The Hidden Limitations of FEC

By Adi Rozenberg, CTO, VideoFlow

Introduction

The need for forward error correction (FEC)

Packet loss is part of the "normal" behavior of packet networks and in particular of unmanaged networks like the Internet. Packets can be dropped, for example, when one of the nodes in the path between the sender and receiver is congested. Packets can be lost because of the physical connection causing bits to flip from "1" to a "0" (or vice versa). A single bit error in a packet will cause the whole packet to be dropped by the first node receiving the packet.

The Transmission Control Protocol (TCP) is addressing the need for packet loss recovery using the Automatic Receive reQuest (ARQ) mechanism. ARQ requires the receiver to send acknowledge to the sender for each packet received. If the sender is not receiving acknowledge within a pre-defined time, it will retransmit the packet. Acknowledge-based protocols like TCP resolve the packet loss issue but at the cost of lower bit rate compared to User Data Protocol (UDP) and, in particular, when used in error-prone networks like the Internet, cellular, microwave and WiFi.

Since the video quality is directly depended on the video bit rate, the professional broadcast industry is using UDP and particularly the RTP (Real Time Protocol) to deliver video over packet networks. However, UDP/RTP do not include a mechanism like ARQ for lost packets recovery. Packet loss can be a single packet loss event or a burst of packet loss events. Since each IP packet contain up to 7 MPEG packets, the impact of a single IP packet loss on the quality of the video is high. Therefore, several packet recovery technologies are in use including the XOR-based Forward Error Correction (FEC) and the redundant data (Raptor Codes). These technologies designed to recover lost packets by sending redundant information that can be used by the receiver to recover lost packets.

The effect of jitter and lost packets on the video quality

This is how the video should look like

Using error correction technology is essential but which technology to use?

FEC Technologies

The three main FEC technologies used by the broadcast industry are: Pro MPEG FEC (COP3R2), SMPTE 2022 and Digital Fountain's Raptor Codes. The Pro MPEG FEC defines the most efficient row/column XOR FEC for IP video streams. The SMPTE 2022 is also using XOR-based FEC like the Pro MPEG FEC but can use larger matrix sizes for better packet loss recovery. Digital Fountain's Raptor codes are in use by the 3GPP for mobile cellular wireless broadcast and multicast and as well as by the DVB-H standards for IP datacast to handheld devices. The Raptor Codes are more efficient than the Pro MPEG FEC and the SMPTE XOR-based FEC hence requiring fewer amounts of redundant data.

The Hidden Limitations of FEC

Protecting from packet jitter and bit rate fluctuations

Packet loss recovery is mandatory but packet jitter is highly likely to be the root cause for the majority of packets declared lost by the decoder simply because they arrived too late. Pro MPEG FEC, SMPTE 2022, and the Raptor codes cannot protect against packet loss caused by packet jitter and/or bit rate fluctuations. A good example for the effect of packet jitter is a pilot VideoFlow had with a European major basketball league where the game was broadcasted over Ka band satellite. Ka band satellite provides very good bit rate with less than 0.1% of packet loss but the measured packet jitter reached 1.6 seconds! The same phenomenon was observed in cellular networks and WiFi links. In these scenarios, FEC was unable to protect the quality of the video. For that reason, it is a rule of thumb in the broadcast industry to add a 20% bit rate guard band to try and absorb packet jitter and bit rate fluctuations. Unfortunately, adding brute force, extra bit rate cannot protect against the unpredicted momentary magnitude of the packets jitter and bit rate fluctuations in unmanaged networks.

Overhead is expensive

The Pro MPEG FEC, SMPTE 2022, and the Raptor codes require sending recovery information constantly even if there is no packet loss at all. This is waste of link capacity that can be utilized for either increasing the video bit rate (hence increasing video quality) or reducing the cost using lower bit rate link. Based on our experience, a 25x4 Pro MPEG FEC matrix can protect against a packet loss burst of up to 25 packets. A 25x4 matrix protects 100 video packets but require receiving 129 packets comprised of 100 video packets, 25 Row-FEC packets, and 4 Column-FEC packets. The bandwidth overhead will be constantly 29% regardless if there are errors or not. SMPTE 2022 will use 50 FEC packets for the same 100 video packets; a 50% overhead. Digital Fountain's Raptor Codes will need only 26 additional packets to recover a burst of 25 lost packets; a 26% overhead. Together with the 20% bit rate guard band, the total bit rate overhead reaches 50%! The overhead is becoming an issue when broadcasting over expensive connections that may be susceptive to weather conditions, utilization, etc. What if there is a technology that needs no more than 5%-10% overhead?

