No-Frills Realtime Service

Sally Floyd and Lixia Zhang

July'95 IETF

Motivation

- Designed for delay-adaptive & loss-tolerant applications
- Imposes minimal requirements on network components
- Highly efficient implementation and use of network resources
- Intended to be usable in situations ranging from small, centrally managed private IP networks to the global Internet.

A New Generation of Playback Applications

A typical example:

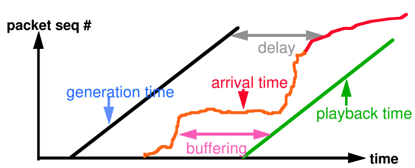

How Applications Set Playback Point

Conventional applications: require guaranteed service

- Ask network for an absolute delay jitter bound & stick to it
- Fidelity: suffer from either large offset delay or otherwise low utilization

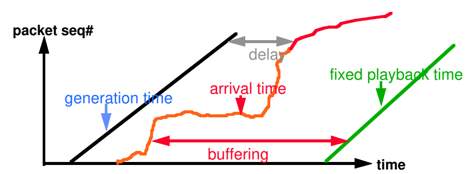

How Applications Set Playback Point (II)

New realtime applications: adaptive & tolerant

- Adjust playback point by observed delay jitter
- Fidelity: much smaller playback delay, but occasional losses due to adjustment errors

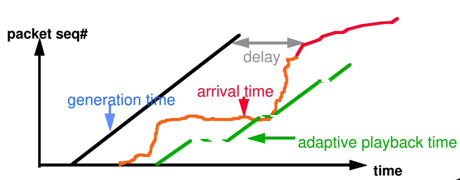

Why Adaptive Applications

Applications can be adaptive

- variation in delay: adjusting playback point
- variation in losses: tolerant to losses (within a limit)

Applications should be adaptive

- more accessible across heterogeneous internet
- more robust in face of unexpected change of conditions

Application adaptivity alone may not be sufficient all the time

- make resource reservations whenever necessary

Resulting service

- a single level of realtime service.
- no quantitative specification for losses or delays

Admission control

- presumably a number of schemes will be applicable
- one possibility: admission decision based on measured utilization
  - + we'll start experimentation with this one first
- encourage others to explore different approaches

Evaluation criteria

- first step in field trial
- collect experience to determine our next step

Should Controlled Delay Service template be used to implement No-Frills Service ?

Yes it can be done

- falls within the CD specification

2 noticeable differences:

- using single level only
- no advertising

Single Level as a Starting Point

Single level seems a simpler first step for experimentation

- answer the question: do we really need more than one delay levels?
- we believe there are complications in providing a multi-level delay service that require further exploration
  - individual routers may set different delay targets
  - when to dial down?
- once we better understand the need and the complication, it would be easy to modify the service definition to add additional levels of service.

Advertising or Not

- It is highly desirable to be able to collect performance statistics
- This can be achieved by following the standard step of MIB definitions
  - desirable to be able to query individual hops along the path
  - as it's currently defined, OPWA provides only the cumulated performance information of the path
  - advertising has little use in one level realtime service; may revisit the issue if move to multi-levels
- Learn from experimentation what MIB entries are needed