VoN Performance Measurements

Jonathan Rosenberg
Bell Laboratories, Lucent Technologies
Columbia University

Overview

- Goals of Measurements
- Measurement Procedure
- Results
  - Network Characteristics
    - Round Trip Times
    - Burst Loss Length Distributions
    - Conditional Loss Probabilities
    - Loss Sample Path
    - Delay Sample Path
  - Application Viewpoint

Goals of Measurements

- Understand Internet behavior as it affects voice
- Characteristics of Internet service
  - Loss
  - Delay
  - Jitter
- Characteristics of Voice Quality
  - Round trip delays
  - “Speech Quality”

Correlating Internet and Voice Quality

- **Some important observations:**
  - **Applications have** *playout buffers*
    - **Extra delay is inserted before first packet is played out**
    - **Packets which arrive after their playout time are lost**
    - **As network delay varies (*jitter*), playout buffers must adapt**
    - **Many algorithms exist for doing this**
    - **These algorithms introduce additional loss and delay which affects perceived quality**
  - **Modern speech coders (G.723.1, G.729, SX7300, etc.) use heavy compression**
    - **Heavy reliance on past state**

- **Losses cause state divergence between encoder and decoder**
- **Result is a *persistence* of errors for some time after a loss**
- **Most speech coders have built in recovery mechanisms to extrapolate missing frames**

Speech Coder Performance

- **Experiments run on G.729 simulator with standard speech segments**
- **Errors introduced into frame sequence before decoding**
- **Question: How long after n consecutive frame losses are errored and error-free decoder output different?**

Burst Size	Mean Convergence Time (frames)	Standard Deviation
1	10.26	7.92
2	8.21	6.51
3	8.22	7.39
4	7.27	6.93
5	9.25	9.39

Speech Coder Performance

- Question: How well does the frame erasure recovery algorithm work in repairing an isolated burst of n frame erasures?
- Informal subjective evaluations among a small set of non-experts:

Burst Length	1	2	3	4	5
Mean MOS	4.21	3.28	2.43	2.08	1.7

- General observation: recovery from single loss works, double losses is noticeable, degradation gradually decreases from there with increasing burst size
- Some rough MSE computations back this hypothesis

Choosing Metrics

- **Include playout buffer algorithm in measurements**
- **Since speech coder performance suffers from "close together" losses, we consider:**
  - *Burst Length Distribution* - **frequency of n consecutive packet losses for different n**
  - *Conditional Loss Probability* - **probability of packet $n+i$ being lost, given packet n is lost, vs. i (should be constant if losses are independent !)**
  - *Loss Sample Path* - **loss probability in non-overlapping windows of time**
  - *Misplay Probability* - **probability that a packet is either lost, or within 10 frames of a loss**
- **Delay Metrics:**

- *RTT Distribution* **- probability that the round trip time** t **is less than** t_0**, as a function of** t_0
- *RTT Sample Path* **- average round trip time in non-overlapping windows**

Collection Procedure

- Endpoints at:
  - Columbia University, New York
  - GMD Fokus, Germany
  - University of California, Santa Cruz
  - University of Massachusetts, Amherst
- Packets are generated as if using G.723.1 at 6.3 kbps, one frame per packet, plus RTP/UDP/IP
- Packet stream is either one way or round trip
- Send and receive times are noted; packets contain sequence numbers for loss determination

