A Survey of Packet-Loss Recovery Techniques

Colin Perkins, Orion Hodson and Vicky Hardman
Department of Computer Science
University College London (UCL)
London, UK

IEEE Network Magazine
Sep/Oct, 1998

Overview

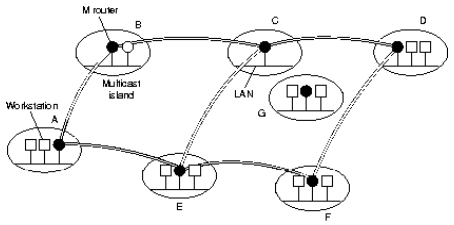

Fig. 7-96. MBone consists of multicast islands connected by tunnels.

- Development of IP Multicast
- "Light-weight session"
  - Scale to 1000's of participants
- How to handle packet loss? (MLC: why doesn't retransmission work?)
  - Repair techniques *beyond* retransmission

Overview

- This paper:
  - Loss characteristics of Mbone
    - (MLC – dated, but not dissimilar from some P2P networks and ad-hoc wireless networks)
  - Techniques to repair loss in a 'light-weight' manner
    - Concentrate on audio
  - Recommendations
- Other papers:
  - Fully-reliable (every bit must arrive), but not real-time
  - Real-time, but do not include receiver based approaches

Outline

- Overview
- *Multicast Channel Characteristics*
- Sender Based Repair
- Receiver Based Repair
- Recommendations

IP Multicast Channel Characteristics

- Group address
  - Client receives on address
  - Sender sends to address, without knowledge of clients
- Loosely coupled connections
  - "Extension" to UDP
  - Not two-way
  - Makes it scalable
  - Allows clients to do *local* repair
- Multicast router shares with unicast traffic
  - Can have high loss
  - Often MBone router 2nd rate

MBone Loss Characteristics

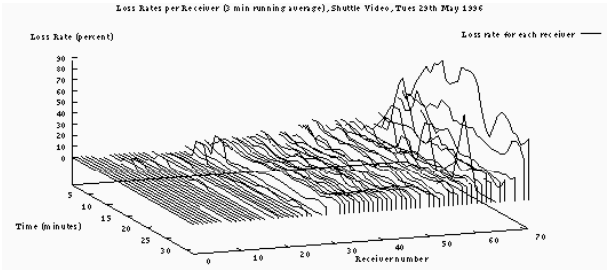

- Some receivers near 0% loss rates
- Most receivers in the 2-5% loss range
- Some see 20-50% loss
- Characteristics differ, so need local decisions

Mbone Jitter Characteristics

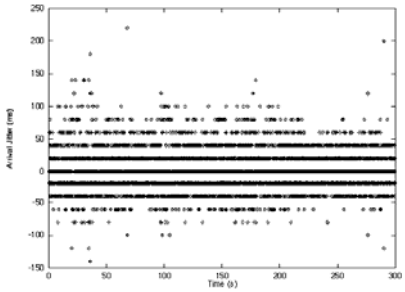

- High jitter
  - If too late, will be discarded and look like loss (e.g. I-policy)
- Interactive applications need low latency
  - Influences repair scheme

Media Repair Taxonomy

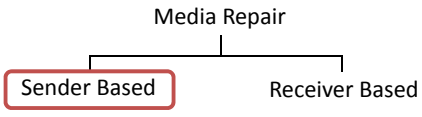

Sender Based Repair Taxonomy

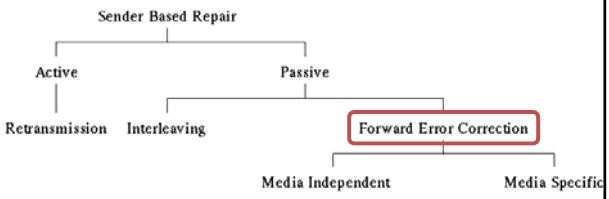

- Work from right to left
- Unit of audio data vs. a packet
  - Unit may be composed of several packets
  - Or one packet may have several units of audio data

Forward Error Correction (FEC)

- Add extra data to stream
- Use extra data to recover lost packets
- Two classes:
  - Media independent (not multimedia specific)
  - Media dependent (knowledge of audio or video)

Media Independent FEC

- Given *k* data packets
- Generate *n-k* check packets
- Transmit *n* packets
- Schemes originally for bits (like *checksums* in packet headers)
  - Applied to packets
  - So, for example *i*'th bit of check packet, checks *i*'th bit of each associated packet

