Generic Uneven Level Protection Algorithm for Multimedia Data Transmission over Packet-Switched Networks

Adam H. Li, Jay Fahlen, Tao Tian, Luciano Bononi, So-Young Kim, Jeong-Hoon Park, and John Villasenor

Abstract

To achieve more efficient usage of channel bandwidth and provide better protection for the media payload transmitted over lossy packet-switched networks, we introduced a new scheme of generic uneven level protection (ULP) forward error correction. The scheme provides different protection levels for data of different significance within a packet. The ULP scheme is designed to be independent from the nature of the media that it protects, and it is very flexible for any protection configuration the user might need without using any out-of-band signaling. Simulation using a video stream transmitted over a lossy packet-switched network shows that the ULP algorithm achieves significant gain for the quality of the transmission over a wide range of network conditions.

Keywords

Forward Error Correction, Uneven Level Protection, Uneven Error Protection, Real-Time Data, Multimedia Transmission, Packet-Switched Network

I. Introduction

WITH the recent phenomenal development of the Internet, transmission of real-time data over packet-switched networks is an increasingly important direction for applications. For example, Internet telephony has become such an important service that it could largely replace the circuit-switched telephone in the future. Another important trend that can not be neglected is the newly emerging Third Generation wireless communication systems. Initially based on the current wireless telephone networks, these systems are evolving steadily from the traditional circuit-switched configuration toward packet-switched architecture. Eventually, there will be large scaled All-IP-based wireless communication systems to provide unified real-time and non-real-time data services.

Transmission of real time data in these circuit-switched network environments has generated much interest in recent research, and there are many remaining challenges to be answered. For Internet transmission, the current best-effort delivery is not designed for transmitting real-time data. The data lost from congestion and other problems in the network can often cause poor speech quality for the Internet telephony services. For wireless networks, the higher transmission error rate related to the wireless links generates much higher data loss compared to the wired connections. Particularly, this loss of data in transmission poses an even greater threat to video streams than to the audio streams, largely due to the intrinsic dependency in time domain among the data packets from the nature of most video coding algorithms. Therefore, the improvement of the transmission schemes to compensate for the end-to-end data loss is a very important problem that must be solved for the successful development of video and audio services over packet-switched networks.

In general, there are two ways to recover lost packets: retransmission and forward error correction (FEC). Due to the real-time nature of many multimedia applications, they have much more strict delay requirements than those of the transmission of normal non-real-time data. As a result, FEC is

Adam H. Li, Jay Fahlen, Tao Tian, and John Villasenor are with the Electrical Engineering Department, University of California, Los Angeles, CA 90095, U.S.A. (email: {adamli, jfahlen, villa}@icsl.ucla.edu).

Luciano Bononi is with the Department of Computer Science, University of Bologna, Italy (email: bononi@cs.unibo.it).

So-Young Kim and Jeong-Hoon Park are with the Samsung Electronics Inc., Suwan, Korea (email: {soyoung_kim, jeonghoon}@samsung.com).

the natural choice for protecting the real-time multimedia data. There are generally three types of FEC techniques [1]: media-independent FEC, media-specific FEC and interleaving. Each technique has its own characteristics, strength and optimal application scenario. Media-independent FECs are usually generic algorithms (i.e., independent to the media format of the payload packet they are protecting) and can provide exact repair for lost packets. However, they usually can not achieve very efficient usage of the channel bandwidth compared to media-specific FECs, which employ knowledge of a media stream format to protect the payload data. By knowing exactly which portions of the stream are most important to protect, the media-specific FEC can apply uneven error protection to the more important bits. Of course, the most obvious disadvantages of media-specific FEC, media dependency and lack of compatibility, also root from the fact that they are tailored for a specific media format. Interleaving generally is not suitable for transmission over packet-switched networks.

In this paper, we describe a generic uneven level protection (ULP) forward error correction scheme. It is a generic algorithm that does not depend on the type of payload it protects. The scheme is very flexible in order to provide various protection levels for data of different importance in the media stream, and it can achieve better channel utilization at the same time. We also performed simulation on the performance of ULP over typical lossy packet-switched channels. The simulation results shows that ULP can bring a significant and consistant gain in objective and subjective qualities of the multimedia payload. In section II, we discuss the background and demonstrate the importance of generality, compatibility, and efficient channel capacity utilization of various FEC schemes. Section III gives the principle of the ULP scheme as well as the usage, advantage of it, and gives an example of how such algorithms can be implemented in protocol. Section IV presents the measurement methods, the simulation conditions, and the results from the experiments. Section V summarizes our discussion.

