WHITEPAPER
IneoQuest Technologies, Inc.

Media Delivery Index

Monitoring and Troubleshooting Streaming Media Delivery

BY INEOQUEST TECHNOLOGIES

iNEOQUEST

WHTP-MDIT-1-0716

SUMMARY

MPEG video transport streams undergo time distortions known as jitter when being transported by packet switched networks such as Ethernet. Identifying and measuring jitter and packet loss in such networks is key to maintaining high quality video delivery. The Media Delivery Index (MDI) is a set of measurements used for monitoring and troubleshooting networks carrying any streaming media type. The MDI can be used to warn or alarm on impairments that result in unacceptable video delivery, and on conditions that result in unacceptable network margin, before video quality is impacted. MDI is defined and discussed and an overview of how MDI can be used to identify typical system issues is presented in this paper.

INTRODUCTION

MPEG video carried in Transport Streams via conventional Ethernet packet switched networks can arrive at its destination node with a time distortion – that is, the time between packets at the destination is different from that at the source. If packets are delayed by the network such as in Figure 1, some packets arrive in bursts with interpacket delays shorter than when they were transmitted, such as those in Group 1, while others are delayed, as in Group 2, such that they arrive with greater delay between packets than when they were transmitted from the source. This is defined as packet jitter or the time difference between when a packet actually arrives and the expected arrival time.

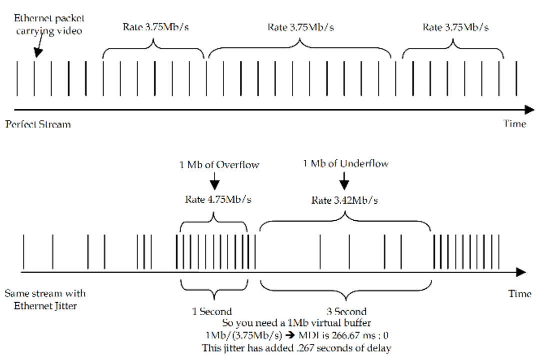

Figure 1

A receiver (decoder) displaying the video at its nominal rate must accommodate the varying input stream arrival times by buffering the data arriving early and assuring that there is enough already stored data when packets arrive late by prefetching data before beginning to display the video. Thus there is a design tradeoff made in all systems between having sufficient buffer memory to smooth all network delays, and encountering excessive delay due to a large buffer memory.

For a given size de-jitter buffer, then, there is a maximum amount of network jitter that can be accommodated without buffer overflow or underflow – either of which would result in a poor video display.

Similarly, the network switch/router infrastructure uses buffers at each node to accommodate instances where multiple input streams destined for a single output port arrive simultaneously. These buffers (queues) must also be sized appropriately to handle network congestion that might be due to the way traffic is routed through the infrastructure to differing link speeds in the infrastructure. If the switches implement methods for delivering Quality of Service (QoS) using packet metering algorithms, they may intentionally hold back packets to meet a QoS transmission specification further using buffer memory. Should such a buffer overflow, packets would likely be lost. Environmental electrical noise might create corrupted packets, also leading to packet loss. Even small packet loss rates result in a poor video display.

Packet delay variation and packet loss have been shown to be the key characteristics in determining whether a network can transport good quality video. These characteristics have been combined to produce the Media Delivery Index (MDI). The MDI can be used to assess how well a network can transport video, and to troubleshoot problems in networks whose performance has deteriorated due to reconfiguration or changing traffic loads.

MDI AND PERFORMANCE VERIFICATION

Viewing a video decoder's output picture is a basic way to assure that acceptable video is being delivered to a customer's location, but is lacking in several important capabilities:

Audit Records: If a video impairment occurs, it may be missed since viewing a picture does not occur 100% of the time. There is no automated way to record the event with an accurate timestamp to correlate its occurrence with other network events.

Subjective Result: The severity of a fault is subject to the viewer's evaluation. Verification results may vary with different viewers.

Timliness: If an impairment is noted, by definition, it's too late to prevent delivery of the fault.

Troubleshooting Help: If an impairment is noted, there are few clues about what went wrong or how to repair it.

Margin Indication: If the video is not impaired, there is no indication of the system's operating margin or the probability of a future impairment.

A better alternative to assure acceptable video delivery is to measure the known, critical network parameters associated with streaming content, namely, packet delivery delay variation and packet loss as described previously. It has been shown that any packet loss causes a video impairment. It is straightforward to detect lost packets in a network infrastructure by monitoring the links between network devices, but it is important to note that even if the infrastructure does not drop packets, it may affect delivery times to a video decoder such that the decoder buffer overflows or underflows resulting in packet loss within the decoder. The MDI accounts for these factors.