FEC Coding

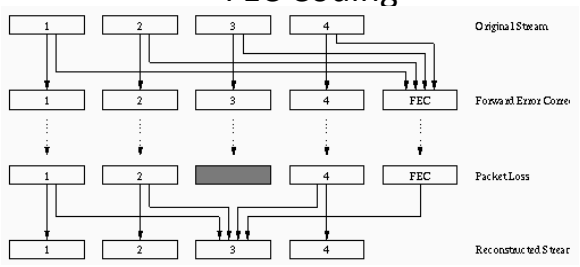

e.g., XOR operation across all packets
Transmit 1 parity packet every *n* data packets
If 1 loss in *n* packets, can fully recover

e.g., Reed-Solomon treat as polynomial, add k packets redundancy
→ If *k-1* loss in *n* packets, can fully recover

Media Independent FEC Advantages and Disadvantages

- Advantages
  - Media independent
    - Audio, video, different compression schemes
  - Computation is small and easy to implement
- Disadvantages
  - Add delay (must wait for all *n* packets)
  - Add to bitrate (causing more loss?)
  - Add decoder complexity

Sender Based Repair Taxonomy

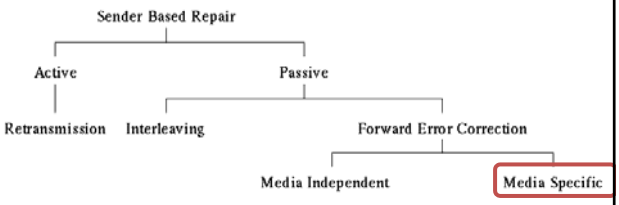

Figure 3: A Taxonomy of Sender Based Repair Techniques

Media Specific FEC

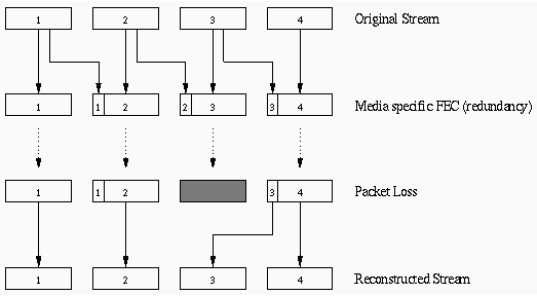

- Multiple copies of data
- "Quality" of secondary frames?

Media Specific FEC Secondary Frame

- Send packet energy and zero crossing rate
  - 2 numbers, so small
  - Coarse, but effective for small loss
    - Better than interpolating across missing packets
- Low bit-rate encoded version of primary
  - Lower number of sample bits audio sample, say
- Full-version of secondary
  - Effective if primary is small (low bandwidth)

Media Specific FEC Discussion

- Typical overhead 20-30% for low-quality
- Media specific FEC can repair various amounts by trading off quality of repair
  - Contrast with media independent FEC has fixed number of bits for certain amount of full repair
- Can have adaptive FEC
  - When speech changes and cannot interpolate
  - Add when increase in loss
  - Delay more than 1 packet when bursty loss

Media Specific FEC Advantages and Disadvantages

- Advantages
  - Low latency
    - Only wait for one additional packet to repair
    - Or multiple if adapted to bursty losses
  - Can have less bandwidth than independent FEC
- Disadvantages
  - Computation may be more difficult to implement
  - Still adds to bitrate
  - Adds decoder complexity
  - Typically lower quality (vs. other methods of repair)

Sender Based Repair Taxonomy

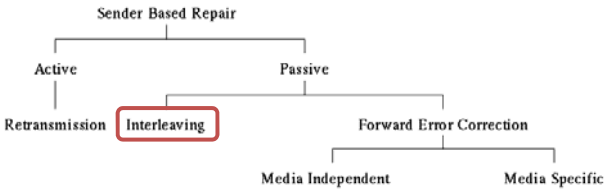

Figure 3: A Taxonomy of Sender Based Repair Techniques

Interleaving

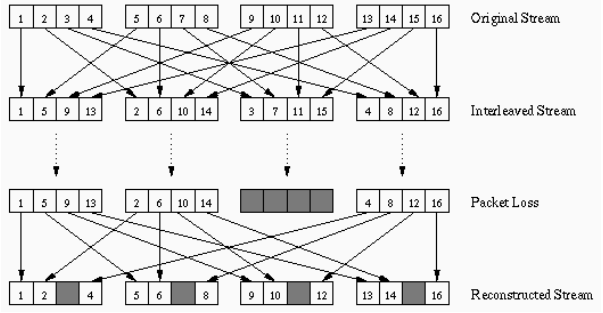

- Doesn't really repair, rather mitigates effects of loss
- Many audio tools send 1 phoneme (40 ms of sound), so most of phoneme intact

