PET - Priority Encoding Transmission: A New, Robust and Efficient Video Broadcast Technology

Bernd Lamparter   Andres Albanese
Malik Kalfane   Michael Luby *

TR-95-046

August 1995

Abstract

This paper presents a new Forward Error Correction Scheme with several priority levels. It is useful for applications dealing with real-time transport streams like video and audio. Those streams consist of several data parts with different importance. PET allows to protect those parts with appropriate redundancy and thus guarantees, that the more important parts arrive before the less important ones.

In the video we show the impacts of losses to an MPEG [2] video stream with and without PET protection. Due to the fragile nature of MPEG the unprotected stream breaks up, the PET protected stream is unaffected by low losses and jerky when high losses are present.
Keywords: Video Transmission, Forward Error Correction, MPEG.

*all International Computer Science Institute, Berkeley, CA 94704, USA

1. Background

Many current and proposed telecommunication based applications rely on broadcast of multi-media information over packet-switched networks. In the past protocols with packet retransmission were used. But in real-time applications retransmission is not a practical solution because of the unpredictable additional delay. Priority Encoding Transmission, or PET, helps overcome two major problems encountered by such applications:

1. Packet losses are unavoidable, irregular, and unpredictable in nature.

2. The wide variety of networks with different capacities, and end-station receivers with different processing power limits the transmission rate to the lowest capacity network and lowest power workstation during a multicast session.

2. What PET does

Using advanced coding techniques, PET improves the quality of transmission over packet networks that have unpredictable losses. PET encodes and stripes the information from long messages into many shorter packets. PET encodes the data using a unique multi-level forward-error-correction scheme that provides graceful acquisition and degradation of information.

PET makes possible something that has not been practical due to bursty losses: good quality video transmission over lossy packet networks (such us Internet) using MPEG-1 and MPEG-2. These bursty losses have forced the data communication industry to broadcast video using motion JPEG, which requires several times the transmission rate of MPEG. PET, by allowing the use of more efficient techniques such as MPEG, will greatly reduce the overall transmission rate requirements despite the fact that it introduces a minimum increase in the MPEG transmission rate.

A version of PET has been implemented and applied to the transmission of MPEG video over the Internet. The demo shows video sequences consisting of a transmission over lossy packet networks using the standard MPEG-1 encoding with and without PET. In typical examples there is a dramatic improvement of picture quality due to PET, and yet PET increases the overall data sent by only a small amount, e. g., 20% over the length of the original MPEG data, Typically, this approach represents a significant reduction in the transmission rate, e. g., five-fold, compared to what would be needed by JPEG to achieve comparable quality.

PET can be used with any packet-switched-network hardware or protocol. It operates as a layer above the transmission of packets (e. g. ATM cells, etc.). PET also makes it possible to broadcast to heterogeneous networks and to receivers with widely different capabilities. This allows the network and receiving stations to discard arbitrary packets which cannot be handled due to limited capacity or processing power.

3. How PET works:

The key idea is an encoding scheme that will recover the most important information of a message no matter which packets are lost or intentionally deleted in a transmission. By the judicious application of erasure coding, PET achieves greatly improved performance, enabling statistical multiplexing of video images. This is done without adding excessive overhead in terms of network capacity or delay, and effectively decreasing transmission cost per channel.

The general idea of this approach is the following. Based on a user (or application) specified priority function, a message is partitioned into blocks of different priorities. For example, in MPEG a message could correspond to a group of pictures and blocks could correspond to frames (I, P, B, ...) of different priorities. Multiple packets are formed by computing a redundant encoding for each block, and then one word from each of the codes is put into each packet (a method also known as striping). Our approach goes beyond forward error correction and striping because it adds multi-level redundancy according to the priority level of the information in each block.

While PET can take advantage of special purpose encoding/decoding hardware, the implemented examples, e. g. real-time video transmission, have demonstrated effective operation without such hardware, i. e. with software implementations using conventional computers.

4. Implementation

PET was implemented within the framework of VIC, a video conferencing tool [3] for the MBone, the Multicast Backbone on top of the Internet. It uses a modification of the software only MPEG decompressor developed by Stefan Eckhart and the MPEG Software Simulation Group at the University of Munich, Germany.

The extended VIC can capture MPEG videos from a sunvideo board or a file. The redundancies for the priority levels may be adjusted by the user according to the expected losses. Each group of picture (GOP) is mapped to one PET message. There is a separate level for the GOP header and the I-, P-, and B-frames with decreasing redundancy. To encode the GOP VIC has to wait for the last frame a delay equal to the duration of a GOP.

This is a specific approach to MPEG but the concept is generic and can be applied to images, audio, and video with many encoding standards. We have implemented a library for PET to support other applications and PET encodings.

With the prioritization used in the demo, the video becomes jerky in the presence of high packet loss. We are working on other ways of using PET to encode MPEG images, e.g., prioritizing the DCT matrix frequency coefficients according to their level of importance. This approach is also useful for JPEG images or movies. With this approach, the video rate is smooth, but instead the quality of single frames varies according to the losses.

5. Conclusion

The demo shows MPEG videos with and without PET protection. As can be seen, it is very disturbing to view an unprotected MPEG sequence that has suffered losses, whereas PET can often prevent any damage, or at least provide a graceful degradation.

In analog (AM) transmission systems, and for channels of equal signal-to-noise ratio, the presence of noise is less noticeable for video than for audio and data. In contrast, video is highly compressed when sent using digital transmission systems, and thus bit errors and packet losses are more noticeable for video than for audio and data. Based on this, the erroneous assumption has been that digital video requires transmission with high quality of service (and high cost). This demo of PET shows that this assumption is a misconception, i.e., it is possible to send digital video over a best effort service (with packet loss and low cost). PET provides a graceful quality degradation when unpredictable packet losses are experienced.

References

[1] A. Albanese, J. Blömer, J. Edmonds, and M. Luby and M. Sudan: "Priority Encoding Transmission", *Proceedings of 35th Annual Symposium on Foundations of Computer Sciences, Santa Fe, New Mexico, Nov. 20-22, 1994, pp. 604-612.*

[2] Didier Le Gall: "MPEG: A Video Compression Standard for Multimedia Applications", *Communications of the ACM 1991*

[3] Steven McCanne and Van Jacobson. *vic:* "A Flexible Framework for Packet Video", *Proceedings of ACM Multimedia'95, Nov. 1995*

[4] The PET project at ICSI: URL: http://www.icsi.berkeley.edu/PET/