Network RTT

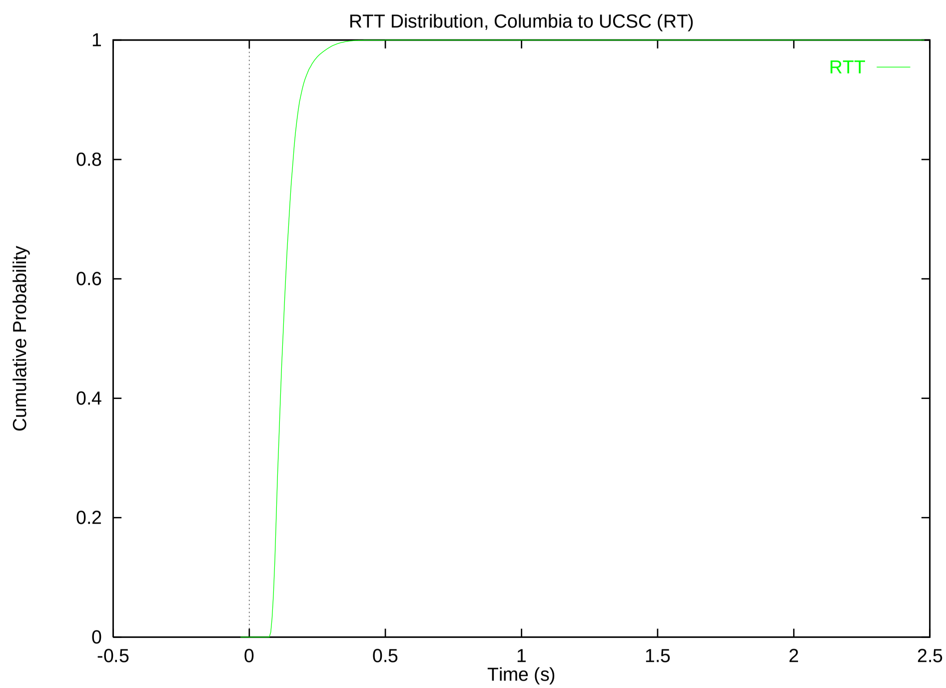

RTT Distribution, Columbia to UCSC (RT)
Cumulative Probability
Time (s)
RTT
0
0.2
0.4
0.6
0.8
1
-0.5
0
0.5
1
1.5
2
2.5