II. Generality of media-independent FEC and efficiency of media-specific FEC

The basic way to execute forward error correction is to send redundant information along with the data that need to be protected (in separated packets or even in separate channels) and to repair the data from these redundancies in case some packets are lost in transmission.

There are many ways to perform media-independent forward error correction [2], [3] for the real-time data which is usually transmitted by using the real-time transport protocol (RTP) [4]. To illustrate the characteristics of these media-independent FEC techniques, we take the generic forward error correction scheme as described in RFC 2733 [2] as an example. The media packets are protected by parity packets generated using simple exclusive OR (XOR) parity operations.

This procedure is illustrated in Figure 1. The sender of the media packets takes a group of N media packets (N = 5 in the illustrated case) that need to be protected and applies the xor operation over the components of the headers and the data payload of the packets. Any packets that are smaller than the largest packet of the group will need to be padded to the length of that packet. This generates an FEC packet which is then sent to the receiver. If any one of the N media packets is found lost at the receiving end, the receiver can recover this lost packet completely from the N-1 received packets and the FEC packet (providing that the FEC packet arrives intact) by reversing the FEC packet generating process. The group size N determines how much protection is applied to the packets. The smaller the N is, the more protection the media packets get, since there is more chance that a lost packet can get recovered. But the tradeoff for a smaller N is the extra bandwidth needed, because more FEC packets are generated.

As the name implies, the media-independent schemes are generic in the sense that they do not care about either the type of media streams they protect or their format. These schemes also have great compatibility so that terminals capable of performing such schemes can communicate well with the terminals that do not support such FEC schemes. The extra FEC packets can simply be ignored, while the media stream can be decoded the same way as before. This may not be such a big issue for point-to-point communication, where the terminals should find out the capacity of the other party during

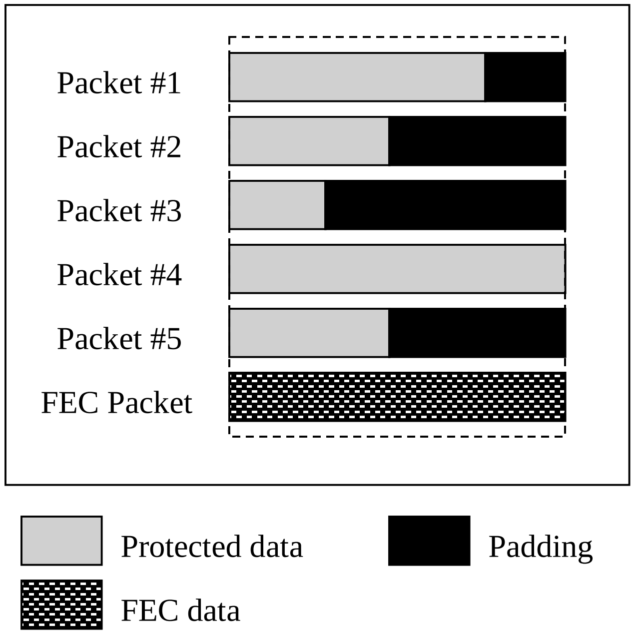

Fig. 1. Generic media-independent FEC as in RFC 2733

capability negotiation, and not send the FEC packets if the other party is not capable of decoding them. However, the compatibility is a great advantage for multicast/conference scenario where one-on-one negotiation is not possible, and terminals in the conference can have different channel conditions, and some of the terminals may have the capability and do need extra FEC protection.

In order to make the scheme media-independent and to achieve complete recovery of the lost packets, padding has to be applied to all packets but the longest one in the group. This practice compromises the efficiency of channel bandwidth usage. For many media payloads, particularly video streams (which are intrinsically variable-rate streams), the packet sizes can vary greatly from one to the next. The padding can be a big waste of the bandwidth resource for many applications, especially in mobile environments, where the channel bandwidth is usually a very scarce resource.

The media-specific FEC generally can achieve much more efficient repair of the media stream by employing the knowledge of the compression scheme of the media payload. With such knowledge, it is possible to apply the FEC protection to the most significant bits of the stream, rather than over the entire packet. These uneven error protection schemes can achieve better bandwidth utilization, and are very suitable for protecting the multimedia streams which have their unique error tolerance characteristics.

