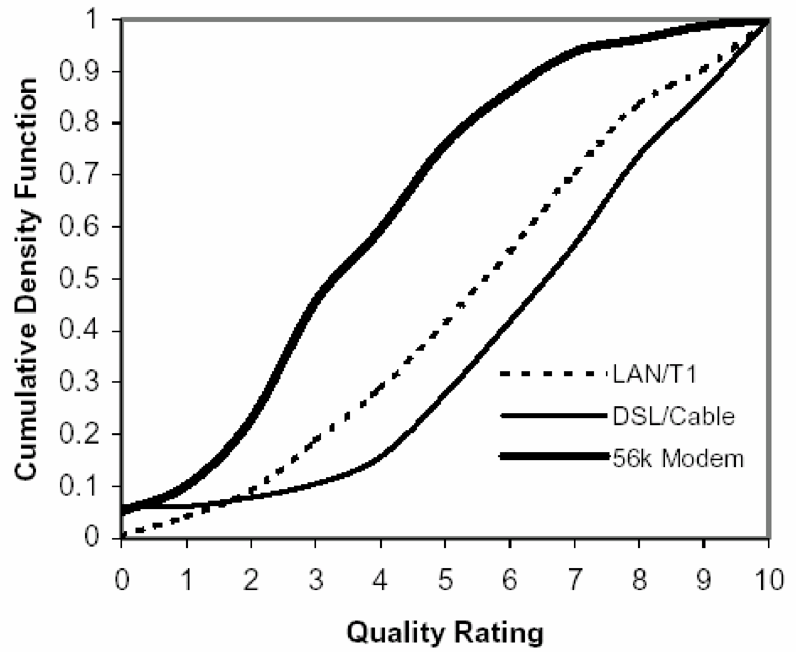

Figure 23. CDF of Quality for Different End-Host Network Configurations

Conclusion

- Empirical measurement of streaming media will be helpful to future development of Internet
- Typical streaming performance is 10 fps aided by delay buffer
- Home broadband connections typically have better performance than corporate LANs
  - This is pushing the bottleneck away from the end user
- Serving streaming media from the diverse countries is about the same
  - Receiving streaming media yields distinct performance differences
    - Dependent on infrastructure

Premise of Study

- Feel that end users are responsible for a large fraction of Internet traffic
- Want to better understand the experience of the masses—feel that other experiments often only dealt with high-speed link bottlenecks
- Undertake large scale study that involves 600 US cities
  - 5000 distinct routers
  - 16 thousand ten minute streaming sessions
  - All over 56K V.90 modems

Server-Client Topology

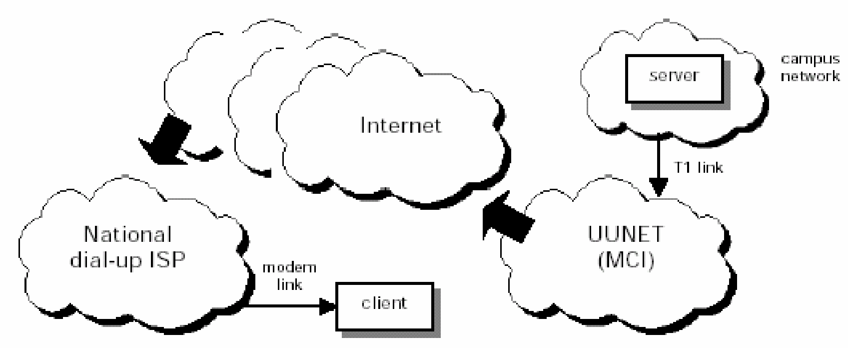

Figure 1. Topology of the experiment.

- Video server was attached to UUNET backbone via T1
- Used 3 nation-wide ISP's for dialup each had 500 V.90 modems
- Dialed long distance to each cities local access number
  - Clients were in lab
  - Assumed infrastructure was mostly digital thus having little effect on connection*
- Implemented a special dialer that used ICMP probes to record the TTL field of each returned TTL expired message
  - This is how they discovered routers and determined if call was clear enough to start streaming session

Real-time Streaming

- Used MPEG-4
  - Created QCIF (176x144) video streams
    - 14 kb/s (1.05 MBytes)
    - 25 kb/s (1.87 MBytes)
- Server split streams into 576-byte IP packets
  - Frames were started on packet boundry
    - With overhead streams were 16 and 27.4 kb/s
- Implemented a strict real-time decoder if packet arrived late underflow even registered
- Used a 2700 ms delay budget for all experiments
- Constant Bit Rate (CBR) enabled all of these hard deadlines to be implemented

