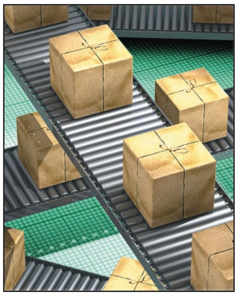

# Not All Packets Are Equal, Part 2

## The Impact of Network Packet Loss on Video Quality

**Jason Greengrass, John Evans, and Ali C. Begen** • *Cisco*

In this second part of a two-part article, the authors highlight the impact that different durations of IP packet loss have on the quality of experience for IP-based video streaming services. They describe the visual impairments that result from such packet losses and present the results of testing and analysis to compare impairments for different loss durations for both MPEG-2-encoded standard and high-definition services.

For consumer video services, real-time streaming applications such as IPTV have stringent service-level agreement (SLA) requirements. We can describe key service-level requirements — commonly specified in SLAs — for IP-based video transport services in terms of delay, jitter, and loss. As we described in the first part of this series,[1] Differentiated Services (Diffserv) IP quality-of-service (QoS) mechanisms can help ensure that service providers meet SLA requirements for delay and jitter for video-streaming services. Such mechanisms can't control packet loss due to lower-layer errors or network connectivity issues, and these network events can result in visual impairments to the video service.

In this final part of our two-part series, we analyze the impact that packet loss has on viewers' quality of experience (QoE) specifically. We also describe the results of video-quality testing we conducted to compare video impairments corresponding to different packet loss durations. Our testing examined both standard-definition (SD) and high-definition (HD) services that were encoded in MPEG-2.

### Video Artifacts

We can define a video artifact as a visible distortion on a displayed video image relative to the original image. In digital video, visual artifacts typically occur due to errors or loss in the transported video bit stream, encoding errors, or too much compression, which discards information required to reconstruct the source image acceptably. We undertook testing to determine which artifacts result from IP packet loss when video is delivered over an IP transport network. We used three video clips from the Society of Motion Picture Television Engineers (SMPTE) for all testing and analysis: a low-motion clip ("Susie") and two high-motion clips ("Football" and "Tennis"). We encoded the clips with a constant video bit rate of 4 Mbps using MPEG-2 Main Profile/Main Level (see http://en.wikipedia.org/wiki/MPEG-2#Video_profiles_and_levels) at a 704 × 480 resolution; we used a 15:2 group of pictures (GoP) at 29.97 frames per second (fps). (See the "MPEG Overview" sidebar for a brief explanation of MPEG video encoding.) Let's look at the artifacts we observed when we introduced IP packet loss into a single program transport stream (SPTS) carrying the Susie, Football, and Tennis clips.

#### Slice Error

Figures 1a and 1b show a slice error, which we see when the network drops at least one IP packet within an I- , P-, or B-frame; Figure 1a shows how a viewer would see the image, whereas Figure 1b shows the output from a video-quality


1089-7801/09/$25.00 © 2009 IEEE

## MPEG Overview

The Motion Picture Expert Group (MPEG) has produced several standards for video compression that providers can use for IP-based video services, including MPEG-2.[1] Standard-definition (SD) television that uses MPEG-2 has a video bit rate reduced to approximately 3.75 Mbps, whereas high-definition (HD) television has a bit rate of approximately 18 Mbps.

During video streaming, an MPEG encoder converts and compresses a video signal into a series of pictures or frames. MPEG-2 has three frame types. Intra, or *I-frames*, carry a complete video picture and are coded without reference to other frames. A received I-frame provides the reference point for decoding a received MPEG stream.

Predictive-coded, or *P-frames*, predict the frame to be coded from a preceding I-frame or P-frame using temporal compression. Finally, bi-directionally predictive-coded, or *B-frames*, use the previous and next I-frame or P-frame as their reference points for motion compensation.

In MPEG encoding, frames are arranged into groups of pictures (GoPs) that include an I-frame and all subsequent frames leading up to the next I-frame. We can describe a regular GoP structure by its total number of frames (that is, the GoP size) as well as by how many B-frames occur between its P-frames — for example, a 15:2 GoP structure has a GoP size of 15 frames and P-frame spacing of two.

### Reference

1. ISO/IEC 13818, Generic Coding of Moving Pictures and Associated Audio Information (MPEG-2), Int'l Standards Organization/Int'l Electrotechnical Commission, 2007.


