

Colin Perkins
Orion Hodson
University College London

Options for Repair of Streaming Media
draft-ietf-avt-info-repair-01.txt

Status of this memo

This document is an Internet-Draft. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as ``work in progress''. To learn the current status of any Internet-Draft, please check the ``lid-abstracts.txt'' listing contained in the Internet-Drafts Shadow Directories on ftp.is.co.za (Africa), ftp.nordu.net (Europe), munnari.oz.au (Pacific Rim), ds.internic.net (US East Coast), or ftp.isi.edu (US West Coast).

Distribution of this document is unlimited.

Comments are solicited and should be addressed to the author(s) and/or the IETF Audio/Video Transport working group's mailing list at rem-conf@es.net.

Abstract

This document summarizes a range of possible techniques for the repair of continuous media streams subject to packet loss. The techniques discussed include redundant transmission, retransmission, interleaving and forward error correction. The range of applicability of these techniques is noted, together with the protocol requirements and dependencies.

1 Introduction

A number of applications have emerged which use RTP/UDP transport to deliver continuous media streams. Due to the unreliable nature of UDP packet delivery, the quality of the received stream will be adversely affected by packet loss. A number of techniques exist by which the effects of packet loss may be repaired. These techniques

have a wide range of applicability and require varying degrees of protocol support. In this document, a number of such techniques are discussed, and recommendations for their applicability made.

2 Terminology and Protocol Framework

A unit is defined to be a timed interval of media data, typically derived from the workings of the media coder. A packet comprises one or more units, encapsulated for transmission over the network. For example, many audio coders operate on 20ms units, which are typically combined to produce 40ms or 80ms packets for transmission.

The framework of RTP [15] is assumed. This implies that packets have a sequence number and timestamp. The sequence number denotes the order in which packets are transmitted, and is used to detect losses. The timestamp is used to determine the playout order of units. Most loss recovery schemes rely on units being sent out of order, so an application must use the RTP timestamp to schedule playout.

The use of RTP allows for several different media coders, with a payload type field being used to distinguish between these at the receiver. Some loss recovery schemes send some units multiple times, using different encoding schemes. A receiver is assumed to have a 'quality' ranking of the differing encodings, and so is capable of choosing the 'best' unit for playout, given multiple options.

3 Network Loss Characteristics

If it is desired to repair a media stream subject to packet loss, it is useful to have some knowledge of the loss characteristics which are likely to be encountered. A number of studies have been conducted on the loss characteristics of the Mbone [8,9] and although the results vary somewhat, the broad conclusion is clear: in a large conference it is inevitable that some receivers will experience packet loss. Packet traces taken by Handley [5] show a session in which most receivers experience loss in the range 2-5%, with a somewhat smaller number seeing significantly higher loss rates. Other studies have presented broadly similar results.

It has also been shown that the vast majority of losses are of single packets. Burst losses of two or more packets are around an order of magnitude less frequent than single packet loss, although they do occur more often than would be expected from a purely random process. Longer burst losses (of the order of tens of packets) occur infrequently. These results are consistent with a network where small amounts of transient congestion cause the majority of packet loss. In a few

cases, a network link is found to be severely overloaded, and large amount of loss results.

The primary focus of a packet loss repair scheme must, therefore, be to correct single packet loss, since this is by far the most frequent occurrence. It is desirable that losses of a relatively small number of consecutive packets may also be repaired, since such losses represent a small but noticeable fraction of observed losses. The correction of large bursts of loss is of considerably less importance.

4 Loss Mitigation Schemes

In the following sections, four loss mitigation schemes are discussed. These schemes have been discussed in the literature a number of times, and found to be of use in a number of scenarios. Each technique is briefly described, and its advantages and disadvantages noted.

4.1 Forward Error Correction

Forward error correction (FEC) is the means by which repair data is added to a media stream, such that packet loss can be repaired by the receiver of that stream with no further reference to the sender. There are two classes of repair data which may be added to a stream: those which are independent of the contents of the stream, and those which use knowledge of the stream to improve the repair process.