Server-Client Architecture

- Server
  - Strictly MPEG-4 using selective negative acknowledgement (NACK) for requests from client
  - Streaming was implemented in bursts of packets 340-500 ms to make overhead of server as low as possible
- Client
  - Designed to recover lost packets through NACK requests
    - Sending NACK and receiving re-transmission
  - Collected all of the RTT data 100 ms resolution

Experiment Success

- Client group 1 performed 16,783 long distance connections
  - Completed 8,429
  - 962 dialup points 637 cities
  - Average hop count 11.3
- Client group 2 performed 17,465 long distance calls
  - Completed 8,423
  - 880 dialup points 575 cities
  - Average hop count 11.9
- Two dialing attempts required
- Discovered 5266 routers
  - DNS and WHOIS lookups done offline

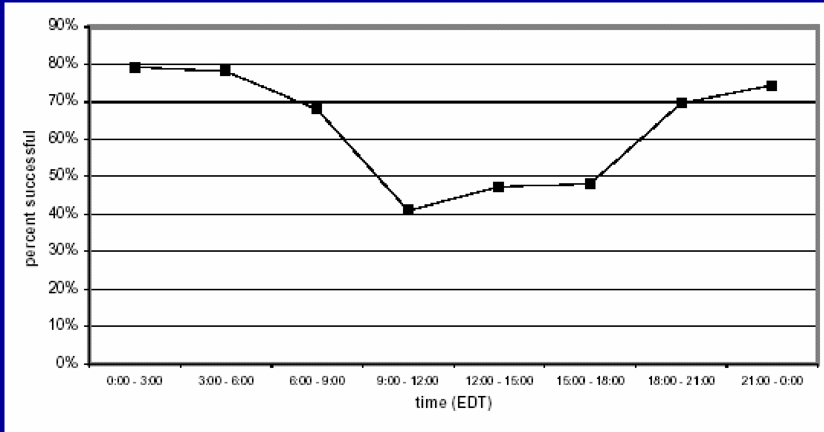

Figure 2. Success of streaming attempts during the day.

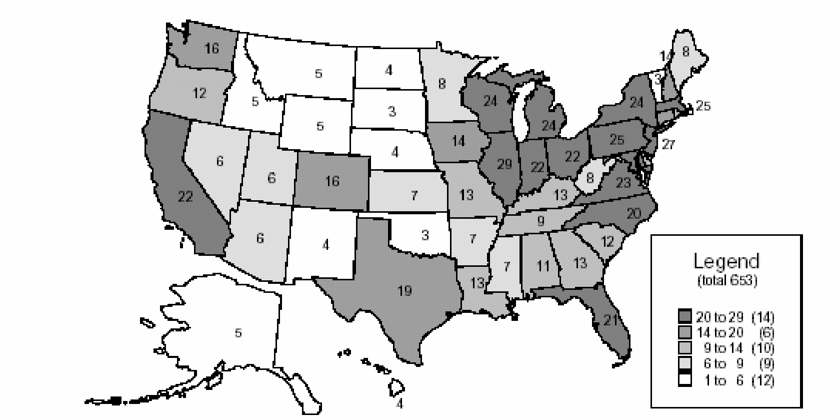

Figure 3. The number of cities per state that participated in either D_1 or D_2.

Packet Loss

- Sources cited by authors suggest 11- 23% packet loss average for Internet
- This experiment lost .53% and .58% respectively at the two bitrates
- 38% had no packet loss
- 75% below .3%
- 91% below 2%
- 2% greater than 6% loss

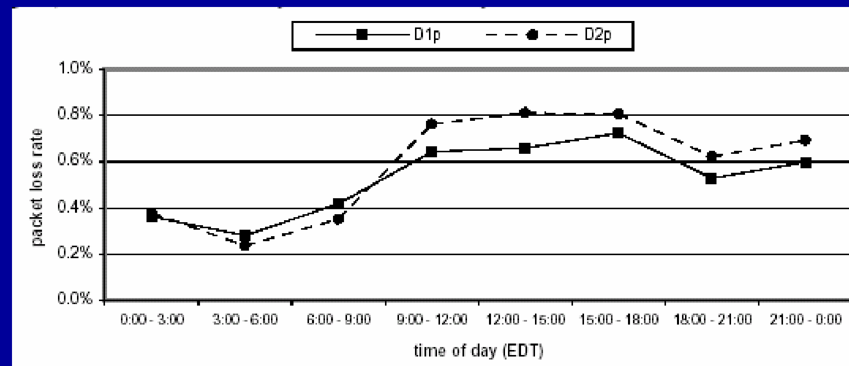

Figure 5. Average packet loss rates during the day.

