The Performance of Two-Dimensional Media Scaling for Internet Videoconferencing

Peter Nee
Kevin Jeffay

University of North Carolina at Chapel Hill

Gunner Danneels

Intel Corporation

http://www.cs.unc.edu/Research/dirt

Overview
Method and Goals

- Experimental videoconferencing application
  - » Based on Intel's ProShare™ version 1.8
  - » Incorporates adaptation of both packet-rate and bit-rate
- Run repeated trials on a real Internet path
  - » 20 hop path between UNC and University of Virginia
  - » Daytime, weekday traffic
- Assess the sustainability of adaptive videoconferences on a lengthy Internet path
- Compare our method with a conventional adaptive method: temporal video scaling

Packet-Rate Vs. Bit-Rate
Two types of constraints

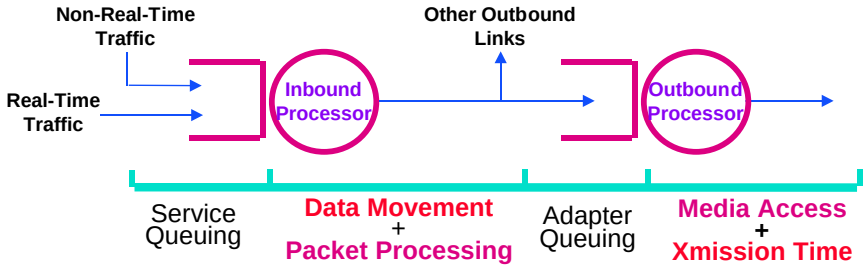

- Capacity constraints
  - » the network is incapable of supporting the desired bit rate in any form
- Access constraints
  - » the network cannot support the desired bit rate with this packaging scheme

Two Types of Constraints
Two dimensions of adaptation

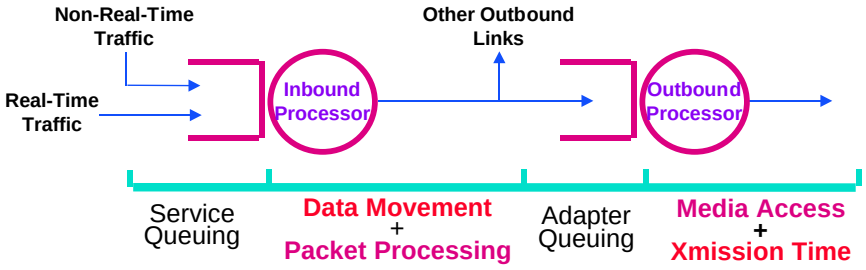

- Adapting to an access constraint
  - » Reduce packet-rate: change packaging or send fewer video frames
  - » Primary Trade-off: higher latency
- Adapting to a capacity constraint
  - » Reduce bit-rate: fewer video frames or fewer bits per video frame
  - » Primary Trade-off: lower fidelity

Example Operating Point Set
Requires Large MTU

- Change video quality to scale bit-rate
- Change video frame rate to scale both packet-rate and bit-rate
- These two adaptations allow coverage of the entire operating point set
- This is a rich operating point set covering a large area in the bit-rate/packet-rate plane

