

Network Working Group
Request for Comments: 2448
Category: Informational

M. Civanlar
G. Cash
B. Haskell
AT&T Labs-Research
November 1998

AT&T's Error Resilient Video Transmission Technique

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (1998). All Rights Reserved.

Abstract

This document describes a set of techniques for packet loss resilient transmission of compressed video bitstreams based on reliable delivery of their vital information-carrying segments. The described techniques can be used over packet networks without packet prioritization. These techniques are related to AT&T/Lucent patents [1, 2].

1. Introduction

It is well known that every bit in a compressed video bitstream is not equal. Some bits belong to segments defining vital information such as picture types, quantization values, parameter ranges, average intensity values for image blocks, etc. When transporting compressed video bitstreams over packet networks, packet losses from such segments cause a much longer lasting and severe degradation on the output of a decoder than that caused by packet losses from other segments. We will call the vital information-carrying segments "High Priority (HP)" segments. The rest of the bitstream consists of "Low Priority (LP)" segments. Clearly, the video outputs resulting from transport techniques that protect the HP segments against packet losses are more resilient to packet losses in general.

Protection of the HP segments can be accomplished in many ways. These include:

- redundant transmission of the HP segments as described in [3] for MPEG RTP payloads

- using forward error correction (FEC) techniques
- transmitting HP segments over reserved channels or using differentiated services.

Both redundant transmission and FEC techniques increase the bandwidth needed to transmit the compressed video bitstream. FEC techniques increase the effectiveness of this additional bandwidth for packet loss protection at the expense of increased processing at the receiver and the transmitter ends and increased overall delay. Using channel reservations or differentiated services based approaches may be the best solutions for protecting the HP segments but, they require network infrastructure changes.

This document outlines another set of HP segment protection techniques based on AT&T/Lucent patents [1, 2] that can be used for reliable video transmission over packet networks without a built-in prioritization mechanism. These techniques use reliable transport protocols and "out-of-band" delivery approaches. In this context, the term "out-of-band" is used to imply information transmission means other than those used for transmitting the main video stream. The details of these techniques are discussed in the following sections. An implementation of these, as applied to MPEG-2 video transmission over IP networks, is described in [4].

The IESG/IETF take no position regarding the validity or scope of any intellectual property right or other rights that might be claimed to pertain to the implementation or use of the technology, or the extent to which any license under such rights might or might not be available. See the IETF IPR web page at <http://www.ietf.org/ipr.html> for any additional information that has been forwarded to the IETF.

2. Identification of the HP segments

The classification of a part of a video bitstream as an HP segment depends on two factors. The first one is the encoding algorithm used in compressing the video data. It is impossible to segment a compressed video bitstream without knowing the syntax and the semantics of the encoding algorithm. The second factor is the determination of a compromise between the HP segment size and the corresponding loss resilience. As the segment size increases, so does the loss resilience. On the other hand, it may not be feasible to deliver large HP segments reliably.

As an example, the "data partitioning" method of the MPEG-2 standard [5] defines the syntax and semantics for one particular way of partitioning an MPEG-2 encoded video bitstream into HP and LP segments. In data partitioning, the smallest useful HP segment can be selected to contain only the header information, which is usually

less than two percent of the video data. HP segments defined this way contain vital information including picture type, quantization factor, motion vector ranges, etc. without which the rest of the bitstream is not decodable. As an alternative, the DC coefficients (the average values) for each picture macroblock may be included in the HP segment increasing its size to about 40% of the bitstream. This way HP segments can be made to carry somewhat usable video information also; however, their reliable transmission may become a demanding task.

Since it is not possible to formulate a general technique that can be used for identifying the HP segments in any encoded video bitstream, we will assume that such segments are identified some way prior to the transmission. For example, some encoders can generate HP and LP segments separately, a stored bitstream can be in the partitioned format, etc. Also, consistent with most of the popular coding techniques, we assume that the HP segments (HP1, HP2, ...) are dispersed on the entire bitstream over time as shown in Fig. 1.




Figure 1

HP segments dispersed on an encoded video bitstream over time

3. Transmission of HP data using a reliable transport protocol [1]

In this approach, one or more of the HP segments are transmitted using a reliable transport protocol prior to starting the transmission of the LP segments. For point-to-point applications, TCP, for multipoint applications, an appropriate reliable multicast protocol [6] may be used for transporting the HP segments. The number of HP segments to be sent before starting the transmission of the LP segments depends on the application's tolerance to the start-up delay. Depending on the HP segment size and the path-MTU [7], one or more HP segments can be put in each packet carrying the HP data.