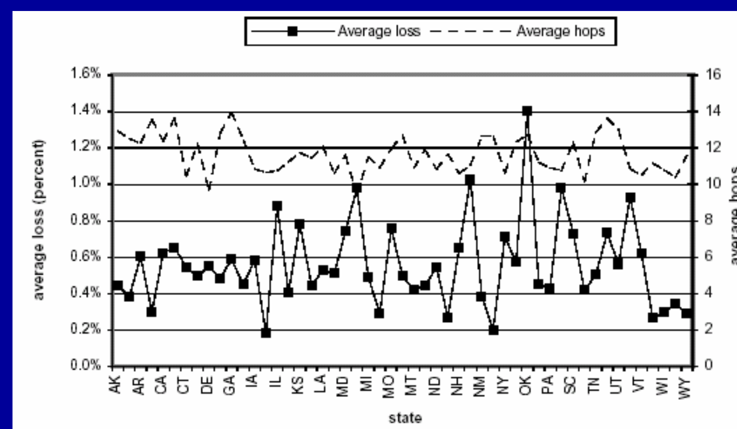

Figure 6. Average per state packet loss rates.

Loss Burst Length/Duration

- How bursty Internet packet loss was during the experiment
- 99% of loss covered by 20 packets
  - 207,384 loss bursts and 431,501 lost packets
- Suspected culprit for loss is overflowing queues in Internet routers
- Lost burst duration 98% under 1 second tails out to 36 seconds
- Future streaming protocols should consider bursty packet losses

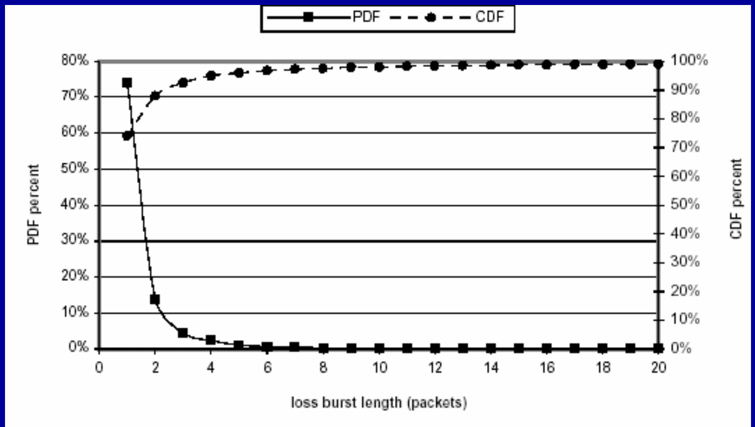

Figure 7. PDF and CDF functions of loss burst lengths in $\{D_{1p} \cup D_{2p}\}$.

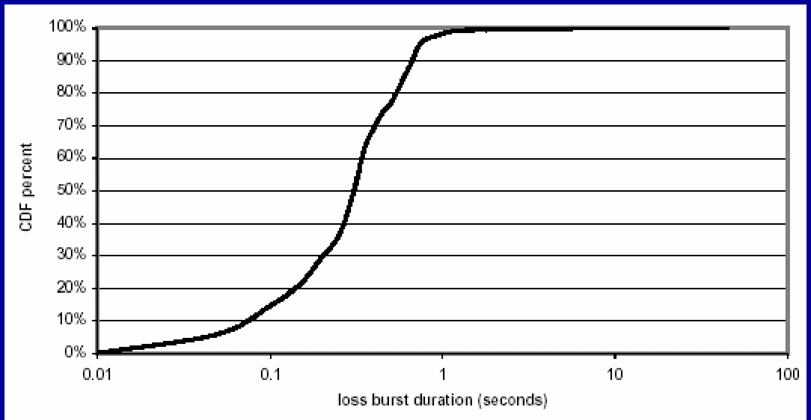

Figure 8. The CDF function of loss burst durations in $\{D_{1p} \cup D_{2p}\}$.

Underflow Events

- Any frame that lost a packet registers an underflow event
  - Did not study error concealment
- Two types of delays
  - High RTT—make packet late for decoding
  - Jitter was responsible for 90 time more underflow events than packet loss and RTT
- 94% of packets recoverable with NACK
  - On average 4.98 frames from the same Group of Pictures (GOP) were recovered

Variation of Round Trip Time

- RTT is not the best measure jitter should probably be used instead as it has the highest impact
- An additional 10.3 seconds in addition to the 2700 ms already budgeted would have accounted for 99% of retransmissions and 84% of completely late packets
- Why didn't they add a larger buffer?