Live video is defined by a low delay between the sender and the receiver

As mentioned earlier in this article, the Pro MPEG FEC will require adding 29 packets for protecting a stream of 100 packets per seconds. The way the FEC packets are sent requires receiving both the stream packets and FEC packets before completing the building of the FEC row/column matrix. This means that, the receiver will be able to start recovering a lost packet only after receiving all packets (129 packets in our example). Waiting for packets to arrive to build the FEC matrix adds delay to the path between the sender and the receiver. In low bit rate links, this additional delay can be substantial reaching 0.5 second to 1 second depending on the line's bit rate and quality. The delay added by the SMPTE2022 will be significantly higher than the delay added by the Pro MPEG FEC because SMPTE2022 uses larger row/column FEC matrices for better protection. The Raptor Codes are not different in this regard with additional delay similar to the Pro MPEG FEC.

The statistics of the lost packets is working against the ability of the FEC matrix to recover packets. The probability of not resolving the FEC matrix is increasing above 5% packet loss. The reason being is permutations of lost packets that will prohibit resolving the FEC matrix. Let us take as an example the packet loss statistics on the Internet. A good Internet connection has a packet loss of 1%-2%. However, there will be times where the packet loss will increase to 5% with peaks of 7%-10% during busy hours such as before people go to work, when they break for lunch, when they get home, and during the evening when chatting with friends, uploading new pics to Instagram or viewing a YouTube video. Networks like cellular, WiFi, and Satellite are also sensitive to weather conditions, utilization, and geography amongst other factors hence packet loss ratios vary and may reach more than 10%. For example, in the pilot VideoFlow did with a European major basketball league a severe thunder storm had developed over the venue causing severe packet loss to the satellite link.

The impact of packet loss on the ability of FEC to recover lost packet is demonstrated in the below examples. After receiving the packets, the receiver builds the row/column FEC matrix with the packets received. Each row FEC packet can correct one lost packet in a row; each column FEC packet can correct one lost packet in a column.

The simple examples are using a 4 x 4 row/column FEC matrix. A 4 x 4 FEC matrix will generate 8 FEC packets – 4 packets for the rows and 4 packets for the columns. Video packets are designated "V", lost packets "L", recovered packets "R" and FEC packets "FEC".

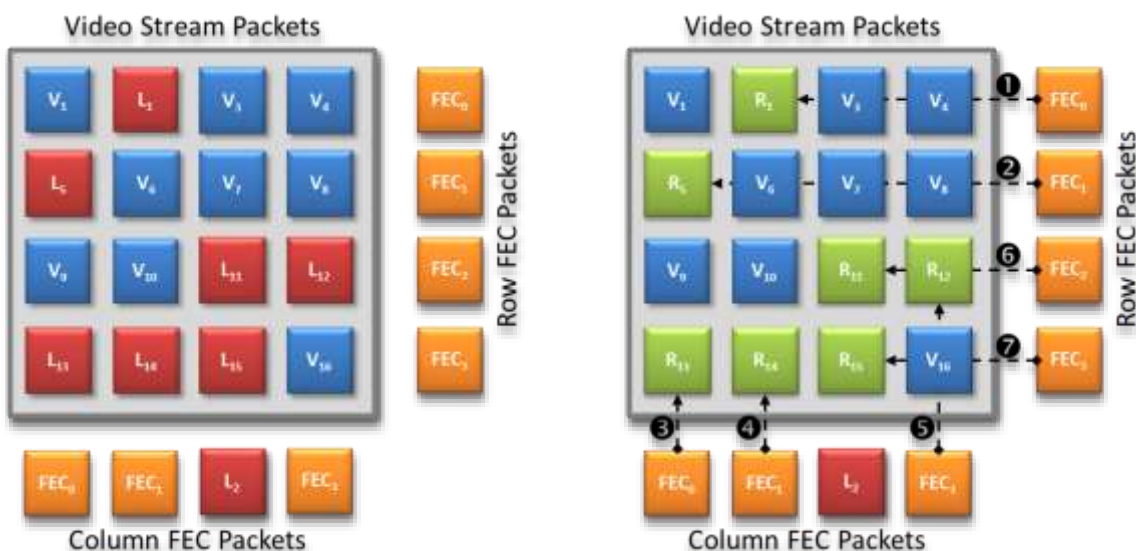The below is an example for FEC successful packet recovery.

Figure 1: Example 1 – Recovering Lost Video Packets at the Receiver

Figure 1 above shows two matrices. The matrix on the left is the FEC matrix arranged by the receiver. The matrix on the right is the FEC matrix after attempting packet recovery. The algorithm starts by scanning the rows first. If there is a single lost packet then the packets are recovered. If there is more than one lost packet then the row is skipped until completing all rows scan. After completing the row scan, the algorithm continues with the column scan. Note that some of the lost packets were recovered in the row scan (R1 and R5) hence it is possible to recover lost packets L13 and L14. Once the column scan search is completed, the algorithm continues with rows scan. Now, that packets L13, L14, L12 are recovered (R13, R14, R12), the row scan can complete recovering lost packets L12 and L15. In this example it took the algorithm 7 iterations to recover 7 lost packets.