While uneven error protection used in the media-specific FECs can provide more efficient protection to the payload media stream by applying the knowledge on the specific stream format it protects, its shortcoming also roots from exactly the same aspect. Namely, each of media-specific schemes is designed specifically for the particular media format it protects. These schemes lack the generality that the media-independent FEC can provide. They also do not have the same compatibility as the latter, i.e., all the senders and receivers are required to be capable of the same protection scheme, regardless of the channel condition. This may not be the most efficient and is a quite strict requirement for terminals in the multicast/conference scenario.

So would there be a way to design a forward error correction scheme that is generic and has the compatibility that the first type of scheme provides, while at the same time is also more efficient in bandwidth usage by applying uneven error protection using the knowledge on the characteristic of the media streams?

III. Principle and advantage of the ULP FEC scheme

In many cases for multimedia streams, we have some very important knowledge about the stream that is valid for almost all media stream forms. *In general, the more important parts of the data are always at the beginning of the data packet.* This is a common practice for most codecs, and there are many reasons behind it. One of the obvious reasons is that the beginning of the packet is closer to the re-synchronization marker at the header and thus is more likely to be correctively decoded, especially if the data is variable length coded. The following analysis of some media formats will illustrate how the uneven error protection with the emphasis at the beginning of the packet can more effectively protect the media payload.

For video streams, most modern formats have optional data partitioning modes to improve error resilience, where the video macroblock header data, the motion vector data and DCT coefficient data are separated in their own individual partitions. In ITU-T H.263 version 3 [5] when the optional data partitioned syntax of Annex V is enabled, and in MPEG-4 Visual [6] Simple Profile when the optional data partitioning mode is enabled, video macroblock (MB) headers and motion vectors are transmitted in the partition(s) at the beginning of the video packet while DCT coefficients of the residue error are transmitted in the partition close to the end of the packet. Because of the intrinsic inter-frame predictive coding scheme used in most video formats, the motion vector data is much more important for the quality of video frame reconstruction. Studies on video coding have shown that when transmission loss is inevitable, the pictures obtained from motion vectors alone (without DCT coefficient) are far superior than pictures reconstructed with DCT coefficient but have motion vector data lost [7]. We can see that the data in these video coding formats are arranged in the order of more important data to less important data. Thus, the image quality should increase from applying more protection to the beginning part of the packet.

In case of audio streams, many new audio codecs do encode into bitstream data of different importance classes and transmit the data in the order of more important to less important. Applying more protection to the beginning of the packet would help. Even for the uniform-significance audio streams, special stretching techniques can be applied to the partially recovered audio data packets. Also, for the audio streams with audio redundancy coding, it makes sense to have more protection applied to the original data which is at the first half of the packet and little or no protection for the redundant copies which are at the trailing half of the packet.

The real time data transmission would benefit from unequal error protections schemes with more emphasis on the beginning part of the packets. Using this general knowledge of the real time media streams, we have designed a generic uneven level protection (ULP) scheme which would possess the generality and compatibility while offering the efficiency of the unequal error protection.

An example of ULP is illustrated in Figure 2. This is a simplified case of single level ULP, where there is only one protection level. The data of length L in the beginning of the packets are protected by the ULP FEC packet, which is the XOR result of all the protected portions of the packets.

The protection of the media stream data can be tuned by adjusting L for the amount of data to be protected, and by adjusting N for the resilience of the protection. With carefully selected N and L values, the bandwidth can be more effectively utilized to protect the more important parts of the data with little or no bandwidth for protection of the less important portion of the packet and the padding bits which carry no information. The determination of proper N and L is out of the scope of this paper and will be studied separately. The scheme is very flexible to accommodate the different configurations and requirement tradeoff for vastly different scenarios.

Because it is always possible to select the length of longest packet in the group as the protection length L, the ULP scheme should always be able to offer an equal or better channel utilization and protection tradeoff than the generic media-independent FEC scheme shown in Section II. The intelligent choice for these values is important for the most efficient utilization of the bandwidth for the protection. As a general guideline, the L should be chosen so that the important portion of data is

