Internet Engineering Task Force AVT WG
INTERNET-DRAFT M. Handley, C. Perkins
draft-ietf-avt-rtp-format-guidelines-04.txt ACIRI, UCL
15th Oct 1999
Expires: Apr 1999

Guidelines for Writers of RTP Payload Format Specifications

Abstract

This document provides general guidelines aimed at assisting the authors of RTP Payload Format specifications in deciding on good formats. These guidelines attempt to capture some of the experience gained with RTP as it evolved during its development.

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of Section 10 of RFC2026. 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.''

The list of current Internet-Drafts can be accessed at
http://www.ietf.org/ietf/1id-abstracts.txt

The list of Internet-Draft Shadow Directories can be accessed at
http://www.ietf.org/shadow.html.

1. Introduction

This document provides general guidelines aimed at assisting the authors of RTP [9] Payload Format specifications in deciding on good formats. These guidelines attempt to capture some of the experience gained with RTP as it evolved during its development.

The principles outlined in this document are applicable to almost all data types, but are framed in examples of audio and video codecs for clarity.

2. Background

RTP was designed around the concept of Application Level Framing (ALF), first described by Clark and Tennenhouse[2]. The key argument underlying ALF is that there are many different ways an application might be able to cope with misordered or lost packets. These range from ignoring the loss, to re-sending the missing data (either from a buffer or by regenerating it), and to sending new data which supersedes the missing data. The application only has this choice if the transport protocol is dealing with data in ''Application Data Units'' (ADUs). An ADU contains data that can be processed out-of-order with respect to other ADUs. Thus the ADU is the minimum unit of error recovery.

The key property of a transport protocol for ADUs is that each ADU contains sufficient information to be processed by the receiver immediately. An example is a video stream, wherein the compressed video data

in an ADU must be capable of being decompressed regardless of whether previous ADUs have been received. Additionally the ADU must contain ``header'' information detailing its position in the video image and the frame from which it came.

Although an ADU need not be a packet, there are many applications for which a packet is a natural ADU. Such ALF applications have the great advantage that all packets that are received can be processed by the application immediately.

RTP was designed around an ALF philosophy. In the context of a stream of RTP data, an RTP packet header provides sufficient information to be able to identify and decode the packet irrespective of whether it was received in order, or whether preceding packets have been lost. However, these arguments only hold good if the RTP payload formats are also designed using an ALF philosophy.

Note that this also implies smart, network aware, end-points. An application using RTP should be aware of the limitations of the underlying network, and should adapt its transmission to match those limitations. Our experience is that a smart end-point implementation can achieve significantly better performance on real IP-based networks than a naive implementation.

3. Channel Characteristics

We identify the following channel characteristics that influence the best-effort transport of RTP over UDP/IP in the Internet:

- Packets may be lost

- Packets may be duplicated
- Packets may be reordered in transit
- Packets will be fragmented if they exceed the MTU of the underlying network

The loss characteristics of a link may vary widely over short time intervals.

Although fragmentation is not a disastrous phenomenon if it is a rare occurrence, relying on IP fragmentation is a bad design strategy as it significantly increases the effective loss rate of a network and decreases goodput. This is because if one fragment is lost, the remaining fragments (which have used up bottleneck bandwidth) will then need to be discarded by the receiver. It also puts additional load on the routers performing fragmentation and on the end-systems re-assembling the fragments.

In addition, it is noted that the transit time between two hosts on the Internet will not be constant. This is due to two effects - jitter caused by being queued behind cross-traffic, and routing changes. The former is possible to characterise and compensate for by using a playout buffer, but the latter is impossible to predict and difficult to accommodate gracefully.

4. Guidelines

We identify the following requirements of RTP payload format specifications:

+ A payload format should be devised so that the stream being transported is still useful even in the presence of a moderate amount of packet loss.

+ Ideally all the contents of every packet should be possible to be decoded and played out irrespective of whether preceding packets have been lost or arrive late.