Interleaving Advantages and Disadvantages

- Advantages
  - Most audio compression schemes can do interleaving without additional complexity
  - No extra bitrate added
- Disadvantages
  - Delay of interleaving factor in packets
    - Even when not repairing!
  - Gains to quality can be modest

Sender Based Repair Taxonomy

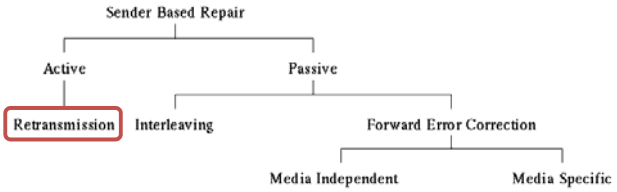

Figure 3: A Taxonomy of Sender Based Repair Techniques

Retransmission

- If delays less than 250 ms, can do retransmission
  - Effective for LAN or fast Internet connection
  - But wide-area wireless & inter-continetnal connection can be 200ms +
- Scalable Reliable Multicast (SRM)
  - Hosts time-out based on distance from sender
    - To avoid implosion
  - Mcast repair request (and repair) to all
  - All hosts can reply (timers based on distance stop implosion)

Retransmission Discussion

- In typical multicast session, can have every packet usually lost by *some* receiver
  - Will *always* retransmit at least once
  - FEC may save bandwidth
- Typically, crossover point to FEC based on loss rate
- Some participants may not be interactive
  - Use retransmission
  - Others use FEC

Retransmission Advantages and Disadvantages

- Advantages
  - Well understood
  - Only add additional data 'as needed'
- Disadvantages
  - Potentially large delay
    - Not usually suitable for interactive applications
  - Large jitter (different for different receivers)
  - Implosion (setting timers difficult)

Media Repair Taxonomy

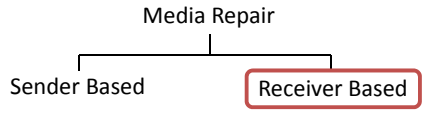

- Do not require assistance of Sender
  - Receiver recovers as best it can
- Often called *Error Concealment*
- Can work well for small loss (up to 15%), small packets (4-40 ms)
- Not substitute for sender-based
  - Rather use both
  - Receiver based can conceal what is left

Taxonomy of Error Concealment

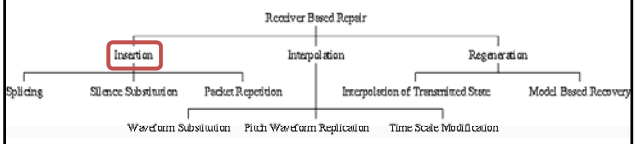

- When packet is lost, replace with fill-in

Splicing

- Splice together stream on either side
  - Do not preserve timing
- Advantages
  - Easy
  - Works ok for short packets of 4-16 ms
- Disadvantages
  - Poor quality for losses above 3%
  - Can interfere with delay buffering

Silence Substitution

- Fill gap left by lost packet with silence
  - Preserve timing
- Advantages
  - Still Easy
  - Works well for low loss (< 2%)
  - Works ok for short packets of 4-16 ms
- Disadvantages
  - Poor quality for higher losses (3%+)
  - Ineffective with 40 ms packets (typical)

Noise Substitution

- Human psych says can repair if sound, not silence (*phonemic restoration*)
  - Replace lost packet with "white noise"
    - Like static on radio
  - Still preserve timing
- Similar to silence substitution
- Sender can send "comfort noise" so receiver gets white-noise volume right

Repetition

- Replace missing packet with previous packet
- Can "fade" if multiple repeats over time
  - Decrease signal amplitude to 0
- Still pretty easy, but can work better than nothing
- A step towards interpolation techniques (next)

Noise Substitution and Repetition

- Advantages
  - Easy to implement
  - Works well for small loss (up to 5%)
- Disadvantages
  - Still doesn't work well for larger losses
  - Does not work well for larger packets

Taxonomy of Error Concealment

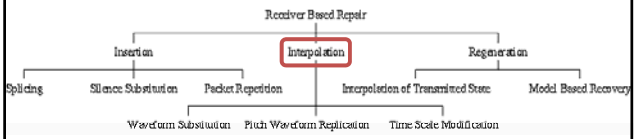

- When packet is lost, reproduce packet based on surrounding packets.