*Figure 1. Video artifacts. With a slice error, we can see the image (a) as a viewer would see it and (b) with the parts in error highlighted. With a blocking or pixelization error (c), the effect occurs when a loss occurs in either an I- or P-frame. (Source material copyright the Society of Motion Picture Television Engineers, used with permission.)*

analyzer, which highlights the parts of the picture that are in error. If a slice error impacts a reference frame, it might propagate through the rest of the GoP. The error propagation will always clear up once the MPEG decoder receives an unimpaired I-frame, and could clear sooner if, for example, only a B-frame is affected.

### Blocking or Pixelization

We saw a blocking or pixelization effect (see Figure 1c) when a loss of or loss within a reference frame occurred – either an I- or P-frame.

When the video stream loses a complete I- or P-frame and movement is occurring within the sequence, the subsequent frames won't have the information necessary for the decoder to display the correct picture, and pixelization can occur. The more movement a frame has, the greater the degree of blocking around the area in motion. The error will clear once the decoder receives an unimpaired I-frame.

### Ghosting

Packet loss that removes an I-frame or a large chunk of slices within an I-frame near a scene change results in ghosting, as Figure 2 shows.

When the stream loses an I-frame on a scene change, the pre-scene-change picture will continue to appear. As the decoder receives the P- and B-frames from the new GoP, it will add this new information to the old scene, creating the ghosting effect in the old displayed frame. The frame displayed in Figure 2 has the old picture displayed with a blocky green image representing the other

*Figure 2. Ghosting. In this image, I-frame loss near a scene change continues to display the pre-scene-change image. As the decoder receives the P- and B-frames from the new group of pictures, it adds this new information to the old scene, creating the ghosting effect. (Source: the Society of Motion Picture Television Engineers; used with permission.)*

table tennis player ghosting into the sequence. Once the decoder receives

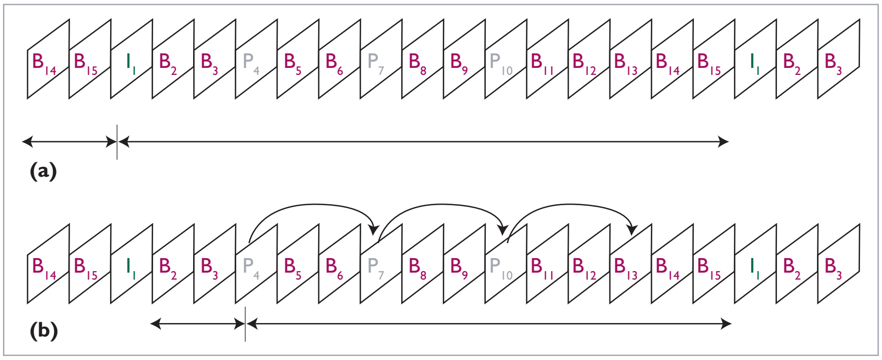


*Figure 3. Affected frames. We can see (a) the impact from I-frame loss on an open GoP and (b) the impact of a P-frame loss.*

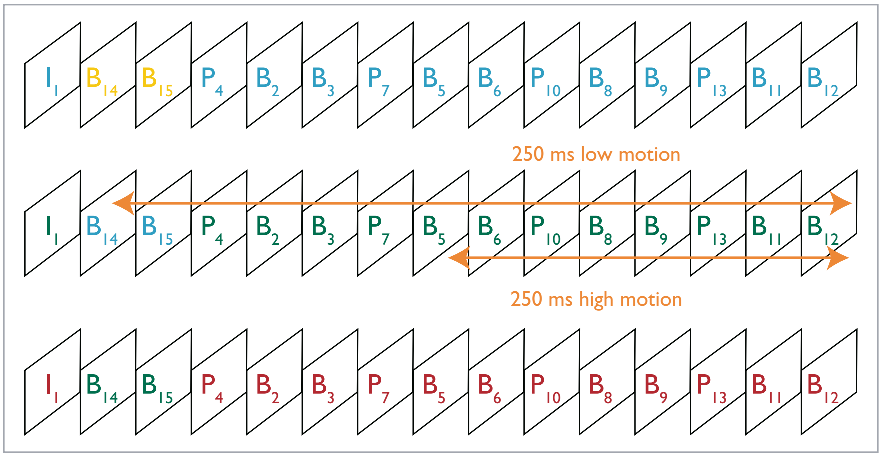


*Figure 4. GoP sequence. The frames in this sequence are impacted for a 250-ms best-case scenario.*

an unimpaired I-frame, the error will clear.