The first of these requirements is based on the nature of the Internet. Although it may be possible to engineer parts of the Internet to produce low loss rates through careful provisioning or the use of non-best-effort services, as a rule payload formats should not be designed for these special purpose environments. Payload formats should be designed to be used in the public Internet with best effort service, and thus should expect to see moderate loss rates. For example, a 5% loss rate is not uncommon. We note that TCP steady state models[3][4][6] indicate

that a 5% loss rate with a 1KByte packet size and 200ms round-trip time will result in TCP achieving a throughput of around 180Kbit/s. Higher loss rates, smaller packet sizes, or a larger RTT are required to constrain TCP to lower data rates. For the most part, it is such TCP traffic that is producing the background loss that many RTP flows must coexist with. Without explicit congestion notification (ECN)[8], loss must be considered an intrinsic property of best-effort parts of the Internet.

When payload formats do not assume packet loss will occur, they should state this explicitly up front, and they will be considered special purpose payload formats, unsuitable for use on the public Internet without special support from the network infrastructure.

The second of these requirements is more explicit about how RTP should cope with loss. If an RTP payload format is properly designed, every packet that is actually received should be useful. Typically this implies the following guidelines are adhered to:

+ Packet boundaries should coincide with codec frame boundaries. Thus a packet should normally consist of one or more complete codec frames.

+ A codec's minimum unit of data should never be packetised so that it crossed a packet boundary unless it is larger than the MTU.

+ If a codec's frame size is larger than the MTU, the payload format must not rely on IP fragmentation. Instead it must define its own fragmentation mechanism. Such mechanisms may involve codec-specific information that allows decoding of fragments. Alternatively they might allow codec-independent packet-level forward error correction[5] to be applied that cannot be used with IP-level fragmentation.

In the abstract, a codec frame (i.e., the ADU or the minimum size unit that has semantic meaning when handed to the codec) can be of arbitrary size. For PCM audio, it is one byte. For GSM audio, a frame corresponds to 20ms of audio. For H.261 video, it is a Group of Blocks (GOB), or one twelfth of a CIF video frame.

For PCM, it does not matter how audio is packetised, as the ADU size is one byte. For GSM audio, arbitrary packetisation would split a 20ms frame over two packets, which would mean that if one packet were lost,

partial frames in packets before and after the loss are meaningless. This means that not only were the bits in the missing packet lost, but also that additional bits in neighboring packets that used bottleneck bandwidth were effectively also lost because the receiver must throw them away. Instead, we would packetise GSM by including several

complete GSM frames in a packet; typically four GSM frames are included in current implementations. Thus every packet received can be decoded because even in the presence of loss, no incomplete frames are received.

The H.261 specification allows GOBs to be up to 3KBytes long, although most of the time they are smaller than this. It might be thought that we should insert a group of blocks into a packet when it fits, and arbitrarily split the GOB over two or more packets when a GOB is large. In the first version of the H.261 payload format, this is what was done. However, this still means that there are circumstances where H.261 packets arrive at the receiver and must be discarded because other packets were lost - a loss multiplier effect that we wish to avoid. In fact there are smaller units than GOBs in the H.261 bit-stream called macroblocks, but they are not identifiable without parsing from the start of the GOB. However, if we provide a little additional information at the start of each packet, we can reinstate information that would normally be found by parsing from the start of the GOB, and we can packetise H.261 by splitting the data stream on macroblock boundaries. This is a less obvious packetisation for H.261 than the GOB packetisation, but it does mean that a slightly smarter depacketiser at the receiver can reconstruct a valid H.261 bitstream from a stream of RTP packets that has experienced loss, and not have to discard any of the data that arrived.