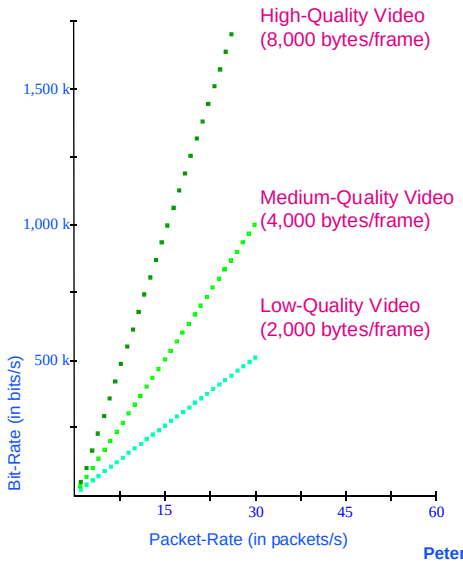

Adaptation Based on Receiver Feedback
Recent Success Heuristic

- End-to-end method
  - » Receiver reports average latency and loss
  - » Sender can detect congestion from the feedback
  - » Sender cannot determine the type of constraint causing congestion
  - » Sender uses a heuristic to find the correct adaptation
- When sender detects congestion
  - » Adapt in the dimension that successfully relieved the last congestion episode
  - » If that fails, try the other dimension
- When sender detects absence of congestion
  - » Wait for a few seconds
  - » Probe in the dimension orthogonal to the last constraint
  - » Back off if probe causes congestion

Example: Recent Success

- Initial operating point (high quality, 12 fps)
- First adaptation: (high quality, 10 fps) congestion persists
- Second adaptation: (medium quality, 10 fps) congestion relieved
- First probe: (medium quality, 12 fps)
- Second probe: (medium quality, 14 fps)

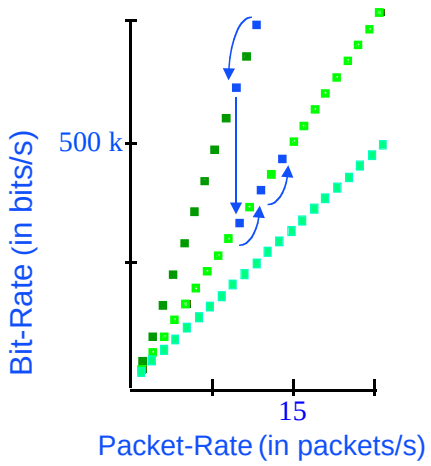

ProShare Operating Points
Ethernet MTU

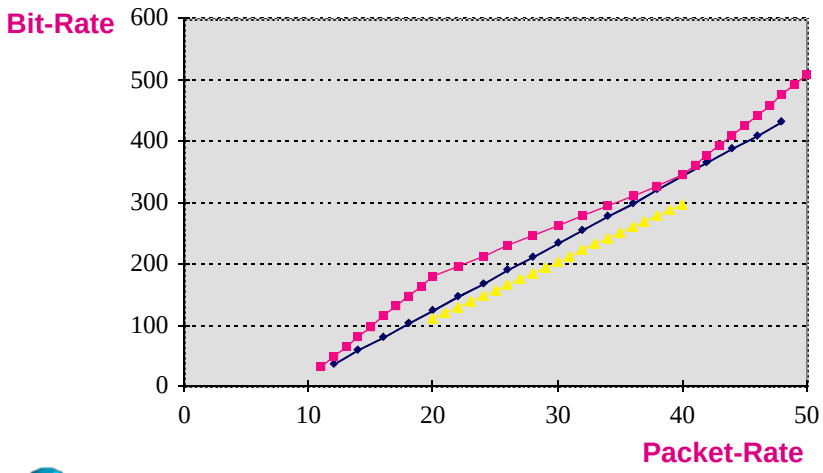

Experimental Setup

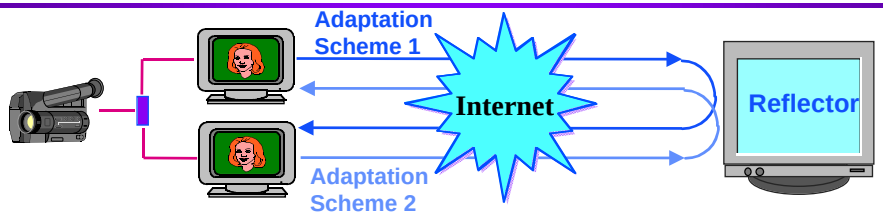

- Bi-directional Audio/Video and feedback
- Head to head comparison of both adaptive methods
- Running over the same path to and from reflector
- Capturing the same video source
- Subjective assessment by human observer
- Objective quality measures logged and analyzed off-line