### Freeze Frame

Losing a continuous string of MPEG frames will result in a freeze-frame artifact; the last frame the decoder receives will usually be displayed until it receives information to replace the existing display.

## Video Impairment Analysis

For our video impairment analysis, we examined packet loss in I-, P-, and B-frames.

### I-Frame Impairment Analysis

An MPEG decoder uses the I-frame either directly or indirectly as the reference frame for all other frames within a GoP. Impairing the I-frame in any way will result in artifacts that continue through that GoP, and the video will recover only when the decoder receives an unimpaired I-frame. This restores video quality because the new I-frame isn't referencing any other frame. Dropping a single IP packet at the beginning of the I-frame and removing the frame header information will result in pixelization that will continue through the GoP. Losing this single IP packet has the same effect as removing the whole I-frame, which might span multiple IP packets. Packet loss within the I-frame that doesn't affect the frame header will result in slice errors that will continue through the GoP.

As we discussed in the last issue, frames B14 and B15 in an open 15:2 GoP will also reference the following I-frame; so, for such a GoP, a packet loss in the I-frame will affect 17 frames, as Figure 3a shows. At a frame rate of 29.97 fps, this equates to 567 ms of impact.

The perceivable impact on a viewer's video quality experience depends on the GoP length and motion within the clip. So, a trade-off exists in deciding on which GoP structure to use to support a streaming video service; larger GoP size provides a larger compression ratio, which enables more or higher-quality video content to be transmitted for a given encoded bit rate; however, the duration of visual artifacts due to IP packet loss might be longer, and user interactivity could be affected. If the GoP was 30:2, for example, then the artifact resulting from an I-frame loss could run for 32 frames (that is, 1,067 ms) in an open GoP structure.

The degree of pixelization that happens with an I-frame loss depends on how much motion or change exists in the clip when the loss occurs; the more motion within the clip, the larger the degree of pixelization.

### P-Frame Impairment Analysis

P-frames are coded using motion-compensated prediction from previous reference frames. The first P-frame takes reference from the I-frame, and subsequent P-frames take reference from their preceding P-frame. Impairments to the P-frame will extend through the remainder of the GoP, and the video will recover only when the decoder receives an unimpaired I-frame. As with an I-frame, therefore, a single IP packet loss at the P-frame header will affect the image as significantly as removing the entire P-frame, and pixelization will be visible through the rest of the GoP. Packet loss within the P-frame will result in slice errors that will continue through the GoP. In Figure 3b, we can see that frames B2 and B3 might reference the P4 frame, which is directly or indirectly providing reference to the remaining frames in the GoP; so, packet loss in frame P4 impairs frames B2 through B15. This highlights that the impact of a P-frame loss or P-frame impairment can be almost as significant as an I-frame loss; losing frame P4, for example, will result in 14 impaired frames, which equates to almost half a second at 29.97 fps.

The impact on the video quality the viewer experiences depends on which P-frame is impaired and how much motion or change exists in the clip. P-frames earlier in the GoP create a longer impairment because more subsequent frames are directly or indirectly impaired until the decoder receives the next I-frame. As with an I-frame loss, the higher the motion, the greater the pixelization that results from a P-frame loss.

### B-Frame Impairment Analysis

In MPEG-2 encoding, B-frames don't provide reference to any other frames, so losing part or all of a B-frame will affect only that particular frame. At a frame rate of 29.97 fps, this equates to a minor glitch that the viewer might see for roughly 1/30th of a second. Losing a B-frame can have more impact with MPEG-4 Part 10, if the encoder is using hierarchical GoPs, because, in that case, B-frames can also be used as reference frames.

## Video Quality Impact Testing and Analysis

We undertook testing to assess the impact of different packet loss durations. In practice, such loss might occur due to transient loss of connectivity following network link or node failures. We identified best- and worst-case loss scenarios for different periods of IP packet loss during MPEG stream transmission. A worst-case scenario would, for example, drop the first IP packet that contains the I-frame header information from the MPEG stream, thus causing impairments until the next GoP; this would impair 15 frames in a closed 15:2 GoP. We tested each scenario to verify the type of artifacts displayed and the impact on viewer QoE. We conducted our testing for periods of consecutive IP packet loss ranging from 10 ms to 500 ms for both the Susie low-motion and Football high-motion video streams, and for both SD and HD resolutions.