RTT Sample Path

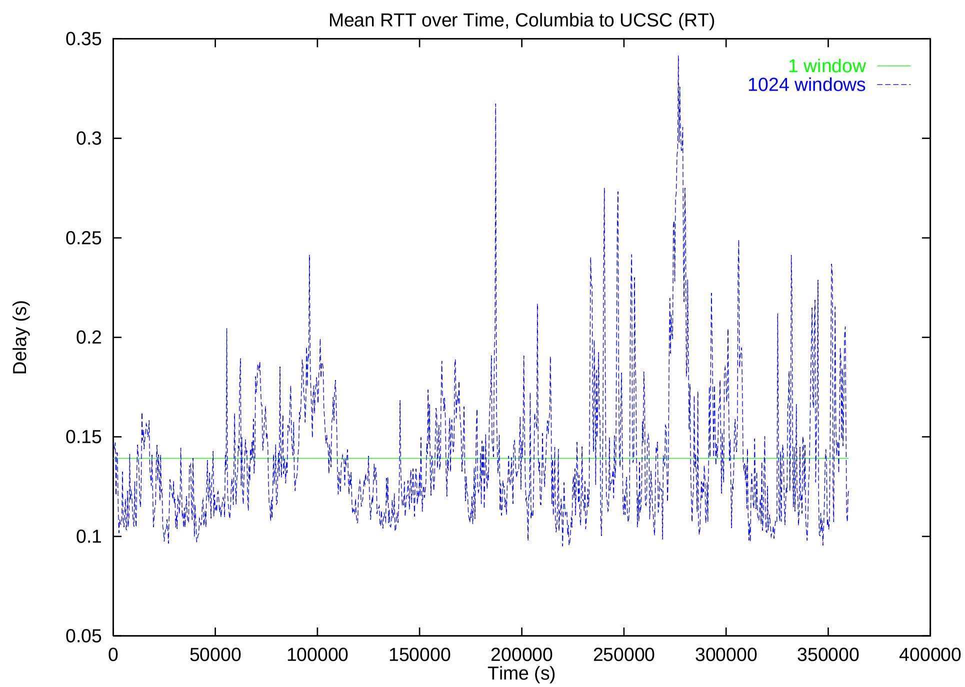

Network Burst Length Distribution

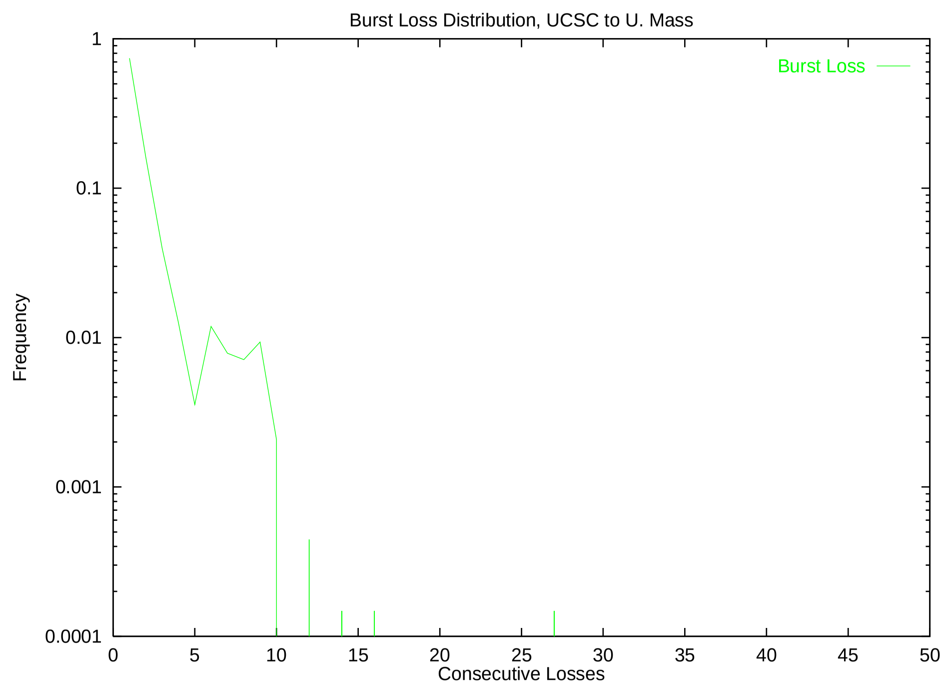

Network Burst Length Distributions

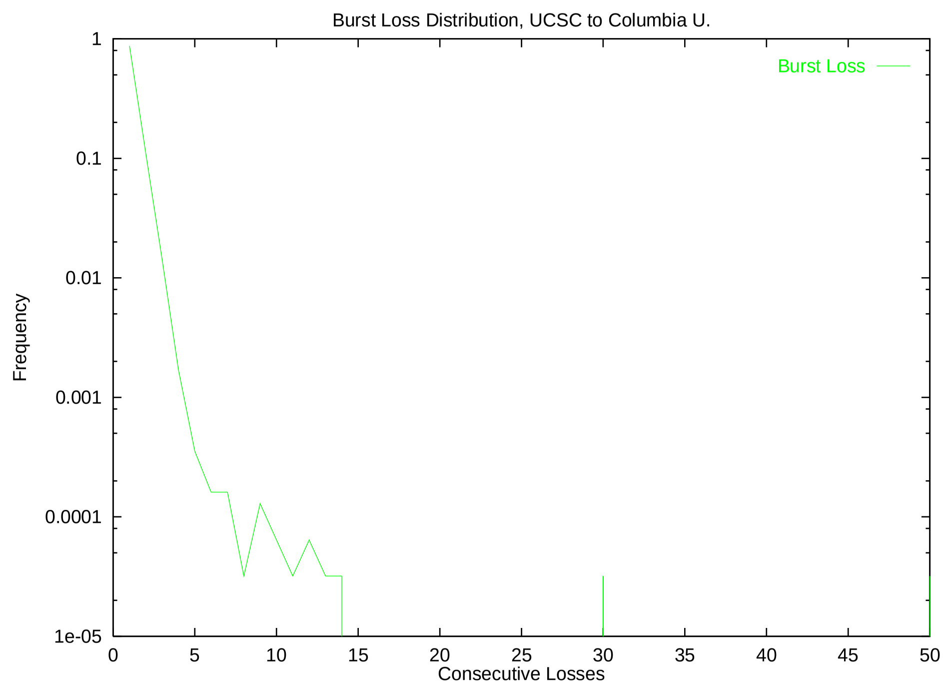

Largest burst is 1206 packets!