An additional guideline concerns codecs that require the decoder state machine to keep step with the encoder state machine. Many audio codecs such as LPC or GSM are of this form. Typically they are loss tolerant, in that after a loss, the predictor coefficients decay, so that after a certain amount of time, the predictor error induced by the loss will disappear. Most codecs designed for telephony services are of this form because they were designed to cope with bit errors without the decoder predictor state permanently remaining incorrect. Just packetising these formats so that packets consist of integer multiples of codec frames may not be optimal, as although the packet received immediately after a packet loss can be decoded, the start of the audio stream produced will be incorrect (and hence distort the signal) because the decoder predictor is now out of step with the encoder. In principle, all of the decoder's internal state could be added using a header attached to the start of every packet, but for lower bit-rate encodings, this state is so substantial that the bit rate is no longer low. However, a compromise can usually be found, where a greatly reduced form of decoder state is sent in every packet, which does not recreate the encoders predictor precisely, but does reduce the magnitude and duration of the distortion produced when the previous packet is lost. Such compressed state is, by definition, very dependent on the codec in question. Thus we recommend:

+ Payload formats for encodings where the decoder contains internal data-driven state that attempts to track encoder state should

normally consider including a small additional header that conveys the most critical elements of this state to reduce distortion after

packet loss.

A similar issue arises with codec parameters, and whether or not they should be included in the payload format. An example is with a codec that has a choice of huffman tables for compression. The codec may use either huffman table 1 or table 2 for encoding and the receiver needs to know this information for correct decoding. There are a number of ways in which this kind of information can be conveyed:

- Out of band signalling, prior to media transmission.
- Out of band signalling, but the parameter can be changed mid-session. This requires synchronization of the change in the media stream.
- The change is signaled through a change in the RTP payload type field. This requires mapping the parameter space into particular payload type values and signalling this mapping out-of-band prior to media transmission.
- Including the parameter in the payload format. This allows for adapting the parameter in a robust manner, but makes the payload format less efficient.

Which mechanism to use depends on the utility of changing the parameter in mid-session to support application layer adaptation. However, using out-of-band signalling to change a parameter in mid-session is generally to be discouraged due to this problems of synchronizing the parameter change with the media stream.

4.1. RTP Header Extensions

Many RTP payload formats require some additional header information to be carried in addition to that included in the fixed RTP packet header. The recommended way of conveying this information is in the payload section of the packet. The RTP header extension should not be used to convey payload specific information ([9],section 5.3) since this is inefficient in its use of bandwidth; requires the definition of a new RTP profile or profile extension; and makes it difficult to employ FEC schemes such as, for example, [7]. Use of an RTP header extension is only appropriate for cases where the extension in question applies across a wide range of payload types.

4.2. Header Compression

Designers of payload formats should also be aware of the needs of RTP header compression [1]. In particular, the compression algorithm functions best when the RTP timestamp increments by a constant value between consecutive packets. Payload formats which rely on sending packets out of order, such that the timestamp increment is not constant, are likely to compress less well than those which send packets in order. This has most often been an issue when designing payload formats for FEC information, although some video codecs also rely on out-of-order transmission of packets at the expense of reduced compression. Although in some cases such out-of-order transmission may be the best solution, payload format designers are encourage to look for alternative solutions where possible.

5. Summary

Designing packet formats for RTP is not a trivial task. Typically a

detailed knowledge of the codec involved is required to be able to design a format that is resilient to loss, does not introduce loss magnification effects due to inappropriate packetisation, and does not introduce unnecessary distortion after a packet loss. We believe that considerable effort should be put into designing packet formats that are well tailored to the codec in question. Typically this requires a very small amount of processing at the sender and receiver, but the result can be greatly improved quality when operating in typical Internet environments.

Designers of new codecs for use with RTP should consider making the output of the codec ''naturally packetizable''. This implies that the codec should be designed to produce a packet stream, rather than a bit-stream; and that that packet stream contains the minimal amount of redundancy necessary to ensure that each packet is independently decodable with minimal loss of decoder predictor tracking. It is recognised that sacrificing some small amount of bandwidth to ensure greater robustness to packet loss is often a worthwhile tradeoff.