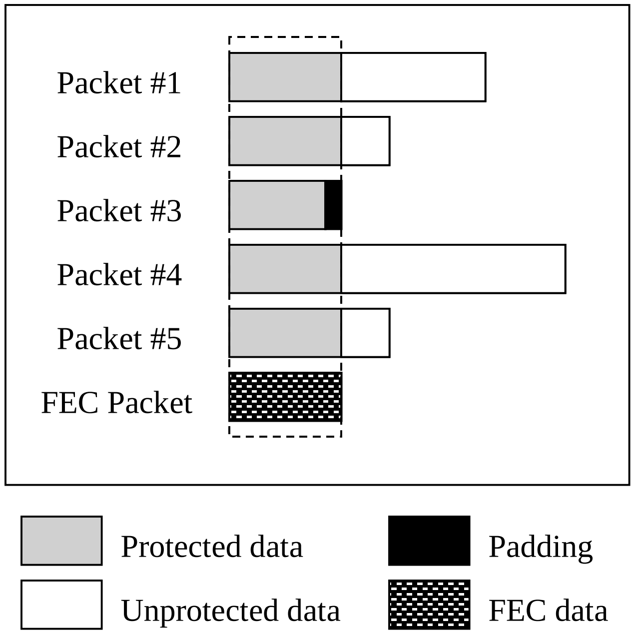

Fig. 2. Generic uneven level protection FEC (Single Level)

covered in most of the cases (and could be different from stream to stream). For channels with higher error rates, the N value could be chosen smaller to give higher level protection. These parameters in the scheme should be left to be determined by users or application (dynamically) for the specific transmission channel and the specific real-time media content.

The ULP scheme uses exclusive OR operation, and is very similar to the media-independent FEC scheme illustrated in Figure 1 except that only the data of length L from the beginning of the packet is protected. Thus, it is also completely generic to the media payload format. This scheme is transparent to the codecs. Every packet is treated as a generic data block.

The complete algorithm of the ULP FEC includes multi-level protection, with the one-level protection (as illustrated in Figure 2) as a special case. Protection of higher levels could be used to protect the less important data that is further back from the beginning of the packet. In general, the higher protection levels use larger grouping of N to generate the ULP protection compared to lower protection levels. Thus the high levels provide less error resilience capability and take less channel capacity, because the data they protect are generally of lesser importance. This multi-level scheme makes the ULP complete and more flexible in order to meet the requirements of different applications. Special care needs to be taken with the algorithm scheme for multi-level protection. The grouping information can be transmited in the grouping table in-band within the ULP FEC packet [8]. For each protection level in each ULP packet, a (bitmap) table is included to signal which packets are protected by the current ULP at this level. The protection length of each protection level in every ULP packet is also specified within the packet. This results a very flexible scheme which can accommodate varieties of grouping choices for different application scenario. It not only avoids the problem of transmitting the extra out-of-band information, it also enables dynamic determination of the protection level in order to adapt to different media streams.

Figure 3 is an example which illustrates the flexibility of such a complete multi-level ULP scheme. In this example, five media data packets are sent through the channel. The data packets #1, #2 and #3 are protected by ULP packet #1 at level 1. The data packets #4 and #5 are protected by ULP packet #2 at level 1. At the same time, ULP packet #2 also provides protection at level 2 to data packets #1, #2, #4, and #5.

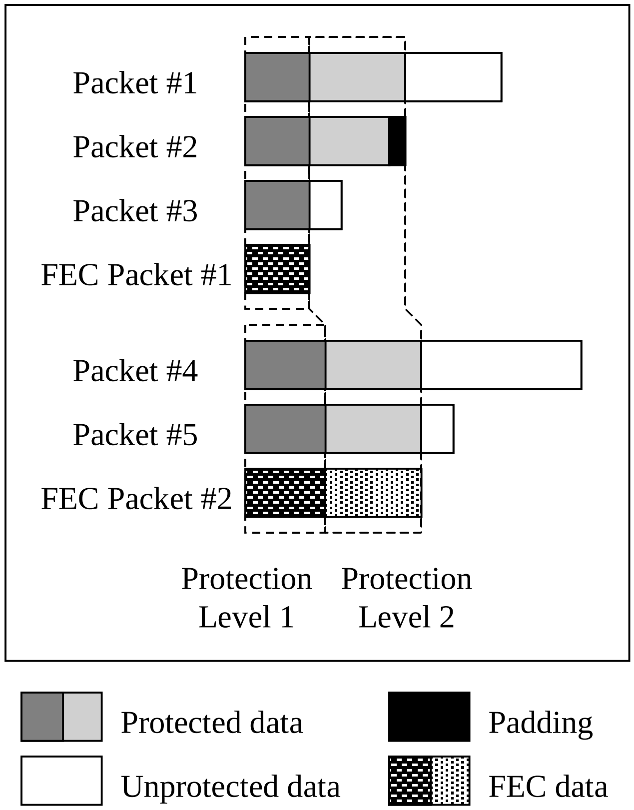

Fig. 3. Generic uneven level protection FEC (Multi-Level)

IV. Simulation

In order to verify the effectiveness of the ULP algorithm in a lossy packet-switched networks, we performed simulations to measure the performance of the ULP scheme.