Network Burst Length Distributions

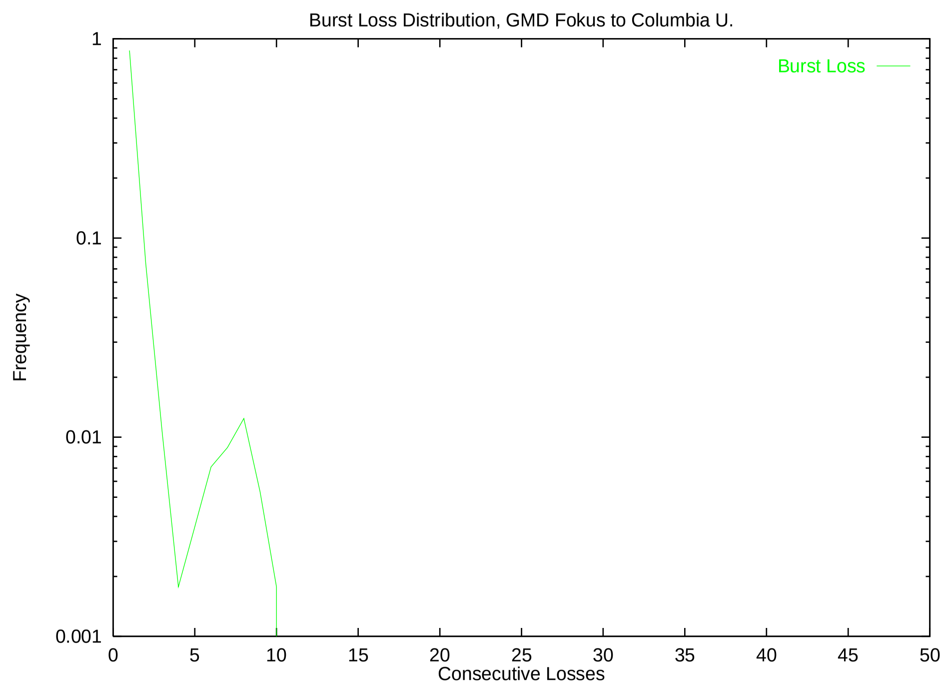

Network Conditional Loss Probabilities

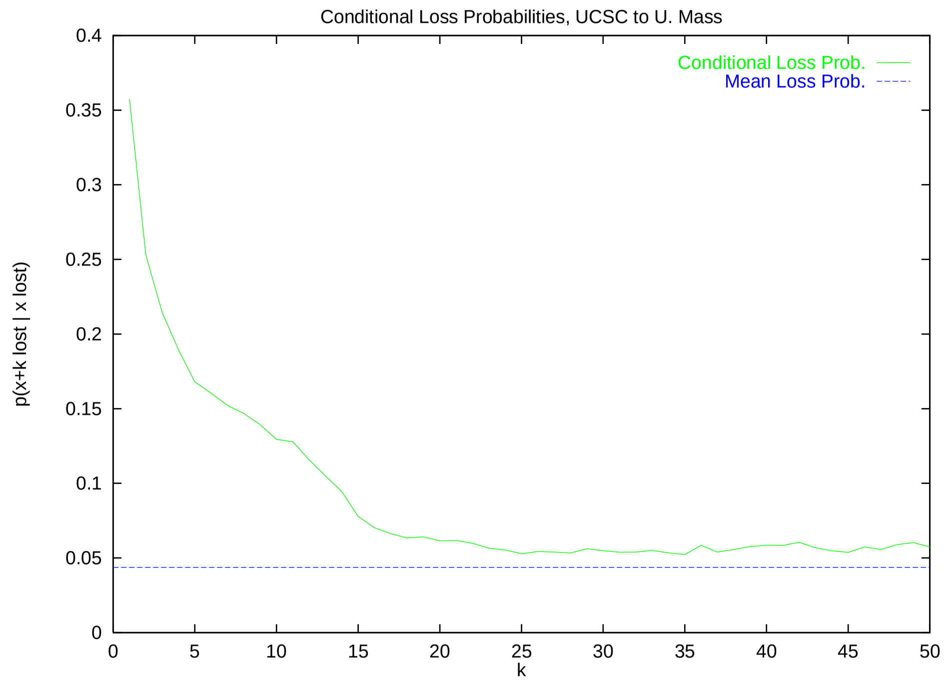

Network Conditional Loss Probabilities

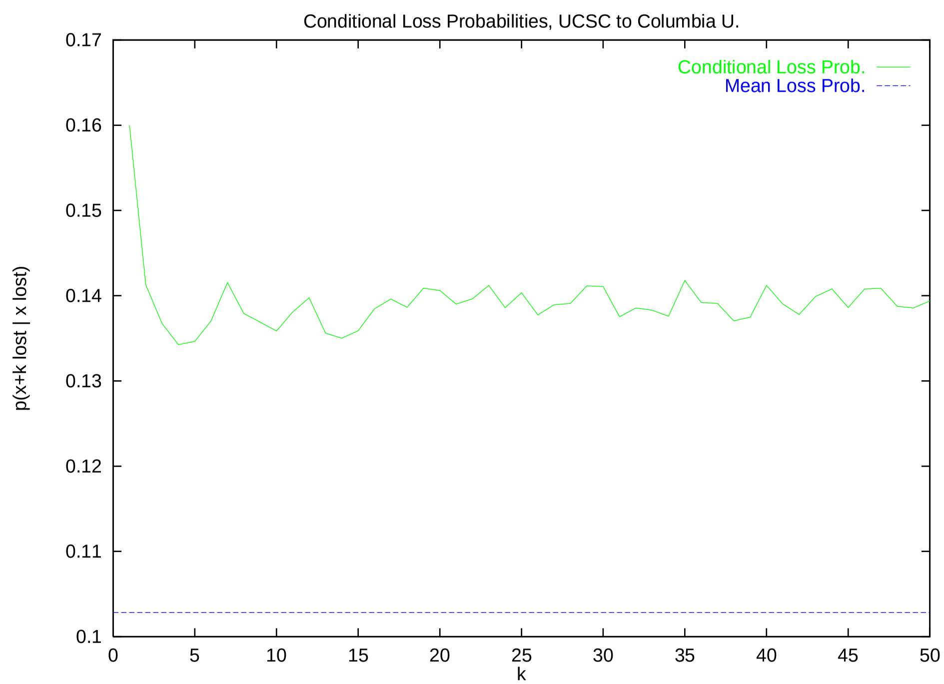

Network Conditional Loss Probabilities