The first step in the test methodology was to create and record a reference MPEG-2 encoded video sequence using the Susie and Football clips. We used an Ineoquest Singulus digital video analyzer (see http://ftp.ineoquest.com/pub/docs/Datasheets/PDS_SING1T_001.pdf) to record the original clips at the IP layer, which we then edited to remove IP packets and create the defined packet loss in the video stream. We then streamed the original clips and impaired clips over an IP network to a set-top box that fed the uncompressed output to a Video Clarity ClearView Classic digital video analyzer (see www.videoclarity.com/PDF/ClearViewDataSheet.pdf). The analyzer performed a side-by-side (A/B), frame-by-frame analysis to compare the original unimpaired video with the impaired one.

We'll next examine the results.

### SD Testing Results and Analysis

Figure 4 shows the frames that were affected across a sequence of three GoPs (blue, green, and red) for the 250 ms best-case scenario for low- and high-motion clips; the frames within the GoPs are shown in transmission order.

Table 1 details the other scenarios for different packet loss durations, referencing the GoP sequence from Figure 4.

From the results, we can draw the following conclusions:

- The impact from 50 ms of IP packet loss can be as significant as that from 500 ms of packet loss when we consider the duration of visible impairment.
- The type of clip — high-motion versus low-motion — affects the number of frames impacted and, hence, the duration of visible impairment for the same loss period. We can see this in our analysis when we compare the 250 ms loss in high-motion and low-motion clips; six more frames are affected in the low-motion clip over the high-motion clip, resulting in a 200 ms difference in impairment duration.
- Impairment duration depends on which frames are impacted: an impacted I-frame results in a longer duration impairment than an impacted P-frame; an impacted first P-frame results in a longer impairment than the ones caused by an impairment in subsequent P-frames; and an impacted B-frame results in a shorter duration impairment than for other frame types.
- Freeze frames are the most common artifact seen due to the continuous loss of MPEG frames.

We also tested several HD scenarios.

### HD Testing Results and Analysis

We undertook HD testing and analysis using the GoP sequence Figure 4 shows. We encoded the 1080i HD video with a constant video bit rate of 18 Mbps including all service information (SI) and audio/video elementary streams using MPEG-2 Main Profile/High Level; we used a 15:2 GoP at 29.97 fps.

Table 2 details the scenarios for different durations of packet loss, referencing the GoP sequence from Figure 4.

From the testing conducted, we determined that the same conclusions we identified for SD also applied to HD. We saw no significant difference between SD impairments compared to HD ones for an equivalent packet loss duration.

| Content type | Scenario | Number of MPEG frames impacted | MPEG frames subject to IP packet loss | Number of displayed frames impacted | Displayed artifacts (number of frames impacted) | Total artifact duration |
|---|---|---|---|---|---|---|
| Low motion | 10 ms best | 1 | B12, (Green) | 1 | Minor glitch | 33 ms |
| High motion | 10 ms best | 1 | B12 (Green) | 1 | Minor glitch | 33 ms |
| Low motion | 10 ms worst | 2 | B12 (Green), I1 (Red) | 19 | Blocking/pixelization | 634 ms |
| High motion | 10 ms worst | 2 | B12 (Green), I1 (Red) | 19 | Blocking/pixelization | 634 ms |
| Low motion | 50 ms best | 3 | P13, B11, B12 (Green) | 5 | Freeze (2), slice (3) | 167 ms |
| High motion | 50 ms best | 2 | B11, B12 (Green) | 2 | Slice/freeze (2) | 67 ms |
| Low motion | 50 ms worst | 4 | P13, B11, B12, (Green) I1 (Red) | 20 | Freeze (3), blocking/ pixelization (17) | 667 ms |
| High motion | 50 ms worst | 3 | B11, B12, (Green) I1 (Red) | 20 | Freeze (3), extensive blocking/pixelization (17) | 667 ms |
| Low motion | 100 ms best | 6 | P10, B8, B9, P13, B1, B12 (Green) | 8 | Freeze (6), blocking/ pixelization (2) | 267 ms |
| High motion | 100 ms best | 3 | P13, B11, B12 (Green) | 5 | Freeze (2), slice/ pixelization (3) | 167 ms |
| Low motion | 100 ms worst | 7 | P10, B8, B9, P13, B1, B12 (Green), I1 (Red) | 23 | Freeze (6), blocking/ pixelization (17) | 767 ms |
| High motion | 100 ms worst | 4 | P13, B11, B12 (Green), I1 (Red) | 20 | Freeze (3), extensive blocking/pixelization (17) | 667 ms |
| Low motion | 200 ms best | 9 | B3, P7, B5, B6, P10, B8, B9, P13, B11, B12 (Green) | 13 | Freeze (11), minor blocking/pixelization (2) | 434 ms |
| High motion | 200 ms best | 6 | P10, B8, B9, P13, B1, B12 (Green) | 8 | Freeze/slice (6), blocking/pixelization (2) | 267 ms |
| Low motion | 200 ms worst | 10 | B3, P7, B5, B6, P10, B8, B9, P13, B11, B12 (Green), I1 (Red) | 28 | Freezes (13), blocking/ pixelization (15) | 934 ms |
| High motion | 200 ms worst | 7 | P10, B8, B9, P13, B1, B12 (Green), I1 (Red) | 23 | Freeze (6), extensive blocking/pixelization (17) | 767 ms |