Unfortunately, this is not always the case. The below example has the same amount of lost packets as in the previous example, only this time the packet loss statistics is not in favor of the FEC matrix.
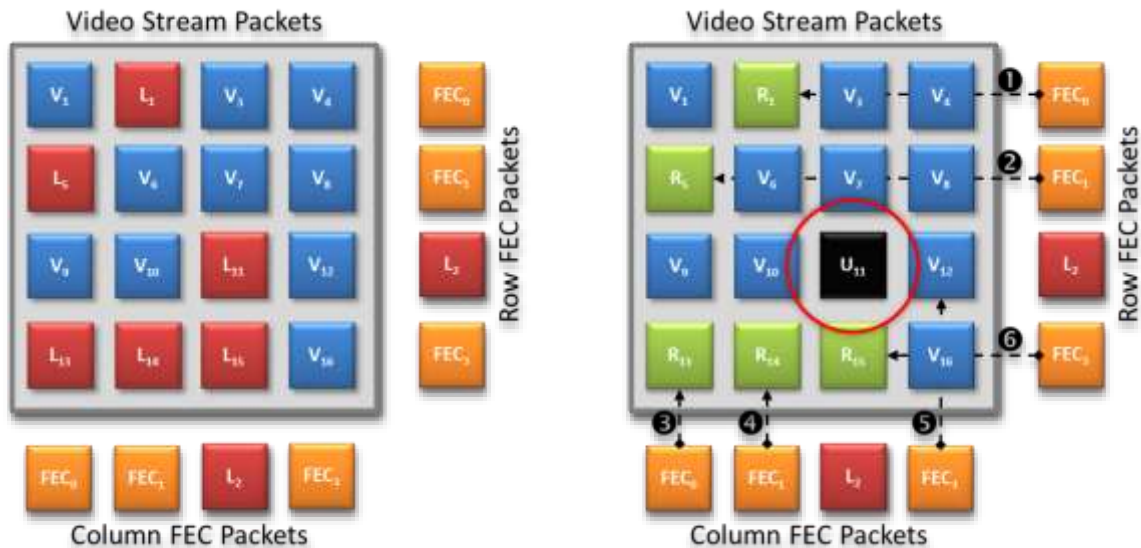
Figure 2: Example 2 – Packet Loss Statistics Yields Unresolved FEC Matrix

The same row/column algorithm is used only this time, lost packet L11 cannot be recovered because both the row FEC and column FEC packets were lost. Same packet loss ratio, different loss statistics and FEC was unable to recover all lost packets.

In the next example, the FEC is unable to resolve the matrix even though **all** FEC packets were received.
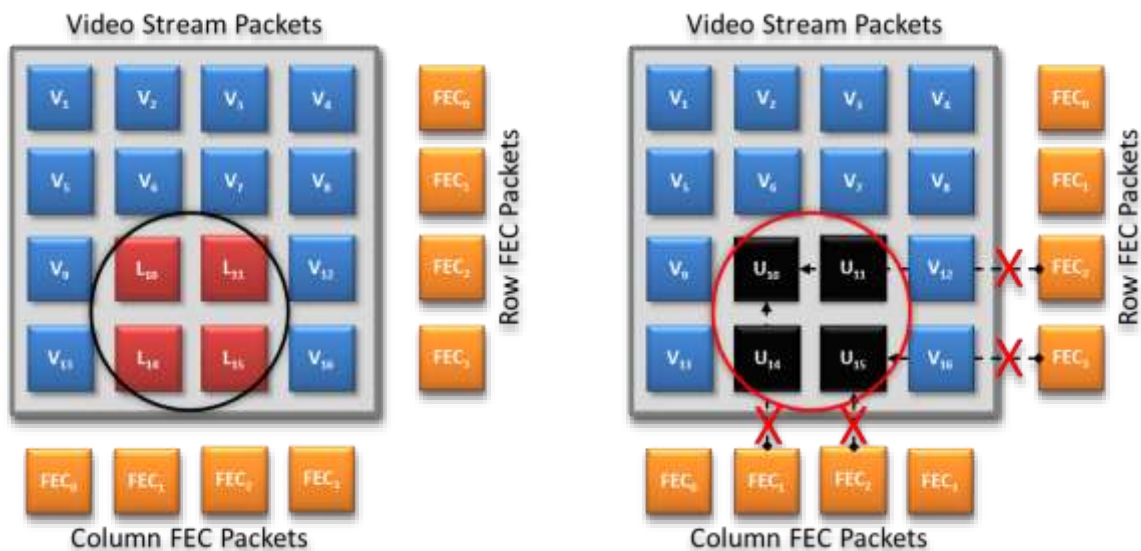