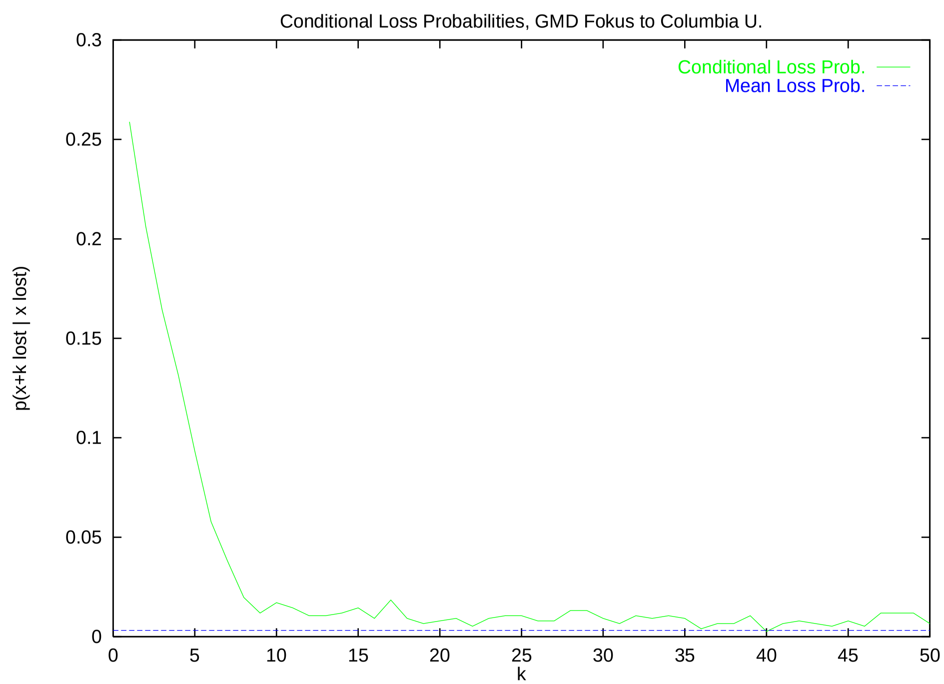

Network Loss Sample Path

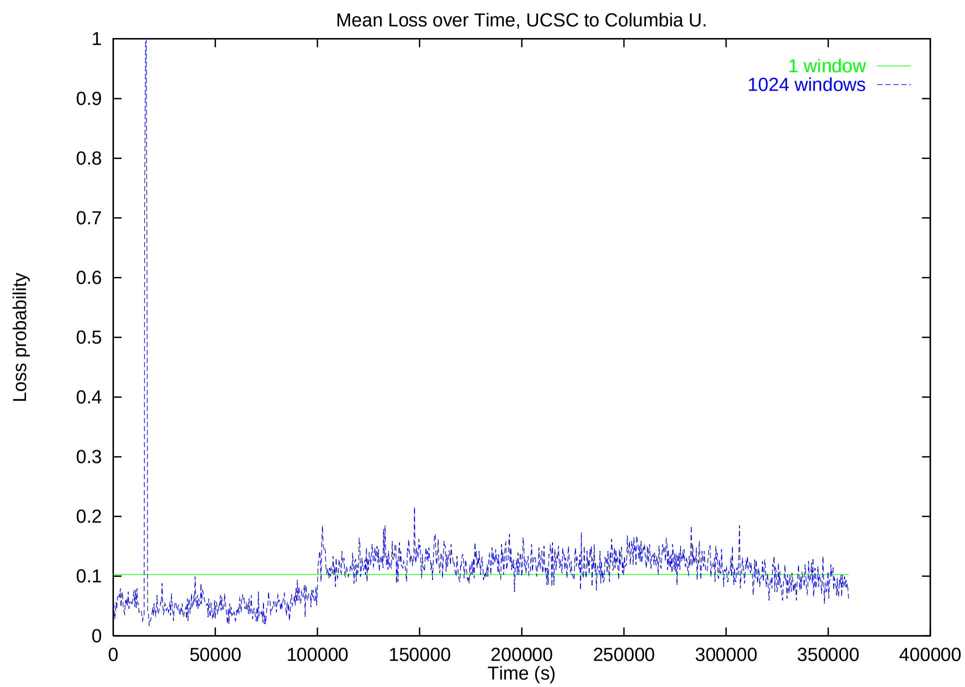

Network Loss Sample Path

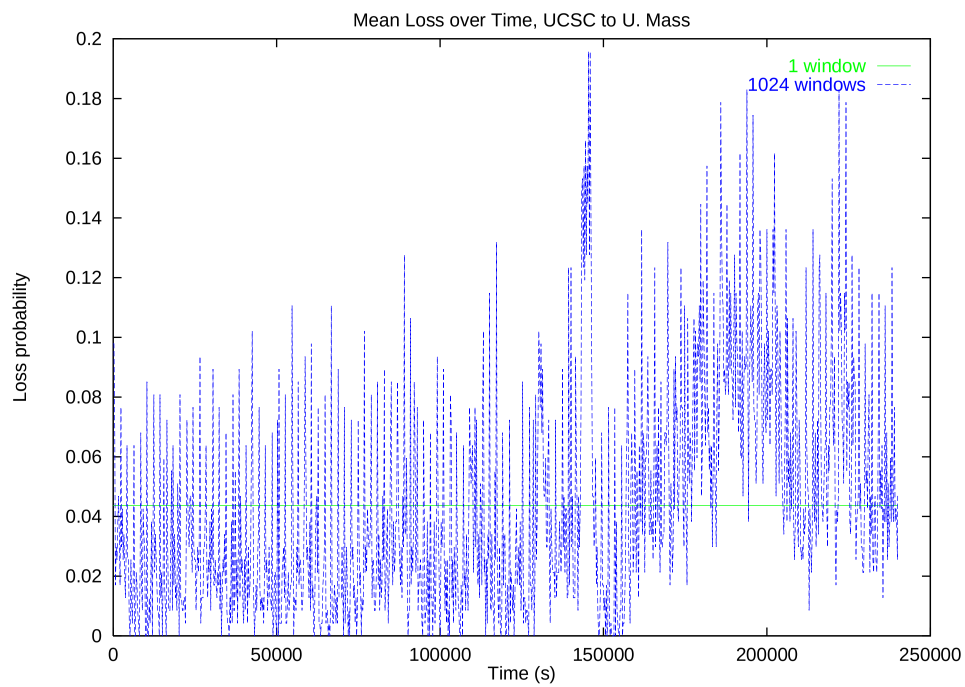

Misplay Probabilities

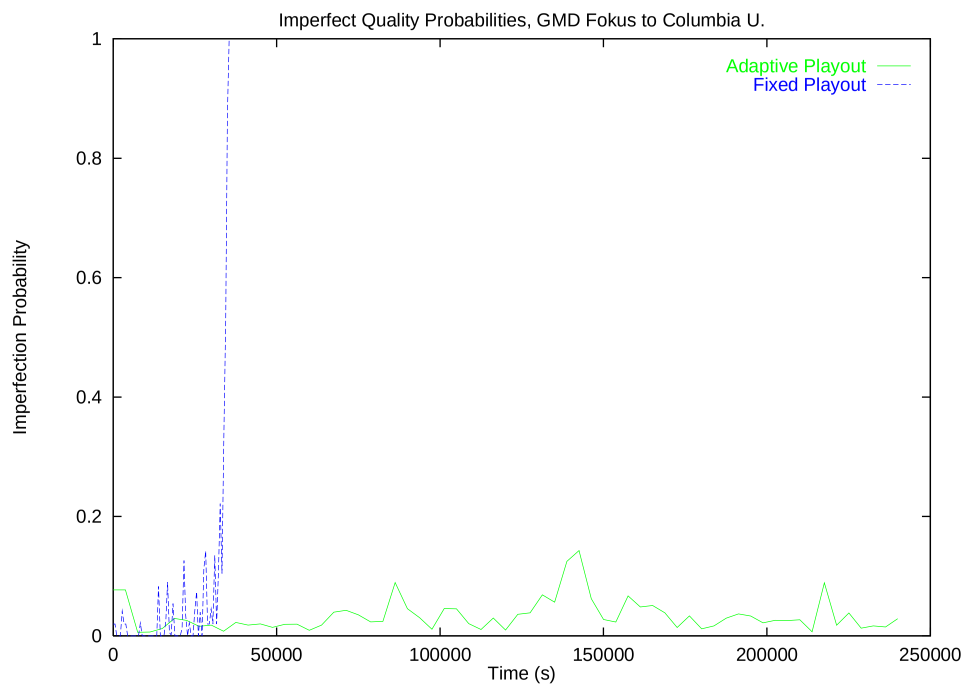

Misplay Probabilities