*Table 1. Best- and worst-case packet loss scenarios for 10 to 500 ms loss periods (standard definition).*

## Further Conclusions

When looking at video impairments resulting from different durations of IP packet loss, we can draw the following conclusions from the testing results and analysis:

- The temporal compression MPEG encoding employs means that without loss-concealment techniques, even short durations of IP packet loss can result in significant visual impairment. Various factors determine the extent of the impairment, including the type of frames impacted, where in the frames the loss occurs, and how much motion or complexity the video has. This is different than voice-over-IP (VoIP) services, in which loss-concealment techniques might conceal up to around 50 ms of packet loss.
- Differences in the motion or complexity within a sequence affect the number of frames impacted for a given IP packet loss duration because frame sizes will vary. Less motion means more frames impacted and, hence, a greater duration of visible impairment for the same loss peri-

Table 1. Best- and worst-case packet loss scenarios for 10 to 500 ms loss periods (standard definition).

| Content type | Scenario | Number of MPEG frames impacted | MPEG frames subject to IP packet loss | Number of displayed frames impacted | Displayed artifacts (number of frames impacted) | Total artifact duration |
|---|---|---|---|---|---|---|
| Low motion | 250 ms best | 14 | B14, B15 (Blue), P4, B2, B3, P7, B5, B6…B12 (Green) | 17 | Freeze (5), freeze (10), blocking/pixelization(2) | 567 ms |
| High motion | 250 ms best | 8 | B5, B6, P10, B8, B9, P13, B11, B12 (Green) | 11 | Freeze (2), freeze (7), blocking/pixelization (2) | 367 ms |
| Low motion | 250 ms worst | 14 | B15(Blue), P4, B2, B3, P7, B5, B6…B12 (Green), I1 (Red) | 31 | Freeze(14), blocking/ pixelization (17) | 1034 ms |
| High motion | 250 ms worst | 9 | B5, B6, P10, B8, B9, P13, B11, B12 (Green), I1 (Red) | 26 | Slice/freeze (9), extensive blocking/ pixelization (17) | 868 ms |
| Low motion | 300 ms best | 15 | I1 (Green), B14, B15 (Blue),P4, B2, B3…B12 (Green) | 17 | Freeze (15), blocking/ pixelization (2) | 567 ms |
| High motion | 300 ms best | 9 | P7, B5, B6, P10, B8, B9, P13, B11, B12 (Green) | 11 | Freeze (2), freeze (7), blocking/pixelization (2) | 367 ms |
| Low motion | 300 ms worst | 16 | I1 (Green), B14, B15 (Blue), P4, B2, B3…B12 (Green), I1 (Red) | 32 | Freeze (6), freeze (10), blocking/pixelization (17) | 1067 ms |
| High motion | 300 ms worst | 10 | P7, B5, B6, P10, B8, B9, P13, B11, B12 (Green), I1 (Red) | 26 | Slice/freeze (9), extensive blocking/ pixelization (17) | 868 ms |
| Low motion | 400 ms best | 15 | I1 (Green), B14, B15(Blue), P4, B2, B3…B12 (Green) | 17 | Freeze (5), freeze (10), blocking/pixelization (2) | 567 ms |
| High motion | 400 ms best | 13 | B15(Blue), P4, B2, B3, P7, B5, B6, P10, B8, B9, P13, B11, B12 (Green) | 16 | Freeze (4), freeze (11), blocking/pixelization(2) | 533 ms |
| Low motion | 400 ms worst | 16 | I1 (Green), B14, B15(Blue), P4, B2, B3…B12 (Green),I1 (Red) | 32 | Freeze(6), freeze(10), blocking/pixelization (17) | 1067 ms |
| High motion | 400 ms worst | 12 | P4, B2, B3, P7, B5, B6, P10, B8, B9, P13, B11, B12 (Green), I1 (Red) | 30 | Freeze(3), freeze(11), blocking/pixelization(16) | 1001 ms |
| Low motion | 500 ms best | 16 | B12 (Blue), I1 (Green), B14, B15(Blue), P4, B2, B3…B12 (Green) | 19 | Freeze (17), blocking/ pixelization (2) | 634 ms |
| High motion | 500 ms best | 15 | I1, B14, B15, P4, B2, B3…B12 (Green) | 17 | Freezes (7), freeze (10) | 567 ms |
| Low motion | 500 ms worst | 17 | B12 (Blue), I1(Green), B14, B15 (Blue), P4, B2, B3…B12 (Green), I1 (Red) | 34 | Freeze(7), freeze (10), blocking/pixelization (17) | 1.134 s |
| High motion | 500 ms worst | 16 | I1, B14, B15, P4, B2, B3…B12 (Green), I1 (Red) | 33 | Freeze(6), freeze (11), blocking/pixelization (16) | 1.101 s |