A. Performance Measurement

An objective measurement method is necessary to measure the performance of the protection schemes. Take for example, if we choose SNR to be the measurement for the performance, a protection scheme can improve the SNR of the result in most cases. However, a protection scheme will also take additional bandwidth to transmit the extra protection information. So, a rate-distortion curve or rate-SNR curve is necessary to measure the performance gain for such schemes.

As illustrated in Figure 4, suppose we have a rate-distortion curve for the unprotect payload data. The horizontal axies denotes the bandwidth used by the payload, and the vertical axies denotes the quality measurement of the payload data (such as SNR). The curve should be monotonically increasing, since the more bandwidth the payload takes, the better quality it can get.

Now, let us take a point (for example, Point A) from the curve. When the protection scheme is applied to the payload, the point will be shifted toward right-hand side (because of the extra bandwidth the protection is taking) and upward (because the quality of the payload will be improved from the pretection it gets).

In the example case of the example algorithm 1 in Figure 4, the resulting point (Point B) lies below the original curve. This shows that even though that example algorithm 1 improves the quality of the payload to some degree in this error channel, it is not efficient enough to be beneficial. For the case of the example algorithm 2, the result point (Point C) is above the original curve. This can demonstrate that example algorithm 2 is an efficient algorithm that can improve the quality of payload in error environment. The approximate gain can be measured vertically form the curve to the resulting point.

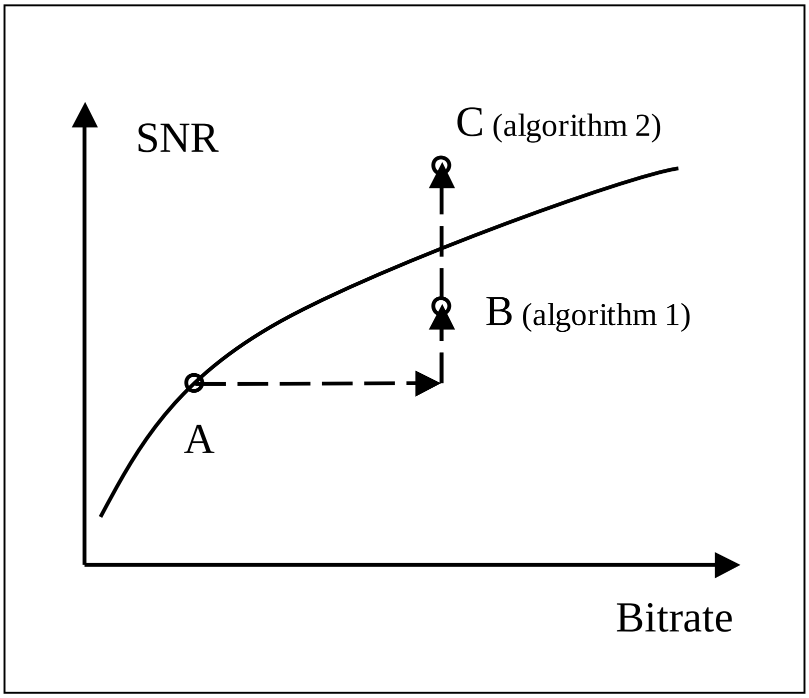

Fig. 4. Rate-Distortion curve for measurement of the protection scheme performance

In conclusion, in order to show the effectiveness of a protection scheme, over an error channel, it is not enough just to show that the scheme improves the quality of the payload, but it is necessary to demonstrate that that the scheme can improve the quality above the original rate-distortion curve.

B. Simulation Environment

The simulations were performed over two types of packet-switched channels with MTU (Maximum Transfer Unit) size of 1500 bytes and 125 bytes. The 1500-byte MTU channels are what usually can be found in typical ethernet-Internet applications. The 125-byte (or 1000-bit) MTU channels are used to simulate a typical wireless network where the packets usually have smaller sizes ranging from a few hundred bits to a few thousand bits. For both of these two common channel types, we simulated the typical channel conditions with packet-loss rate of 1%, 2%, 5% and 10%.