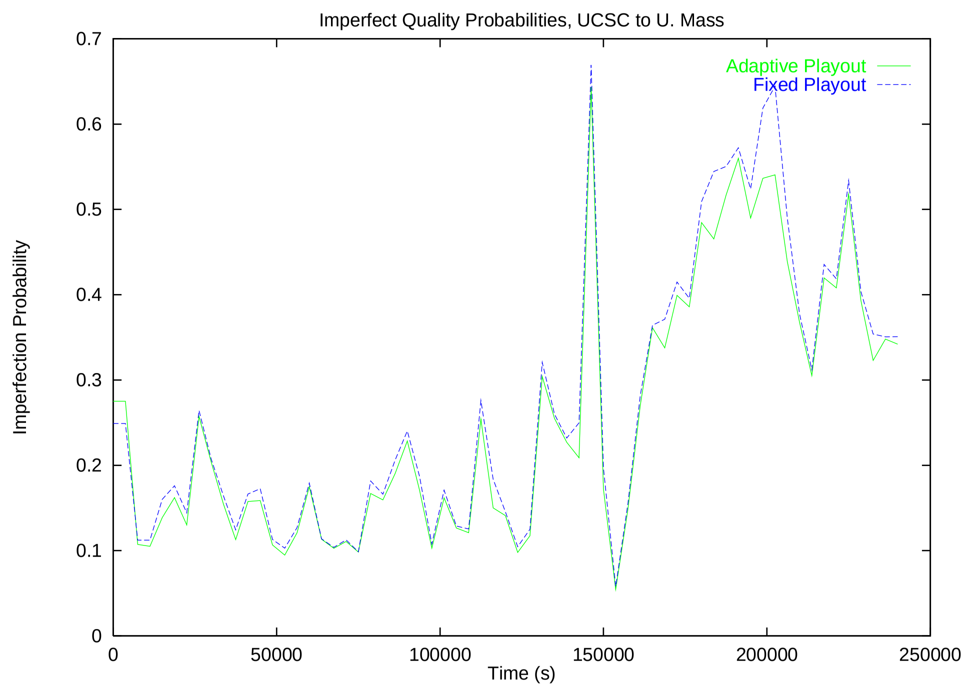

Conditional Loss Probabilities

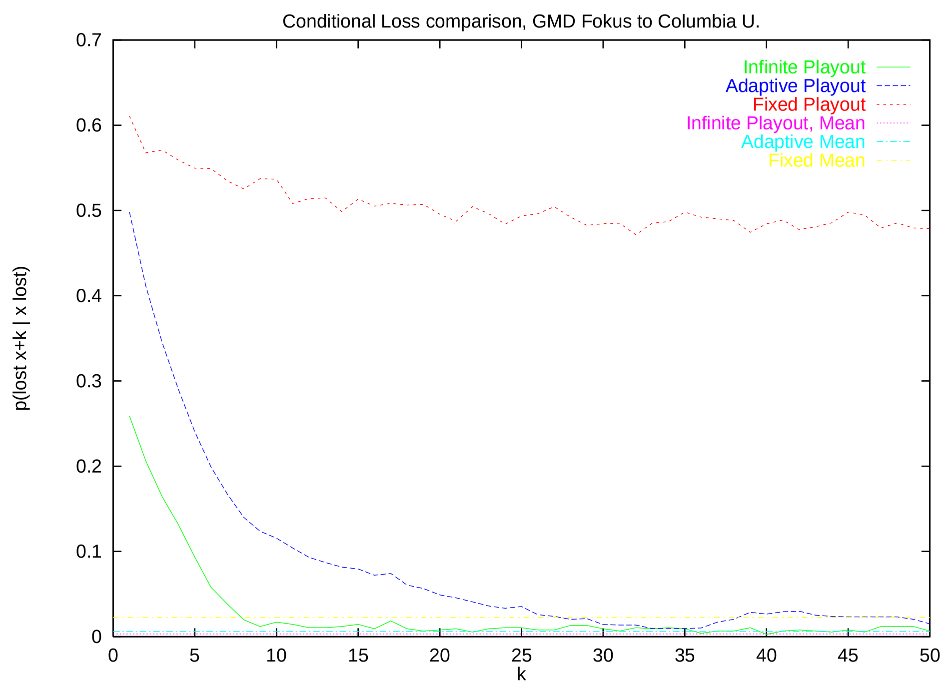

Burst Loss Comparison

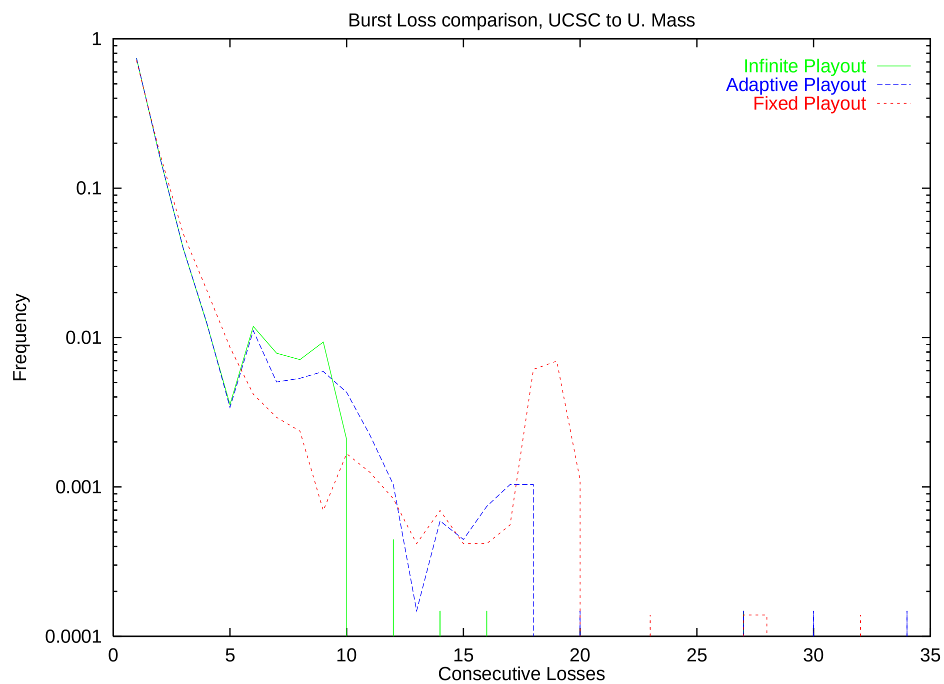

Burst Loss Comparison

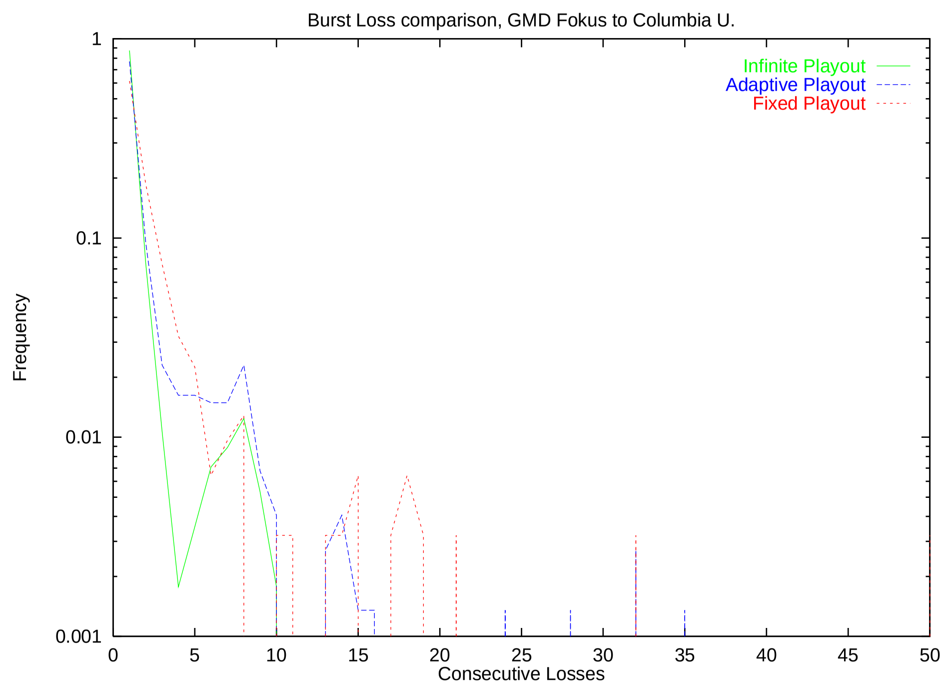