Results of a Typical Run
Audio and video throughput

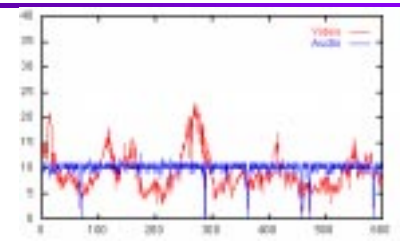

1D Frames Received x Time (in s)

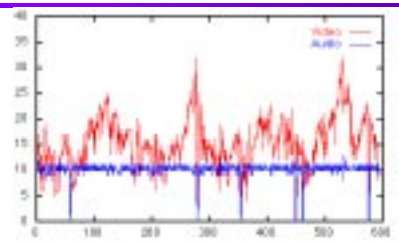

2D Frames Received x Time (in s)

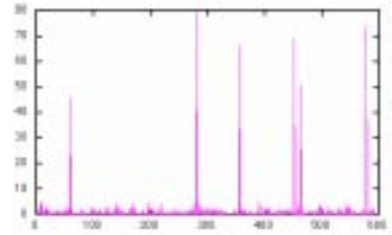

1D Packet Loss x Time (in s)

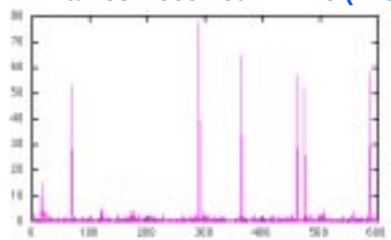

2D Packet Loss x Time (in s)

Sustainability Results
Adaptive methods - Internet traffic

- Used real Internet traffic
  - » Repeated trials from 10 am to 7 PM weekdays
  - » Trials separated by at least two hours
  - » Scattered over three months

Time Slot	Sustainable	Not Sustainable	% Sustainable
10:00-12:00	6	3	67%
12:00-14:00	4	4	50%
14:00-16:00	1	11	8%
16:00-18:00	3	9	25%
18:00-20:00	4	5	44%
Percentage	36%	64%	

Summary and Conclusions

- Conference quality
  - » Two-dimensional scaling delivers more video throughput
  - » Most noticeable in the ten to twenty frame range
  - » In this range, the operating point set is richest
  - » No significant difference in loss or latency between the two methods
- Sustainability results
  - » No significant difference between the two methods
  - » We do not have a rich set of low-end operating points
  - » Sustainability depends primarily on successful audio - same for both systems
- Two-dimensional scaling: better conferences
  - » The richer the operating point set, the greater the win

Future Work

- Improve the audio component
  - » Change the audio codec to allow scaling and aggregation of audio
  - » Experiment with distinct audio and video feedback and adaptation
  - » Provide better audio when network resources are available
  - » Provide same or better sustainability
- Understand the interaction with TCP
  - » Controlled experiments with a population of TCP applications and adaptive videoconferences
  - » Does two-dimensional scaling do enough to avoid congestion collapse?
  - » Is two-dimensional scaling unfair to TCP?

Adaptive Videoconferencing
Network Citizenship

- Adaptive scaling methods do perform congestion control
  - » Avoid packet loss - i.e. wasted resources
  - » Minimize successfully transmitted but useless packets
- Two-dimensional scaling does not behave exactly like TCP
  - » Takes advantage of richer feedback
  - » Has additional application specific concerns
  - » Potentially unfair to TCP users
- Is this sufficient for good citizenship?
  - » Congestion control is an end system issue
  - » Fairness is best addressed in the network

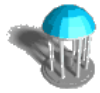

Building an Adaptive ProShare

- Original system
  - » Bi-directional LAN videoconferencing (Ethernet MTU)
  - » Fixed low bit rate audio operating point
  - » Choose one of three fixed video operating points at start
- Experimental system
  - » Add measurement and feedback
  - » Adapt video with internal interface
  - » One-dimensional - change frame rate
  - » Two-dimensional
    - – Change frame rate
    - – Change bytes per frame
    - – Aggregate audio with video tails when possible