It is hoped that, in the long run, new codecs should be produced which can be directly packetised, without the trouble of designing a codec-specific payload format.

It is possible to design generic packetisation formats that do not pay attention to the issues described in this document, but such formats are only suitable for special purpose networks where packet loss can be avoided by careful engineering at the network layer, and are not suited to current best-effort networks.

6. Security Considerations

The guidelines in this document result in RTP payload formats that are robust in the presence of real world network conditions. Designing payload formats for special purpose networks that assume negligable loss rates will normally result in slightly better compression, but produce formats that are more fragile, thus rendering them easier targets for denial-of-service attacks.

Designers of payload formats should pay close attention to possible security issues that might arise from poor implementations of their formats, and should be careful to specify the correct behaviour when anomalous conditions arise. Examples include how to process illegal field values, and conditions when there are mismatches between length fields and actual data. Whilst the correct action will normally be to discard the packet, possible such conditions should be brought to the attention of the implementor to ensure that they are trapped properly.

The RTP specification covers encryption of the payload. This issue should not normally be dealt with by payload formats themselves. However, certain payload formats spread information about a particular application data unit over a number of packets, or rely on packets which relate to a number of application data units. Care must be taken when changing the encryption of such streams, since such payload formats may constrain the places in a stream where it is possible to change the encryption key without exposing sensitive data.

Designers of payload formats which include FEC should be aware that the automatic addition of FEC in response to packet loss may increase network congestion, leading to a worsening of the problem which the use of FEC was intended to solve. Since this may, at its worst, constitute a denial of service attack, designers of such payload formats should take

care that appropriate safeguards are in place to prevent abuse.

Authors Addresses

Mark Handley
AT&T Center for Internet Research at ICSI,
International Computer Science Institute,
1947 Center Street, Suite 600,
Berkeley, CA 94704, USA
mjh@aciri.org

Colin Perkins
Dept of Computer Science,
University College London,
Gower Street,
London WC1E 6BT, UK.

C.Perkins@cs.ucl.ac.uk

Acknowledgments

This document is based on experience gained over several years by many people, including Van Jacobson, Steve McCanne, Steve Casner, Henning Schulzrinne, Thierry Turletti, Jonathan Rosenberg and Christian Huitema amongst others.

References

[1] S. Casner, V. Jacobson, ``Compressing IP/UDP/RTP Headers for Low-Speed Serial Links'', RFC 2508.

[2] D. Clark, D. Tennenhouse, "Architectural Considerations for a New Generation of Network Protocols" Proc ACM Sigcomm 90.

[3] J. Mahdavi and S. Floyd. ``TCP-friendly unicast rate-based flow control''. Note sent to end2end-interest mailing list, Jan 1997.

[4] M. Mathis, J. Semske, J. Mahdavi, and T. Ott. ``The macro-scopic behavior of the TCP congestion avoidance algorithm''. Computer Communication Review, 27(3), July 1997.

[5] J. Nonnenmacher, E. Biersack, Don Towsley, ``Parity-Based Loss Recovery for Reliable Multicast Transmission'', Proc ACM Sigcomm '97, Cannes, France, 1997.

[6] J. Padhye, V. Firoiu, D. Towsley, J. Kurose, ``Modeling TCP Throughput: A Simple Model and its Empirical Validation'', Proc. ACM Sigcomm 1998.

[7] C. Perkins, I. Kouvelas, O. Hodson, V. Hardman, M. Handley, J.C. Bolot, A. Vega-Garcia, S. Fosse-Parisis, ``RTP Payload for Redundant Audio Data'', RFC 2198.

[8] K. K. Ramakrishnan, Sally Floyd, ``A Proposal to add Explicit Congestion Notification (ECN) to IP'' RFC 2481, Jan 1999.

[9] H.Schulzrinne, S.Casner, R.Frederick, V. Jacobson, "Real-Time Transport Protocol", RFC1889, Jan 1996.