Delay Distribution Comparison

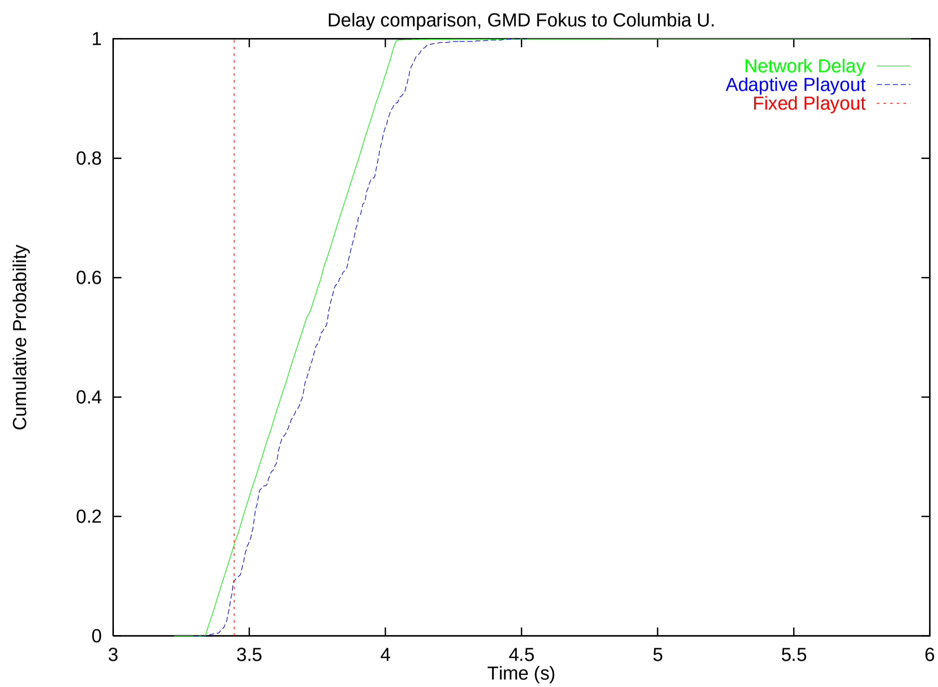

Conclusions

- Playout buffers have an effect on loss process seen by application - increasing the mean and small lag correlations
- Burst length distributions tend to fall off exponentially at first, but with “hiccups” at the tails, increasing overall burstiness
- Conditional loss probabilities verify that network losses are not independent for voice
- Recovery mechanisms and algorithms should target correlated losses - isolated losses are handled well by decoder recovery