The payload data used in this simulation was H.263 version 3 [5] video bitstreams. The encoding options used to generate these video bitstreams included Annex I, J, T and V of H.263 version 3. They are Advanced Intra Mode, Deblocking Filter Mode, Modified Quantization Mode and Data Partitioned Slice Mode. Each sequence used were generated from original video sequence of 30000 frames, which roughly correspond to video clip of over 16 minutes long. Such long video sequence were used to ensure objective measurement of the resulting quality.

The quality of the received payload was measured by calculating the PSNR (Peak Signal-Noise Ratio) of the decoded video sequences. The PSNR wass calculated between each and every frame of the source sequence (at full frame rate) and the corresponding reconstructed frame, as prescribed in the common testing condition for video performance evaluation in Advanced Video Expert Group [10].

The testing video sequences were encoded with bitrate setting of 32, 40, 48, 56, and 64 kbps. The bitrate used to plot the rate-distortion curve were the measured bitrate from the actual encoded bitstream.

C. Simulation Methods and Results

The encoded video sequence data packets were then sent though a simulated channel. Random numbers were generated to simulate packet lost (for both the data packets and FEC packets) with the simulated packet loss rate.

Performance were measured for two scenarios: one-level and two-level ULPs. For one-level ULP, the packets were grouped according to the N value selected, and a length of about 20% of the average packet length was protected. For two-level ULP, the packets were protected at level-1 grouped at the N value selected, and were protected at level-0 with group size of N/3 (which indicates a higher degree of protection). The protection length at level-0 was about 20% of average packet length. The N values used in the ULP were 24, 12, 6, and 3 respectively.

For the packets that arrive correctly, the content of the data packets was transferred to the receiving file. For the packets that were lost in transmission, the simulation program checked to decide if a lost packet can be recovered with the FEC protection. If the lost packet was within the protection, the packet will be recovered (partially if recover at level-0 in this case, and completely if at level-0 and level-1) according to the scheme and protection levels. If the packet lost was beyond the protection, the content was transferred. Any lost packet or part of it that was not recovered was replaced with stuffing bit 1's in the received file.