Interpolation Based Repair

- Waveform substitution
  - Use waveform repetition from both sides of loss
  - Works better than repetition (that uses one side)
- Pitch waveform replication
  - Use repetition during unvoiced speech and use additional pitch length during voiced speech
  - Performs marginally better than waveform
- Time scale modifications
  - "Stretch" audio signal across gap
  - Generate new waveform that smoothly blends across loss
  - Computationally heavier, but performs marginally better than others

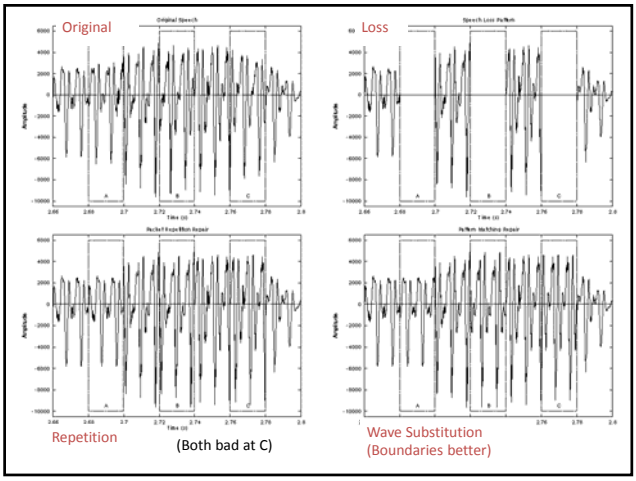

Taxonomy of Error Concealment

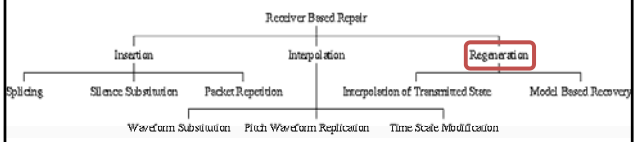

- Use knowledge of audio compression to derive codec parameters, using knowledge of code to regenerate

Regeneration Based Repair

- Interpolation of transmitted state
  - State-based decoding can then interpret what state codec should be in
  - Reduces boundary-effects
  - Typically high processing
- Model-Based recovery
  - Regenerate 'speech' to fit with speech on either side
  - Very complicated, often language dependent

Summary of Receiver Based Repair

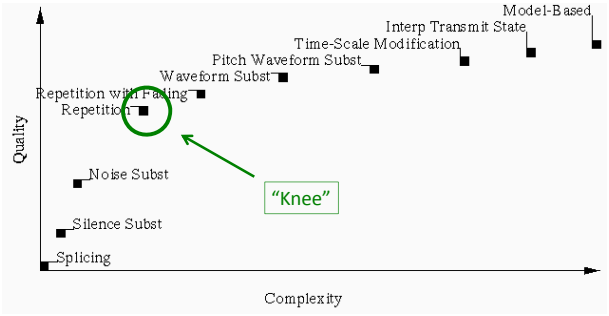

- Quality increase decreases at high complexity
- Repetition is at 'knee' in curve

Groupwork

- Consider:
  - Interactive voice from Asia to U.S.
  - Multicast video of taped lecture
  - Multicast replicated database update
  - Interactive voice across city
- Choose a repair technique and justify:
  - Interleaving
  - Retransmission
  - Media Specific FEC
  - Media Independent FEC

Recommendations: Non-Interactive Applications

- Latency less important
- Bitrate a concern (mcast has varied capacities)
  - → Can use *interleaving*
  - → Use *repetition* for concealment
- Retransmission does not scale
  - Ok for unicast
- Media independent FEC may be ok

Recommendations: Interactive Applications

- Want to minimize delay
  - → *Interleaving* delay is too large
  - → *Retransmission* delay can be large
  - → *Media independent FEC* usually large
    - (Or computationally expensive)
- Can use *media specific FEC*
  - Delay is low
  - Approximate repair is ok
  - Can be tuned (via quality and repair placement) to suit network and user