4.1.1 Media-Independent FEC

A number of media-independent FEC schemes have been proposed for use with streamed media. These techniques add redundant data to a media stream which is transmitted in separate packets. Traditionally, FEC techniques are described as loss detecting and/or loss correcting. In the case of streamed media loss detection is provided by the sequence numbers in RTP packets.

The redundant FEC data is typically calculated using the mathematics of finite fields [1]. The simplest of finite field is $GF(2)$ where addition is just the eXclusive-OR operation.

Basic FEC schemes transmit k data packets with $n-k$ parity packets allowing the reconstruction of the original data from any k of the n transmitted packets. Budge et al [3]) proposed applying the XOR operation across different combinations of the media data with the redundant data transmitted separately as parity packets. These vary the pattern of packets over which the parity is calculated, and hence have different bandwidth, latency and loss repair characteristics.

Luby et al [8] have discussed applying parity in layers. The first layer in their scheme is the parity bits generated from the media. The second layer is calculated as the parity bits of the first layer and so on. This obviously improves the repair properties, but consumes additional bandwidth.

Parity-based FEC based techniques have a significant advantage in that they are media independent, and provide exact repair for lost packets. In addition, the processing requirements are relatively light, especially when compared with some redundancy schemes which use very low bandwidth, but high complexity encodings. The disadvantage of parity-based FEC is that the codings have higher latency in comparison with the media-specific schemes discussed in following section. An RTP payload format for parity-based FEC is defined in [14]. The format is generic, and can specify many different parity encodings.

A number of FEC schemes exist which are based on higher-order finite fields. An example of such are Reed-Solomon (RS) codes which are more sophisticated and computationally demanding. These are usually structured so that they have good burst loss protection. There has been much work conducted in this area, and it is believed that a number of streaming applications use RS codes.

4.1.2 Media-Specific FEC

The basis of media-specific FEC is to employ knowledge of a media compression scheme to achieve more efficient repair of a stream than can otherwise be achieved. To repair a stream subject to packet loss, it is necessary to add redundancy to that stream: some information is added which is not required in the absence of packet loss, but which can be used to recover from that loss.

The nature of a media stream affects the means by which the redundancy is added. If units of media data are packets, or if multiple units are included in a packet, it is logical to use the unit as the level of redundancy, and to send duplicate units. By recoding the redundant copy of a unit, significant bandwidth savings may be made, at the expense of additional computational complexity and approximate repair. This approach has been advocated for use with streaming audio [5,6] and has been shown to perform well. An RTP payload format for this form of redundancy has been defined [12].

If media units span multiple packets, for instance video, it is sensible to include redundancy directly within the output of a codec. For example the proposed RTP payload for H.263+ [2] includes multiple copies of key portions of the stream, separated to avoid the problems of packet loss. The advantages of this second approach is efficiency: the codec designer knows exactly which portions of the stream are

most important to protect, and low complexity since each unit is coded once only.

An alternative approach is to apply media-independent FEC techniques to the most significant bits of a codecs output, rather than applying it over the entire packet. Several codec descriptions include bit sensitivities that make this feasible. This approach has low computational cost and can be tailored to represent an arbitrary fraction of the transmitted data.

The use of media-specific FEC has the advantage of low-latency, with only a single-packet delay being added. This makes it suitable for interactive applications, where large end-to-end delays cannot be tolerated. In a broadcast-style environment, it is possible to delay sending the redundant data, achieving improved performance in the presence of burst losses [7], at the expense of additional latency.

4.2 Retransmission

Retransmission of lost packets is an obvious means by which loss may be repaired. It is clearly of value in broadcast style applications, with relaxed delay bounds, but the delay imposed means that it does not typically perform well for interactive use.