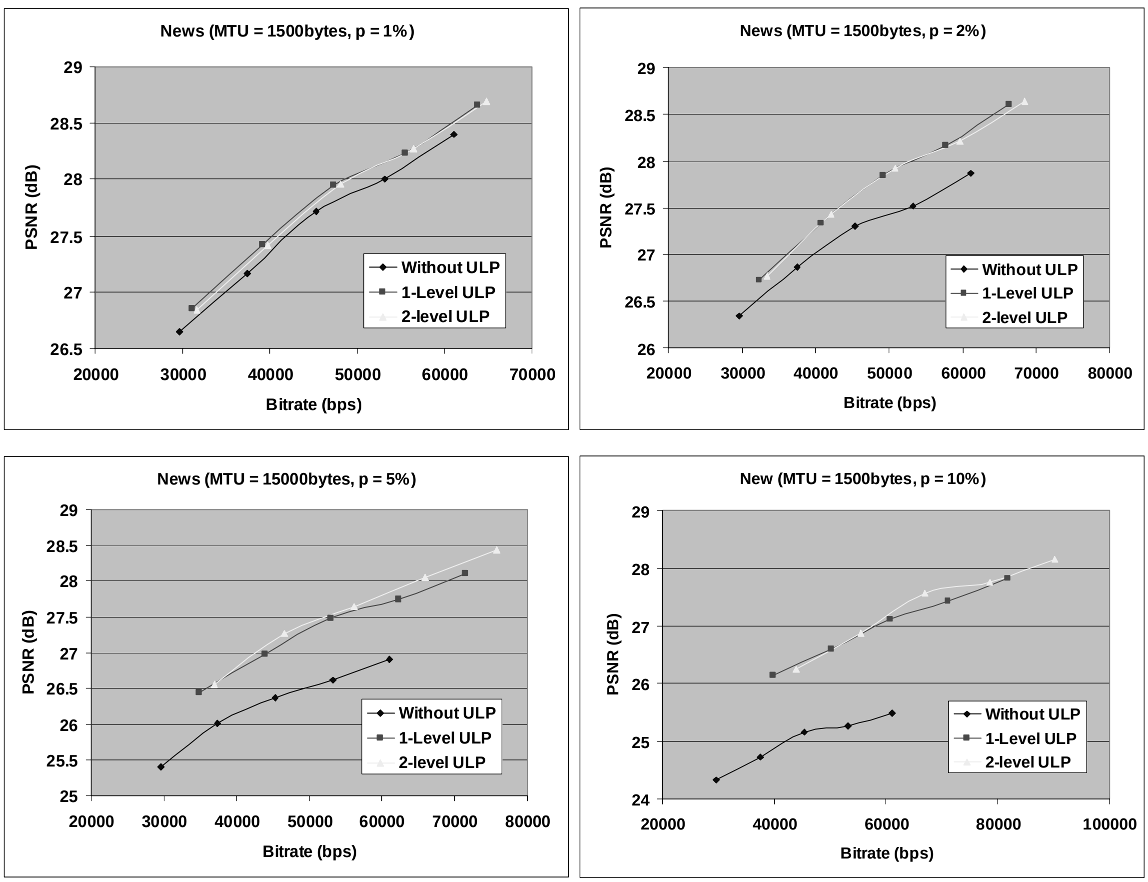

Fig. 5. Rate-Distortion curve for measurement of 1500 byte MTU

We performed the simulation over several video sequences. Due to space limitation here, we show only the results for the "news" sequence. Figure 5 is the rate-distortion curve for the simulations of 1500-byte MTU channels, and Figure 6 is of 125-byte MTU channels. The performance gain from ULP depends on the channel packet-loss rate as one would expect. For the channels tested, the average improvement in PSNR ranges from 0.2 - 0.5 dB at 1% packet loss rate to up to 2 - 3 dB at 10% packet

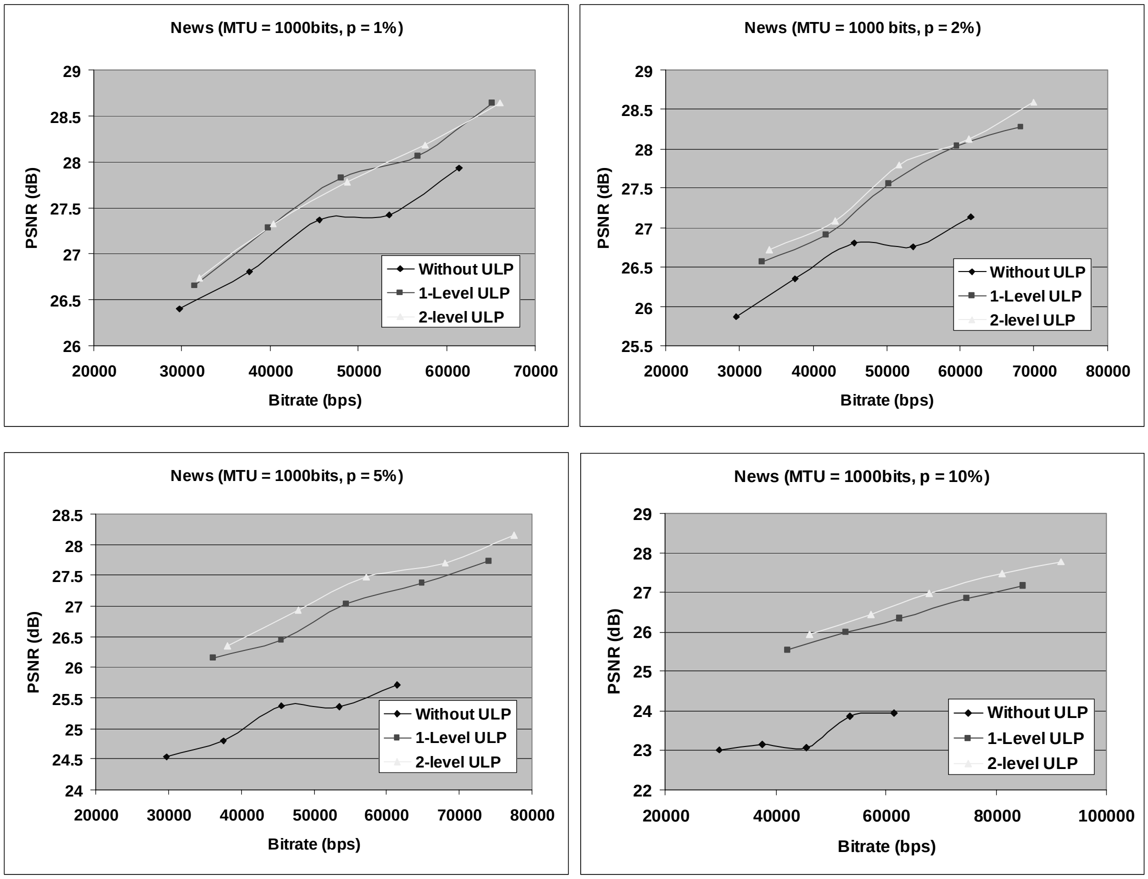

Fig. 6. Rate-Distortion curve for measurement of 125 byte MTU

loss rate, respectively. This is significant gain, and great improvement in the subjective visual quality is obtained. The gain in also very consistent over the range of bitrate tested as shown in the figures. These performance measurement methods, simulation conditions [11] and results [12] have also been summarized in presentations to IETF and technical reports to ITU.

