Application Note – AN112

Understanding QMM
Objective Picture Quality Measurements

Tom Orlowski

VP Product Development, Pixelmetrix Corporation

Digital television broadcast systems based on MPEG can compress video information from hundreds of megabits to less than 5 Mbps with very little user perceived picture quality degradation. This is the case as long as everything goes well. However, given the complexity of today's digital television broadcast chain, problems can give rise to poor picture quality being delivered to the end user. This application note explores picture quality assessment in relation to the Pixelmetrix Quality Monitor Module (QMM) – a component of the DVStation Preventative Monitoring Platform.

Background

Digital television broadcast systems based on MPEG can compress video information from hundreds of megabits to less than 5 Mbps with very little user perceived picture quality degradation. This is the case as long as everything goes well. However, given the complexity of today's digital television broadcast chain, problems can give rise to poor picture quality being delivered to the end user.

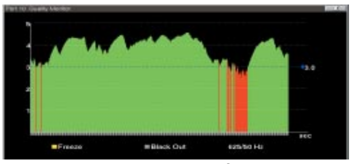

Objective picture quality measurement tools like the QMM can be used to monitor for these problems, set alarms on specific thresholds, and alert operators to investigate and isolate the fault in the network.

In addition to the alarm thresholds, one can configure the QMM to periodically log picture quality over a user-defined interval. For each measurement interval the minimum, maximum, and average picture quality points are recorded. Finally, with a click of the mouse, graphs of picture quality can be created for long term trend analysis.

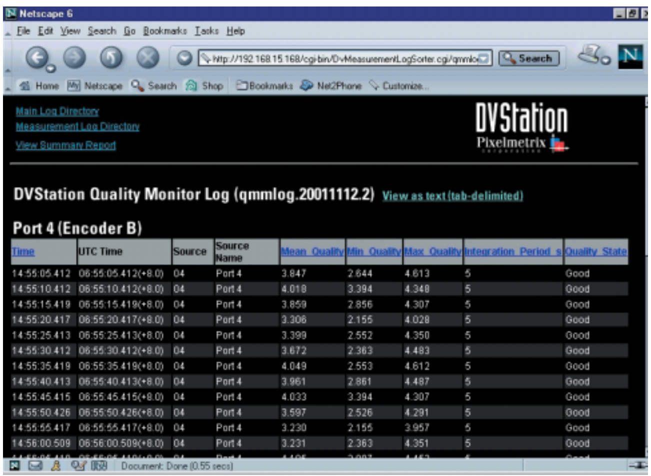

In all, picture quality assessment tools augment the other important components of a total broadcast quality management system – namely RF signal monitoring, MPEG-2 Transport Stream monitoring, and Broadcast Content Validation.

What is Poor Video Quality?

The MPEG-2 video codec (compression and decompression) model exploits spatial and temporal redundancy in video to reduce the amount of information required to describe moving pictures.

Spatial redundancy is the picture property that horizontally and vertically adjacent pixels have a very high probability of being of the same intensity (luminance) and/or color (chrominance). Temporal redundancy is the picture property that similarly located pixels of sequential pictures also have a very high probability of being of the same luminance and/or chrominance.

MPEG-2 uses the discrete cosine transforms (DCT) to modify the pixel spatial data representation to a frequency domain representation. The frequency domain data is quantized (less visible features are truncated and thrown away) and the remaining bits are compressed using highly efficient lossless compression. All pixel data is DCT converted based on fixed size blocks, 8 by 8 pixels in size.

Temporal redundancy is exploited though the use difference coding and motion vectors.

When things go wrong, the most typical visual artifact is tiling (blocking). This tiling effect comes from an imperfect restoration of the original image though the inverse DCT. Since information is removed after the DCT is performed, the artifact preserves the shape of the original data that went into the DCT, an 8 x 8 pixel block.

Tiling is not only the most common type of video defect in MPEG-2 systems, it is also the defect that is most noticeable and annoying to human viewers. As such, a system that can measure the severity of titling can be used as an effective objective picture quality measurement tool.

Copyright © 2003 Pixelmetrix Corporation

Looking for Tiles, Measuring Quality

The main task of the Pixelmetrix QMM (Quality Measurement Module) is to find tiles and measure their “visibility” within a video picture frame.

The QMM looks for tile boundaries by identifying “crank” features. Given four sequential pixels, a tile boundary exists between the two middle pixels if there is little variation between the two pairs of outside pixels and there is a large variation between the two middle pixels.