In addition to the possibly high latency, there is a potentially large bandwidth overhead to the use of retransmission. Not only are units of data sent multiple times, but additional control traffic must flow to request the retransmission. It has been shown that, in a large Mbone session, most packets are lost by at least one receiver [5]. In this case the overhead of requesting retransmission for most packets may be such that redundant transmission is more acceptable. This leads to a natural synergy between the two mechanisms, with a redundant transmission being used to repair all single packet losses, and those receivers experiencing burst losses, and willing to accept the additional latency, using retransmission based repair as an additional recovery mechanism. Similar mechanisms have been used in a number of reliable multicast schemes, and have received some discussion in the literature [10, 6].

In order to reduce the overhead of retransmission, the retransmitted units may be piggy-backed onto the ongoing transmission. This also allows for the retransmission to be recoded in a different format, to further reduce the bandwidth overhead.

The choice of a retransmission request algorithm which is both timely and network friendly is an area of current study. An obvious starting point is the SRM protocol [4], and experiments have been conducted using this, and with a low-delay variant, STORM [17]. This work shows the trade-off between latency and quality for retransmission

based repair schemes, and illustrates that retransmission is an effective approach to repair for applications which can tolerate the latency.

An RTP profile extension for SRM-style retransmission requests is described in [11].

4.3 Interleaving

When the unit size is smaller than the packet size, and end-to-end delay is unimportant, interleaving [13] is a useful technique for reducing the effects of loss. Units are resequenced before transmission, so that originally adjacent units are separated by a guaranteed distance in the transmitted stream, and returned to their original order at the receiver. Interleaving disperses the effect of packet losses. If, for example, units are 5ms in length and packets 20ms (ie: 4 units per packet), then the first packet could contain units 1, 5, 9, 13; the second packet would contain units 2, 6, 10, 14; and so on. It can be seen that the loss of a single packet from an interleaved stream results in multiple small gaps in the reconstructed stream, as opposed to the single large gap which would occur in a non-interleaved stream. In many cases it is easier to reconstruct a stream with such loss patterns, although this is clearly media and codec dependent.

The obvious disadvantage of interleaving is that it increases latency. This limits the use of this technique for interactive applications, although it performs well for broadcast use. The major advantage of interleaving is that it does not increase the bandwidth requirements of a stream.

A potential RTP payload format for interleaved data is a simple extension of the redundant audio payload [12]. That payload requires that the redundant copy of a unit is sent after the primary. If this restriction is removed, it is possible to transmit an arbitrary interleaving of units with this payload format.

5 Recommendations

If the desired scenario is a one-to-many transmission, in the style of a radio or television broadcast, latency is of considerably less importance than reception quality. In this case, the use of interleaving and/or retransmission based repair is appropriate, with interleaving being preferred due to its bandwidth efficiency (provided that approximate repair is acceptable).

In an interactive session (typically defined as a session where the end-to-end delay is less than 250ms, this includes media coding/decoding, network transit and host buffering), the delay imposed by the use of interleaving and retransmission is not acceptable, and a low-latency

FEC scheme is the only means of repair suitable. The choice between media independent and media specific forward error correction is less clear-cut: media-specific FEC can be made more efficient, but requires modification to the output of the codec. When defining the packetisation for a new codec, this is clearly an appropriate technique, and should be encouraged.

If an existing codec is to be used, a media independent redundant transmission scheme is usually easier to implement, and can perform well. If the processing requirements are not excessive, recoding the redundant data using a different codec is an effective means of reducing the bandwidth overhead of a stream. A media stream protected in this way may be augmented with retransmission based repair with minimal overhead, providing improved quality for those receivers willing to tolerate additional delay.

Whilst the addition of error correction data to an media stream is an effective means by which that stream may be protected against packet loss, application designers should be aware that the addition of large amounts of repair data will increase network congestion, and hence packet loss, leading to a worsening of the problem which the use of error correction coding was intended to solve.