We also note that the grouping size used in these simulations for the ULP protection levels were chosen without a thorough search. It is possible to further improve the quality by more careful analysis of the channel characteristics, and then choosing the optimum parameters.

V. Conclusion

In this paper, we described the generic Uneven-Level Protection (ULP) forward error correction scheme. Utilizing the knowledge that almost all video and audio media packet formats have more important data close to the beginning of the packet, the ULP FEC scheme differentially applies more protection to the more important parts of the packet, while less protection is used for the less important data. This provides better allocation of the channel capacity for the protection operation and achieves more efficient use of the bandwidth. The ULP FEC packets that are sent along with the media stream enable recovery of the media packets lost in the transmission, either completely or partially, depending on the situation of which packets are lost and the grouping of the protection operation. Because of the simple parity (exclusive OR) operation used in generating the ULP FEC packets, the scheme described is totally independent of the media format of the protected packets and can have total compatibility with non-FEC-enabled terminals. Its simplicity, generality, compatibility and the capability to achieve more efficient use of the channel capacity could greatly improve the transmission of multimedia content over error-prone channels. Simulations with actual video sequences were carried out to validate the proposed scheme and measure the improvement of the performance the proposal brings. The simulation results show that the quality of the decoded media stream improves significantly by using the ULP forward error correction technique, even for the casually chosen protection parameters. The gain in performance is consistent over the whole range of the channel bitrate used.

Acknowledgments

The authors would like to acknowledge the suggestions from many people, particularly Matthieu Tisserand, Dong-Seek Park, and Stephen Wenger. The authors also would like to thank Samsung Electronics for funding this project.

References

[1] C. Perkins and O. Hodson, *"Options for Repair of Streaming Media"*, Request for Comments (Informational) 2354, Internet Engineering Task Force, June 1998.

[2] J. Rosenberg and H. Schulzrinne, *"An RTP Payload Format for Generic Forward Error Correction"*, Request for Comments (Proposed Standard) 2733, Internet Engineering Task Force, December 1999.

[3] J. Rosenberg and H. Schulzrinne, *"An RTP Payload Format for Reed Solomon Codes"*, Work in progress.

[4] H. Schulzrinne, S. Casner, R. Frederick and V. Jacobson, *"RTP: A Transport Protocol for Real-Time Applications"*, Request for Comments (Proposed Standard) 1889, Internet Engineering Task Force, January 1996.

[5] International Telecommunication Union, *"Video Coding for Low Bitrate Communication"*, ITU-T Recommendation H.263 Version 3.

[6] International Standardzation Organization / IEC *"Information Technology - Generic Coding of Audio-Visual Object"*, ISO/IEC Recommendation 14496-1~6.

[7] A. Li, S. Kittitornkun, Y.H. Hu, D.S. Park, and J. Villasenor, *"Data Partitioning and Reversible Variable Length Codes for Robust Video Communications"*, IEEE Data Compression Conference Preceeding, March 2000.

[8] A. Li, F. Liu, J. Villasenor, J.H. Park, and D.S. Park, *"An RTP Payload Format for Generic FEC with Uneven Level Protection"*, Work in progress.

[9] A. Li, J.H. Park, Y.L. Lee, and J. Villasenor, *"Draft text for Generic Uneven Level Protection (ULP) for Annex I of H.323"*, ITU-T Documentation AVD-2095, 2001.

[10] S. Wenger, *"Common Conditions for Video Performance Evaluation in H.324/M Errror-Prone Systems"*, ITU-T Documentation Q15-E-50, 1998.

[11] A. Li, J.H. Park, Y.L. Lee, and J. Villasenor, *"Simulation conditions for H.323 Annex I transmission over error prone channels"*, ITU-T Documentation AVD-2096, 2001.

[12] A. Li, J. Fahlen, T. Tian, S.Y. Kim, J.H. Park, Y.L. Lee, and J. Villasenor, *"Simulation results of ULP for transmission over error prone channels"*, ITU-T Documentation AVD-2097, 2001.