od. Furthermore, videos with less motion have larger I-frames that constitute a significant proportion of the total bytes in the GoP; so, the probability of a packet loss affecting the I-frame is greater.

- The impairment resulting from 50 ms of packet loss can be as significant as that for 500 ms;

| Table 2. Best- and worst-case packet loss scenarios for 10 to 500 ms loss periods (high definition). | | | | | | |
|---|---|---|---|---|---|---|
| **Content type** | **Scenario** | **Number of MPEG frames impacted** | **MPEG frames subject to IP packet loss** | **Number of displayed frames impacted** | **Displayed artifacts (number of frames impacted)** | **Total artifact duration** |
| Low motion | 10 ms best | 1 | B12 (Green) | 1 | Minor glitch | 33 ms |
| High motion | 10 ms best | 1 | B12 (Green) | 1 | Minor glitch | 33 ms |
| Low motion | 10 ms worst | 2 | B12 (Green) I1 Red | 19 | Pixelization | 634 ms |
| High motion | 10 ms worst | 2 | B12 (Green) I1 Red | 19 | Pixelization | 634 ms |
| Low motion | 50 ms best | 3 | P13, B11, B12 (Green) | 5 | Freeze( 2), slice/ pixelization (3) | 167 ms |
| High motion | 50 ms best | 3 | P13, B11, B12 (Green) | 5 | Freeze (2), slice/ pixelization (3) | 167 ms |
| Low motion | 50 ms worst | 4 | P13, B11, B12(Green) | 20 | Freeze (3), 17 pixelization | 667 ms |
| High motion | 50 ms worst | 4 | P13, B11, B12 (Green) I1 Red | 20 | Freeze (3), 17 pixelization | 667 ms |
| Low motion | 100 ms best | 5 | B8, B9, P13, B11, B12 (Green) | 8 | Slice/freeze (2), freeze (4), pixelization (2) | 267 ms |
| High motion | 100 ms best | 4 | B9, P13, B11, B12 (Green) | 7 | Slice (1), freeze (4), pixelization (2) | 234 ms |
| Low motion | 100 ms worst | 6 | B8, B9, P13, B11, B12 (Green) I1 Red | 23 | Slice/freeze (6), pixelization (17) | 767 ms |
| High motion | 100 ms worst | 5 | B9, P13, B11, B12 (Green) I1 Red | 22 | Slice/freeze (5), pixelization (17) | 734 ms |
| Low motion | 200 ms best | 9 | P7 to B12 Green | 11 | Freeze (2), slice/freeze (7), pixelization (2) | 367 ms |
| High motion | 200 ms best | 8 | B5 to B12 Green | 11 | Slice/freeze (2), freeze (7), pixelization (2) | 367 ms |
| Low motion | 200 ms worst | 10 | P7 to B12 Green,I1 Red | 26 | Freeze (9), pixelization (17) | 868 ms |
| High motion | 200 ms worst | 9 | B5 to B12 Green I1 Red | 26 | Freeze (9), pixelization (17) | 868 ms |

the number of displayed frames subject to impairment can be the same depending on where in the encoded MPEG stream the packet loss occurred.