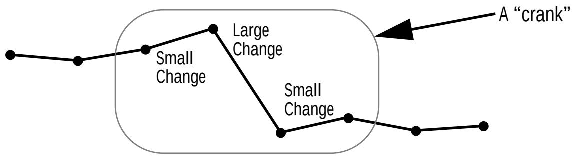

Consecutive Pixel Values

The probability that a crank pixel feature is attributed to an MPEG artifact is enforced if another crank appears exactly 8 (or integral multiple of 8) pixels in the horizontal or vertical direction.

The QMM scans each picture line and looks for cranks about every pixel transition. The “power” of each crank is the difference between the middle two pixels with respect to the difference of the outer pairs of pixels. As a line is scanned, the crank power values are accumulated into eight “buckets” in a round robin fashion. All of the lines are scanned and the results are accumulated into the same 8 buckets.

At the end of the scan the bucket values are sorted and the measure of the tiling effect is the difference/ratio between the lowest and highest bucket values.

This measurement is then linearized against a human visual model to give a 1 to 5 score based on the Single Stimulus Continuous Quality Scale.

5	Excellent	Imperceptible
4	Good	Perceptible but not annoying
3	Fair	Slightly annoying
2	Poor	Annoying
1	Bad	Very annoying

Resolution and Picture Scaling Challenges

The MPEG-2 video codec allows for various picture sampling resolutions at the input to the MPEG encoder.

At very low encoding bit rates, MPEG-2 struggles to maintain high picture quality and generates a significant amount of tiling. For low bit rates of less than 3 Mbps, it is often desirable to reduce the encoder input sampling resolution. This reduction effectively reduces the amount of information the encoder has to compress. It is popular to reduce the horizontal sampling resolution from 720 or 704 pixels down to 640, 544 or 480 pixels.

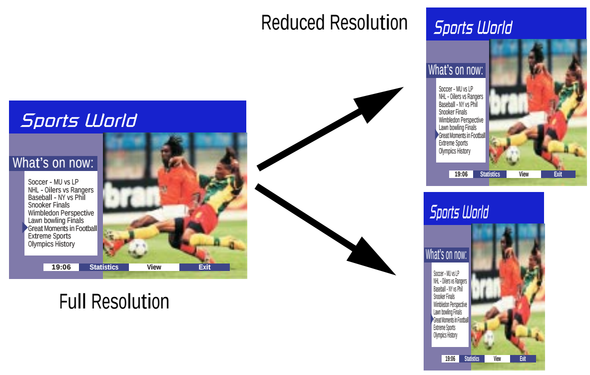

Once transmitted and decoded, the sub-resolution pictures are expanded (scaled) by the MPEG decoder up to proper proportions.

The resulting images are of lower resolution, but due to the decrease of information that needs to be coded, the tiling effects are significantly reduced. This results in overall better perceivable picture quality.

These images still have the potential to show 8 pixel frequency tiling artifacts, but these artifacts will no longer appear on the SDI signal at 8 pixel frequencies because of aspect ratio recovery scaling applied on the output of the MPEG decoder.

This presents a significant challenge to the measurement device looking for tiling effects. Methods to overcome these problems include running sophisticated “learning” algorithms on the SDI stream prior to running the actual tile detection algorithm. This “learning” algorithm tries to guess what the actual encoded resolution was. The pitfall of this approach is that the native MPEG codec resolution of an SDI signal can change without notice. As such, these learning algorithms must run periodically to ensure the tiling detection algorithm is running using the correct encoding resolution assumptions. When these resolutions change in the broadcast, there are interruptions in the tiling measurement accuracy.

Another way to determine these native resolutions is to use MPEG video elementary stream information. However, this requires that the measurement tool have access to the MPEG-2 compressed domain transport stream in addition to the baseband video stream.

The Pixelmetrix QMM uses a much simpler alternate approach to the scaling issue. Instead of measuring tiling effects in the horizontal direction, titling measurements are done in the vertical direction. MPEG-2 allows only 1x or 2x scaling options in the vertical direction. Even with 2x scaling enabled, the edges of the 16 pixel size tiles still fall exactly into the 8 pixel repetition measurement buckets. As such, the measurement algorithm is not affected by scaling.

Copyright © 2003 Pixelmetrix Corporation

Quality of I/P/B Frames Given Motion Vectors

As discussed above, the MPEG-2 video codec uses motion vectors to exploit the temporal redundancy of video sequences. Consecutive pictures that have similar features displaced by a few pixels can be coded by describing the motion of the picture element instead of redundantly re-encoding the picture element in a new location. Motion vectors allow sequences with camera pans to be encoded very efficiently.