At the time of writing, there is no standard solution to the problem of congestion control for streamed media which can be used to solve this problem. There have, however, been a number of contributions which show the likely form the solution will take [9, 16]. This work typically used some form of layered encoding of data over multiple channels, with receivers joining and leaving layers in response to packet-loss (which indicates congestion). The aim of such schemes is to emulate the congestion control behaviour of a TCP stream, and hence compete fairly with non-real-time traffic. This is necessary for stable network behaviour in the presence of much streamed media.

6 Author's Address

Colin Perkins/Orion Hodson
Department of Computer Science
University College London
Gower Street
London WC1E 6BT
United Kingdom

Email: <c.perkins|o.hodson>@cs.ucl.ac.uk

References

- [1] R.E. Blahut. Theory and Practice of Error Control Codes. Addison Wesley, 1983.

- [2] C. Bormann, L. Cline, G. Deisher, T. Gardos, C. Maciocco, D. Newell, J. Ott, S. Wenger, and C. Zhu. RTP payload format for the 1998 version of ITU-T rec. H.263 video (H.263+). IETF Audio/Video Transport Working Group, November 1997. Work in progress.
- [3] D. Budge, R. McKenzie, W. Mills, W. Diss, and P. Long. Media-independent error correction using RTP, May 1997. Work in progress.
- [4] S. Floyd, V. Jacobson, S. McCanne, C.-G. Liu, and L. Zhang. A reliable multicast framework for light-weight sessions and applications level framing. IEEE/ACM Transactions on Networking, 1995.
- [5] M. Handley. An examination of Mbone performance. USC/ISI Research Report: ISI/RR-97-450, April 1997.
- [6] M. Handley and J. Crowcroft. Network text editor (NTE): A scalable shared text editor for the Mbone. In Proceedings ACM SIGCOMM'97, Cannes, France, September 1997.
- [7] I. Kouvelas, O. Hodson, V. Hardman, and J. Crowcroft. Redundancy control in real-time Internet audio conferencing. In Proceedings of AVSPN'97, Aberdeen, Scotland, September 1997.
- [8] G.M. Luby, M. Mitzenmacher, M. Amin Shokrollahi, D.A. Spielman, and V. Stemann. Practical loss-resilient codes. In Twenty-Ninth Annual ACM Symposium on the Theory of Computing, 1997.
- [9] S. McCanne, V. Jacobson, and M. Vetterli. Receiver-driven layered multicast. In Proceedings ACM SIGCOMM'96, Stanford, CA., August 1996.
- [10] J. Nonnenmacher, E. Biersack, and D. Towsley. Parity-based loss recovery for reliable multicast transmission. In Proceedings ACM SIGCOMM'97, Cannes, France, September 1997.
- [11] P. Parnes. RTP extension for scalable reliable multicast, November 1996. Work in progress.
- [12] C. S. Perkins, I. Kouvelas, O. Hodson, V. Hardman, M. Handley, J.-C. Bolot, A. Vega-Garcia, and S. Fosse-Parisis. RTP Payload for redundant audio data. IETF Audio/Video Transport Working Group, 1997. RFC2198.
- [13] J.L. Ramsey. Realization of optimum interleavers. IEEE Transactions on Information Theory, IT-16:338--345, May 1970.

- [14] J. Rosenberg and H. Schulzrinne. An A/V profile extension for generic forward error correction in RTP. IETF Audio/Video Transport working group, July 1997. Work in progress.
- [15] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson. RTP: A transport protocol for real-time applications. IETF Audio/Video Transport Working Group, January 1996. RFC1889.
- [16] L. Vicisano, L. Rizzo, and J. Crowcroft. TCP-like congestion control for layered multicast data transfer. In Proceedings IEEE INFOCOM'98, 1998.
- [17] R. X. Xu, A. C. Myers, H. Zhang, and R. Yavatkar. Resilient multicast support for continuous media applications. In Proceedings of the 7th International Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV'97), Washington University in St. Louis, Missouri, May 1997.