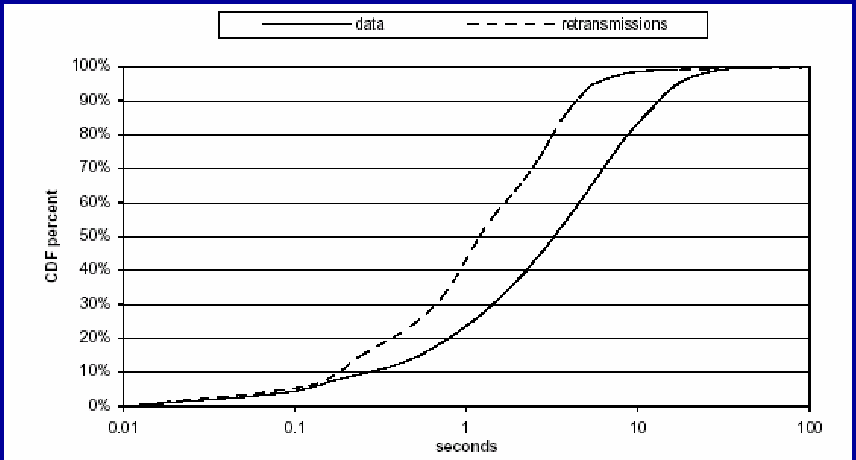

Figure 10. CDF functions of the amount of time by which retransmitted and data packets were late for decoding.

Packet Reordering/Path Symmetry

- Packet re-ordering was in Dataset 1
  - 35% of missing packets were due to re-ordering (0.2% of the number sent)
  - Not dependent on time of day
- Path symmetry
  - ISP implement peering typical user experience will by assymetric
    - Faster down than up

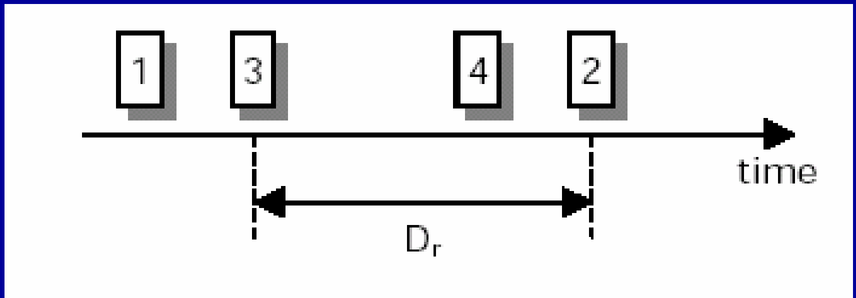

Figure 15. The meaning of reordering delay D_r.

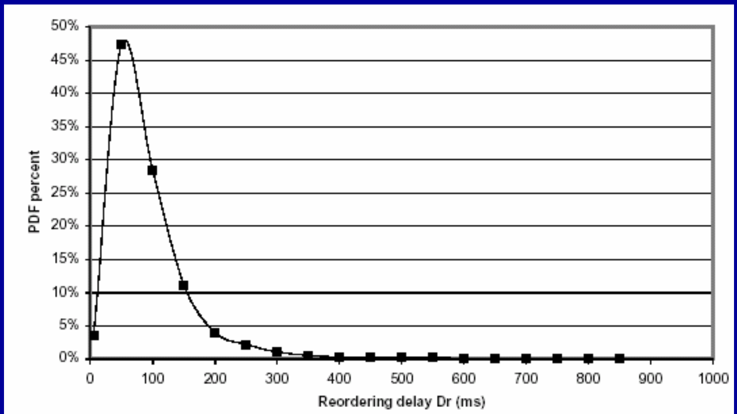

Figure 16. The PDF of reordering delay D_r in $\{D_{1p} \cup D_{2p}\}$.