HP segments can be packetized using RTP with the following definitions for the header fields:

Payload Type: A distinct payload type number, which may be dynamic, should be assigned to HP segments of each video payload.

M Bit: Set for packets containing HP data for key pictures.

timestamp: Uses the same format as that of the video payload. Shows the sampling time for the video data following the first HP segment in the packet.

The SSRC field may be defined following the rules developed for the transmission of layered media streams in [8]. That is:

- A single SSRC space is used for the HP segment packets and the main video stream. Only the latter is used for SSRC allocation and conflict resolution. When a source discovers that it has collided, it transmits an RTCP BYE message on only the main video stream.
- A participant sends sender identification (SDS) on only the main video stream.

Most HP segments are self-identifying and can be packed without any additional headers. For others, techniques used for packetizing generic payload types may be used or special payload types may be defined.

It is possible to send the HP data along with the LP data (i.e., the original, unpartitioned bitstream) in addition to sending the HP segments separately. This way, the separately transmitted HP segments are needed only when packet losses occur.

4. Out-of-band transmission of the HP information [2]

In cases where a certain sequence of HP segments is used periodically for the entire duration of the video bitstream, this sequence may be transmitted once before the start of video transmission using a reliable transport protocol. The receiver can save this information and use it to recover lost HP segments during the main video transmission.

In this approach, the timestamps are not meaningful for the HP data and they may not be included in the transmitted HP segment sequence. In most cases, the synchronization between the stored HP segments and the LP data stream can be accomplished using the key-frames because the HP data sequence usually cover the video segment between two key-frames (e.g. a group-of-pictures (GOP) in MPEG). If the sequence of HP segments covers a video sequence with more than one key-frame, some indicator, e.g. if available the M-bit may be used to indicate a packet which carries the beginning of LP data that follows the first stored HP segment.

5. Security Considerations

RTP packets transmitted according to the techniques outlined in this document are subject to the security considerations discussed in the RTP specification [9]. This implies that confidentiality of the media streams is achieved by encryption. Because the data compression used is applied end-to-end, encryption may be performed after compression

so there is no conflict between the two operations. For certain coding techniques and applications, encrypting only the HP segments may provide sufficient confidentiality.

The described techniques do not introduce any significant additional non-uniformity in the receiver side computational complexity for packet processing to cause a potential denial-of-service threat.

References

- [1] Glenn L. Cash, Mehmet R. Civanlar, "Method Of And Apparatus For The Transmission Of High And Low Priority Segments Of A Video Bitstream Over Packet Networks," United States Patent Number: 5,481,312, Jan. 2, 1996.
- [2] Glenn L. Cash, Mehmet R. Civanlar, "Video Bitstream Regeneration Using Previously Agreed To High Priority Segments," United States Patent Number: 5,510,844, April 23, 1996.
- [3] Hoffman, D., Fernando, G., Goyal, V. and M. Civanlar, "RTP Payload Format for MPEG1/MPEG2 Video", RFC 2250, April 1997.
- [4] M. R. Civanlar, G. L. Cash, "A practical system for MPEG-2 based video-on-demand over ATM packet networks and the WWW," Signal Processing: Image Communication, no. 8, pp. 221-227, Elsevier, 1996.
- [5] ISO/IEC International Standard 13818; "Generic coding of moving pictures and associated audio information," November 1994.
- [6] Overview of Reliable Multicast Protocols Web Page, URL http://gaia.cs.umass.edu/sigcomm_mcast/talk1.html.
- [7] Mogul, J. and S. Deering, "Path MTU Discovery", RFC 1191, November 1990.
- [8] M. F. Speer, S. McCanne, "RTP Usage with Layered Multimedia Streams", Work in Progress.
- [9] Schulzrinne, H., Casner, S., Frederick, R. and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", RFC 1889, January 1996.

Authors' Addresses

M. Reha Civanlar
AT&T Labs-Research
100 Schultz Drive
Red Bank, NJ 07701
USA

Email: civanlar@research.att.com

Glenn L. Cash
AT&T Labs-Research
100 Schultz Drive
Red Bank, NJ 07701
USA

Email: glenn@research.att.com

Barry G. Haskell
AT&T Labs-Research
100 Schultz Drive
Red Bank, NJ 07701
USA

Email: bgh@research.att.com

Full Copyright Statement

Copyright (C) The Internet Society (1998). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