- The GoP length affects impairment duration: the longer the GoP, the longer the potential duration. For an $x$:$y$ GoP structure at $z$ fps, in the worst case $x + y$ frames will be impacted, with a resulting impairment duration of $(x + y)/z$ seconds.
- Receiving a nonimpaired I-frame restores video quality after any impairment due to packet loss.

The graph in Figure 5 plots the best- and worst-case impairment duration for a given packet loss duration; we can use this to estimate the possible range of impairment that we might expect for a given packet loss duration. The worst-case occurs when a loss affects an I-frame — because the I-frame represents a significant proportion of the total GoP size, the worst-case is statistically likely for any duration of loss. We suggest using this plot as a guide for engineering purposes.

Our initial investigation suggests that impairments resulting from an equivalent packet loss duration will likely be greater in MPEG-4 Part 10 deployments. As mentioned, MPEG-4 Part 10 can use B-frames as references in hierarchical GoPs to improve video compression efficiency, which

| Table 2. Best- and worst-case packet loss scenarios for 10 to 500 ms loss periods (high definition). | | | | | | |
|---|---|---|---|---|---|---|
| **Content type** | **Scenario** | **Number of MPEG frames impacted** | **MPEG frames subject to IP packet loss** | **Number of displayed frames impacted** | **Displayed artifacts (number of frames impacted)** | **Total artifact duration** |
| Low motion | 250 ms best | 12 | P4 to B12 Green | 14 | Freeze (2),slice/freeze (10), pixelization (2) | 467 ms |
| High motion | 250 ms best | 9 | P7 to B12 Green | 11 | Freeze (2),slice/freeze (7), pixelization (2) | 367 ms |
| Low motion | 250 ms worst | 13 | P4 to B12 Green, I1 Red | 29 | Freeze (14), pixelization (15) | 968 ms |
| High motion | 250 ms worst | 10 | P7 to B12 Green, I1 Red | 26 | Freeze (9), pixelization (17) | 868 ms |
| Low motion | 300 ms best | 14 | B14 B15 Blue, P4 Green to B12 Green | 17 | Freeze (2), freeze (13), pixelization (2) | 567 ms |
| High motion | 300 ms best | 12 | P4 to B12 Green | 14 | Freeze (2), slice/freeze (10), pixelization (2) | 467 ms |
| Low motion | 300 ms worst | 15 | B14 B15 Blue, P4 Green to B12 Green, I1 Red | 33 | Freeze (2), freeze (15), pixelization (16) | 1101 ms |
| High motion | 300 ms worst | 13 | P4 to B12 Green, I1 Red | 29 | Freeze (12), pixelization (17) | 968 ms |
| Low motion | 400 ms best | 15 | I1 Green, B14 B15 Blue, P4 Green to B12 Green | 17 | Freeze (2), slice/freeze (13), pixelization (2) | 567 ms |
| High motion | 400 ms best | 15 | I1 Green, B14 B15 Blue, P4 Green to B12 Green | 17 | Freeze (2), slice/freeze (13), pixelization (2) | 567 ms |
| Low motion | 400 ms worst | 16 | I1 Green, B14 B15 Blue, P4 Green to B12 Green, I1 Red | 33 | Freeze (2), freeze (2), freeze (16), pixelization (13) | 1101 ms |
| High motion | 400 ms worst | 16 | I1 Green,B14 B15 Blue, P4 Green to B12 Green,I1 Red | 33 | Freeze (2), freeze (2), freeze (16), pixelization (13) | 1101 ms |
| Low motion | 500 ms best | 17 | B11 Blue B12 Blue through to B12 Green | 20 | Slice/freeze (4), freeze (15), pixelization (1) | 667 ms |
| High motion | 500 ms best | 16 | B12 Blue through to B12 Green | 19 | Slice/freeze (3), freeze (15), pixelization (1) | 634 ms |
| Low motion | 500 ms worst | 18 | B11 Blue B12 Blue through to B12 Green, I1 Red | 35 | Freeze (4), slice/freeze (15), pixelization (16) | 1167 ms |
| High motion | 500 ms worst | 17 | B12 Blue through to B12 Green, I1 Red | 34 | Freeze (3), slice/freeze (15), pixelization (16) | 1134 ms |

means that impairments resulting from B-frame loss might potentially be more significant. Furthermore, the greater achievable compression ratio with MPEG-4 Part 10 means that service providers often prefer it in bandwidth-constrained networks. At the same time, these providers might also lengthen the GoPs to improve compression efficiency even further. If a GoP length extends to 30 frames, for example, an I-frame loss will result in a visible impairment for 30 frames — that is, more than a second.