Figure 3: Example 3 – Unresolved FEC Matrix despite of receiving all FEC Packets

Packets L10, L11, L14, and L15 were lost. The burst dropped the packets in a formation that neither the column FEC packets nor the row FEC packets can recover the lost packets because in each row and column there is more than one lost packet. Contrast to example 2, which has no solution, a larger FEC matrix might be able to resolve it in this example but at the cost of higher overhead.

VideoFlow Technology

VideoFlow technology is based on a deep understanding of both video and networking providing the best video quality protection possible. VideoFlow is making use of existing standards and protocols but in a smart way.

It is all about jitter

VideoFlow had realized very early that in many scenarios, the main reason for an annoying video quality is packet jitter and not packet loss. VideoFlow had developed a technology that locks to the stream's bit rate either by using the PCR information (most accurate) or by measuring the packet rate (less accurate). Locking to the stream's bit

rate enables playing out jitter free the video packets at the cost of only 0.5-1sec delay. From the VideoFlow solution's perspective, a decoder is sensing negligible jitter with a delay. The delay can be set to be as low as 20 msec or high as 15 minutes depending on the networks' behavior and/or on the applications' requirements. Now that a buffer for jitter cancelation is set, the VideoFlow solution leverages the same buffer for recovering packets dropped by the network without additional delay.

Leverage the benefits of ARQ

The VideoFlow solution is making use of the ARQ protocol. ARQ requires a communication between the receiver and the sender by which the receiver notifies the sender of any lost packet. Upon receiving the request from the receiver, the sender retransmits these packets to the receiver. The benefit of ARQ over FEC is straightforward. Since only lost packets are resent by the sender, the bit rate overhead is proportional to the network's packet loss. For example 0.1% packet loss will add only 0.1% overhead to the video packet stream while a 7% packet loss will add only 7% overhead to the video packet stream (compare it with a 50% overhead when using Pro MPEG FEC). Since the resent packets can be dropped as well by the network, the actual overhead may be slightly higher. It also means that there might be network conditions where more than one request for resending a lost packet will occur. For that reason, VideoFlow technology sends the request for lost packets in a way that minimizes the number of packet retransmissions. Furthermore, based on the integrated real time analysis, the VideoFlow technology will quickly identify the receipt of recovery packet from the sender and will quickly stop sending requests to the sender to avoid additional retransmissions. In some cases, where the network conditions are bad, it may be worth giving up requests in favor of an overall better video quality by freeing bit rate for the video stream on the expense of the packet recovery stream at the cost of a local glitch to the video quality.

Packet loss can be declared by the receiver either because it was dropped by the network or because it was late to arrive. A normal behavior of unmanaged networks includes a mix of packet jitter, bit rate fluctuations, and packet loss. Therefore VideoFlow technology includes a mechanism that adds a grace time for the packet's arrival time at the receiver. This mechanism reduces the amount of packet requests to the sender because of jitter as it enables some of the late packets to arrive before declared lost by the VideoFlow receiver. On the sender side, the VideoFlow transmitter prioritizes the recovery packets to be sent such that packets with older timestamps will be expedited to ensure they will be received by the decoder before the presentation time is expired.

Conclusion

Pro MPEG FEC, SMPTE 2022, and Raptor Codes can provide a reasonable protection against packet loss. Unfortunately, the overhead required for achieving this protection is constant and quite high ranging from 26% to 70% of the video bit rate. SMPTE 2022 and Pro MPEG FEC in particular will start experiencing issues in recovering packets above 5% packet loss as either FEC packets or video packets can be dropped in a way prohibiting resolving the FEC matrix required for recovering the lost packets. None of the error correction technologies can compensate for jitter hence, in many networks, FEC will simply not protect the video quality. The rule of thumb of adding 20% bit rate overhead to handle jitter or bit rate fluctuations is challenged when delivering video over unmanaged networks.

VideoFlow technology is better suited for protecting the quality of live video over unmanaged networks. It can lock to the video bit rate and to accurately play it out hence cancelling the jitter using a small buffer (low delay). VideoFlow's packet recovery is taking advantage of this buffer for recovering lost packets using ARQ adding a very small overhead with no additional delay. VideoFlow technology provides a complete and robust high quality live video protection at a fraction of the overhead required by FEC.

For more information about all VideoFlow's products please visit our Web site at: www.video-flow.com or contact us at: info@video-flow.com.