In an MPEG-2 video sequence, P and B frames are allowed to carry motion vector information. These motion vectors can describe the current picture frame in terms of previous or future frames. I frames do not carry any motion vector information; they are self-contained descriptions of a complete frame.

The problem that these motion vectors create for tools that look for picture tiling effects is that these motion vectors are usually not of exactly 8 pixels in length. As an example, a camera pan that results in motion vectors of 3 pixels will move the tile edges by 3 pixels and put them in a bitmap location where some algorithms are not looking for edges.

On a frame by frame basis this creates a quality *beating* effect. Using an algorithm with fixed tile edge location expectations results in false reporting of much better quality for P and B frames, when actually these frames are usually slightly inferior to I frames. The P and B frames are reported as having less tiling because the algorithm is looking for the tile edges between the wrong pixel pairs.

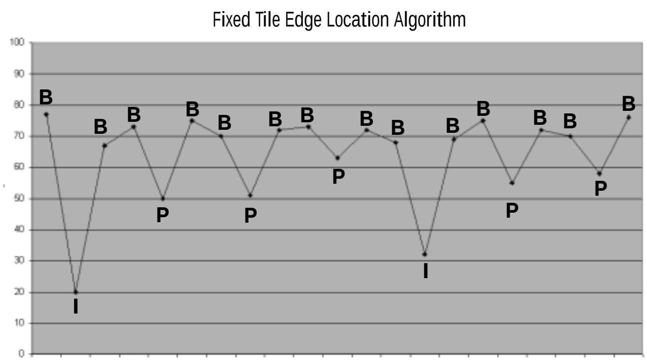

The Pixelmetrix QMM algorithm makes no assumptions about where these edges are, only that these edges (if truly attributed to a tiling effect) will occur at a repetitive rate of 8 pixels. As such it reports very consistent results for similar quality consecutive pictures whether they be I, P, or B coded. This results is measurements that more closely reflect the actual tiling intensity rather than the frequency of I frames, determined by the encoding GOP (group of pictures) length.

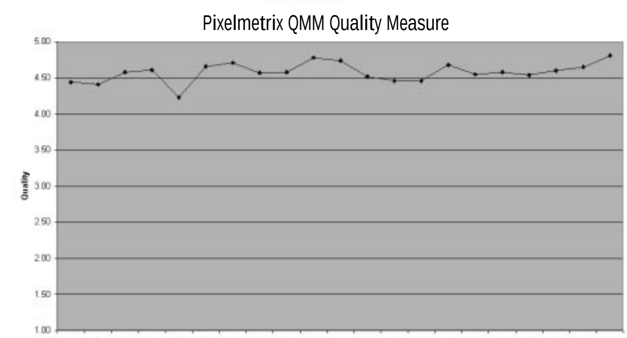

DVStation

Monitoring content, MPEG TS, and RF within an easy-to-use and integrated environment, the system has the highest port density in the industry. Ideal in environments with *many signals in one place* – such as satellite uplink centers, DTH operators, or cable head ends.

DVStation-Remote

Ideal for remote deployments with *a few signals in many places*, the system consists of a 1U control unit and up to four interface adaptors. Remote diagnostics can be conducted simultaneously from several locations, or alternatively staff can access telemetry directly by attaching a standard keyboard and CRT.

DVStation-IP

The Pixelmetrix DVStation-IP, a world first of its kind provides advanced video and content analysis and monitoring functionality for IP networks. The compact 1RU form features a multi-speed 10, 100, and 1000 Mb/s interface (gigabit ethernet).

DVStation-Pod

Featuring the same software and user interface of the DVStation, the DVStation-Pod product line consists of several book sized modules containing the interface circuitry. Each module connects to a laptop or desktop PC.

For More Information

To learn more about the DVStation, request a demo, or learn how Pixelmetrix might help you optimize video network integrity, contact us today!

On the Internet: sales@pixelmetrix.com
www.pixelmetrix.com

North America: 954-472-5445
Europe: +41-79742-7454
Asia Pacific: +65-547-4935

About the Author

Tom Orlowski is VP of Product Development with Pixelmetrix Corporation, manufacturer of the DVStation, a preventative monitoring solution for digital broadcast networks.

Copyright © 2003 Pixelmetrix Corporation