Numerous technology approaches can minimize the packet loss duration experienced due to network failure events in IP and MPLS networks, such as IP routing protocol fast convergence, MPLS TE (Multiprotocol Label Switching Traffic Engineering) Fast Reroute, and Multicast-only-Fast-Reroute, but even with these tools, streams might still experience packet loss, which will affect perceived QoE. Service providers can further augment these network-level technologies using application-level approaches to recover from any loss experienced, thus improving viewers' QoE. Such approaches in-

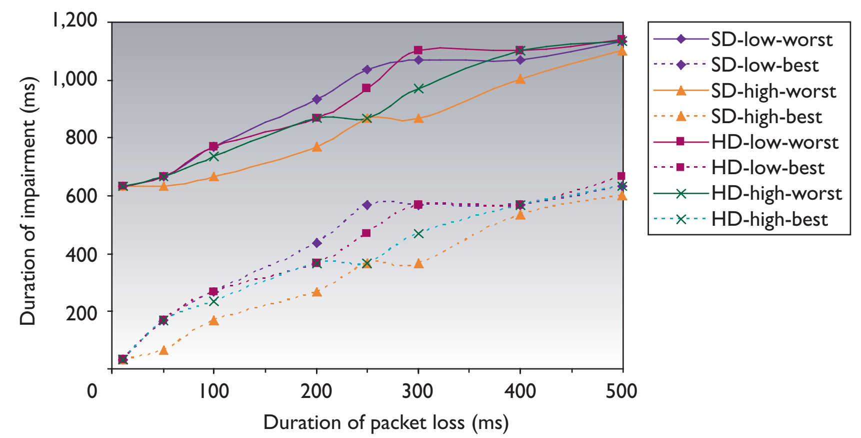


*Figure 5. Impairment duration vs. packet loss duration. We show how the impairment duration varies with the duration of packet loss.*

clude forward error correction (FEC), retransmission, temporal redundancy (that is, sending the same source data twice, separated in time), and spatial redundancy (using redundant video streams that use paths that are spatially diverse at the network level). In a future article, we hope to compare different network technology approaches for minimizing and recovering from video packet loss.

**Acknowledgments**

We thank our colleagues at Cisco for their valuable comments, which have helped in the development of this article.

**Reference**

1. J. Greengrass, J. Evans, and A.C. Begen, "Not All Packets are Equal, Part 1: Streaming Video Coding and SLA Requirements," *IEEE Internet Computing*, vol. 13, no. 1, 2009, pp. 70–75.

**Jason Greengrass** is a technical lead engineer in Cisco's Network Solutions Integration and Test Engineering Organization. He specializes in the testing and deployment of triple-play network solutions, with a focus on IP-based video quality testing and analysis. He's a Cisco Certified Internet Expert at Routing and Switching. Contact him at jasong@cisco.com.

**John Evans** is a distinguished consulting engineer within Cisco's Development Organization, where he works on the definition of network architectures for service providers. His technology focus spans IP routing and core IP/MPLS technologies, traffic engineering/management, quality of service, and video transport. Evans has an MSc in communications engineering from the University of Manchester Institute of Science and Technology. He authored *Deploying IP and MPLS QoS for Multiservice Networks* (Elsevier, 2007) and has filed numerous patents in his focus technology areas. He's a member of the IEEE. Contact him at john@cisco.com.

**Ali C. Begen** is a software engineer in the Video and Content Platforms Research and Advanced Development Group at Cisco, where he participates in video transport and distribution projects. His interests include networked entertainment, multimedia transport protocols, and content distribution. Begen has a PhD in electrical and computer engineering from the Georgia Institute of Technology. He's a member of the IEEE and the ACM and served as the local chair for IEEE ICSC 2008. Contact him at abegen@cisco.com.