Conclusion

- Internet packet loss is bursty
- One way delay jitter is more harmful to quality than big RTT
- While there is some re-ordering its effects are small
- Majority of paths sampled were asymmetric

Paper Overview

- Very much like the first paper but without the video
- Demonstrate that audio flows are different than FTP, HTTP, and TELNET sessions
- Method to identify audio traffic
- Preliminary
  - Just started to analyze the data

Methodology

- Connected to the switch at Broadcast.com
  - Used a sniffer and tcpdump on a separate host to make 5 traces
- Traces
  - 1-2 were music and radio streams
  - 3-5 were talk show type streams

Trace	1	2	3	4	5
Date	Mar 99	Mar 99	Jun 99	Jun 99	Jun 99
Start time, GMT	N/A	N/A	16:02	13:32	13:38
Duration	83 sec	141 sec	5.5 hr	10.5 hr	18.2 hr
Packets	134 K	284 K	5.5 M	1.6 M	5.9M
Bytes	38 M	63 M	1.3 G	0.4 G	1.3 G

Table 1. Summary of Traffic Traces

Identifying Audio Traffic

- Outbound:inbound byte ratios
  - 28:1, 40:1, and 50:1 for the three traces
- Inbound data is mostly control data

	Trace 3	**Trace 4**	**Trace 5**
Audio Data	1,160 MB 3.7M packets	403 MB 1.2M packets	1,268 MB 4.6M packets
Control Data	41.3 MB 1.7M packets	10.3 MB 0.3M packets	22.67 MB 1.2M packets
Other Packets	1.0 MB 90K packets	0.8 MB 44K packets	1.0 MB 98K packets

Table 3 Summary of traffic traces

Identifying Audio Traffic

- Highest rate of 20 Kbps is associated with stream 1 & 2
- Traces 4,5, and 6 data rates vary between 1 to 5 Kbps
- Data rates seem to concentrate in the areas of
  - 244/254, 290/300, and 490/502 bytes for 1, 6.5 and 20 Kbps audio flows

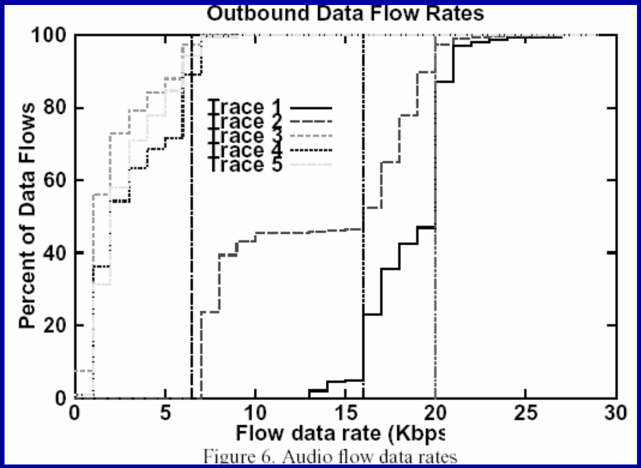

Figure 6. Audio flow data rates

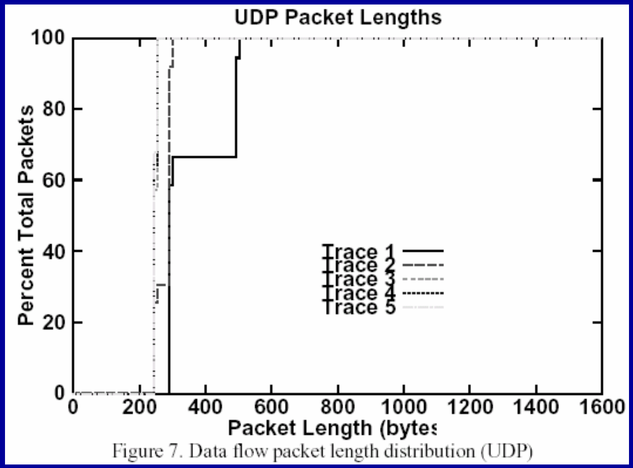

Figure 7. Data flow packet length distribution (UDP)

Not TCP Friendly

- Help develop future routing policies for this type of traffic
- Majority of traffic is UDP
- From previous presentations we know that this does not help the health of a network

	Trace 3	Trace 4	Trace 5
Bytes			
UDP	723 M (60 %)	415 M (79%)	955 M (74 %)
TCP/Non-HTTP	432 M (36 %)	68 M (17%)	304 M (24%)
TCP/ HTTP	47 M (3.9 %)	18 M (4%)	36 M (2%)
Multicast	0	0	0
Packets			
UDP	3.68 M (67 %)	1.26 M (80%)	4.52 M (77%)
TCP/Non-HTTP	1.66 M (30 %)	0.26 M (17%)	1.21 M (21%)
HTTP/TCP	0.14 M (3 %)	0.05 M (3%)	0.12 M (2%)
Multicast	0	0	0

Table 4 Aggregate Traffic

Conclusion

- Find that audio flows differ from Telnet, HTTP, FTP and typically measured traffic
- Audio flows are significantly longer
- By being able to identify these flows perhaps future routing protocols could be designed to address the congestion control issue