DF : MLR

The DF is computed at the arrival time of each packet at the point of measurement and displayed and/or recorded at time intervals. A time interval is typically about 1 second. The DF is updated and displayed at the end of an interval. Given a virtual buffer level X where DF is being measured,

$$X = |\text{Bytes Received} - \text{Bytes Drained}|$$

Then,

$$DF = [\max(X) - \min(X)] / \text{media rate}$$

Where media rate is expressed in Bytes/second and max(X) and min(X) are the maximum and minimum values measured in an interval.

Furthermore, the largest difference is captured and displayed for all intervals in a measurement period. The measurement period can range from a few seconds for monitoring a specific network activity, to indefinitely for a long-term monitor.

The MLR is computed by subtracting the number of media packets received during an interval from the number of media packets expected during that interval and scaling the value to 1 second:

$$MLR = (\text{Packets Expected} - \text{Packets Received}) / \text{Interval Time in Seconds}$$

MLR is the number of media packets lost per second.

DF DISCUSSION

(Bytes Received – Bytes Drained) represents the size of what a virtual buffer would need to hold at the measurement point and time. When this value is divided by the media rate, the result indicates the time it would take to drain the virtual buffer. A faster media rate takes less time to empty the buffer. By taking the largest span (max – min) of the value encountered during a measurement period, the buffer time needed to avoid packet loss is the DF. By displaying the largest span encountered during an interval, the worst case value can be seen. This value can be logged with a time stamp and correlated with other network events to gain insight into the magnitude and possible cause of network upsets. By characterizing the network in advance, the magnitude of the upsets necessary to cause video corruption can be set for various locations within a network infrastructure and, therefore, indicate how much headroom is required to avoid packet loss and provide margin can be determined.

To effectively use a buffer to smooth packets arriving earlier or later than nominal, the buffer must be prefilled before starting the video play (the drain) so that a transient gap in data does not permit the buffer to empty, thus disrupting the video play. Since the DF represents the time it takes to empty the virtual buffer at the video rate, the DF indicates the latency contribution from the network to the total latency in delivering the video from source to destination. Excessive network jitter requiring large buffers to smooth the video may result in unacceptable end to end delay.

If the DF value grows continuously from interval to interval, it indicates that the long term, sustained drain rate is different than the fill rate (either faster or slower.) This is most likely caused by a video source issue, as network upsets are typically transient in nature.

MLR DISCUSSION

Capturing, displaying, and logging the maximum MLR value that has occurred over a measurement period indicates the largest magnitude of upset in an interval over the measurement period. Loss due to buffer overflows that are the result of network congestion or network device misconfiguration will typically be periodic and will result in many intervals with a positive but similar magnitude MLR. Transient effects due, perhaps, to local electrical noise interference may result in only a small number of intervals affected. A constant MLR with varying but low traffic loads (no network congestion) may point to server or other media source issues.

MDI USE

A general approach to using MDI in installing, modifying or evaluating a video network is:

1. First identify, locate and address any packet loss issues using MLR.
2. Identify, locate, measure and address jitter margins using DF.
3. Establish an infrastructure monitor for both DF and MLR for transient upsets exceeding preset thresholds representing network or server changes, misconfigurations, failures, etc. The thresholds may be determined based on the network equipment type, decoder type, margin and quality desired, etc.

OTHER MEASUREMENTS

To troubleshoot concerns identified with the MDI and to aid in system configuration and monitoring, other measurement parameters may also be desired.

Network Utilization: Tracking the instantaneous, minimum and maximum overall network utilization is needed to verify that sufficient raw bandwidth is available for a stream on a network. High utilization level is also an indicator that localized congestion is likely due to queue behavior in network components. The DF provides a measure of the results of congestion on a given stream.

Instantaneous Flow Rate (IFR) and Instantaneous Flow Rate Deviation (IFRD): The measured IFR and IFRD confirm a stream's nominal rate and, if not constant over time, gives insight into how a stream is being corrupted.

Average Rate in Mb/s: This measure indicates whether the stream's rate being analyzed conforms to its specified rate over a measurement time. This is the longer-term measurement of IFR.

Stream Utilization in Percent of Network Bandwidth: This measure indicates how much of the available network bandwidth is being consumed by the stream being analyzed.

CONCLUSIONS AND OBSERVATIONS

The Media Delivery Index (MDI) can be used to monitor both the quality of a delivered video stream as well as to show system margin by providing an accurate measurement of jitter and delay.

The MDI addresses a need to assure the health and quality of systems that are delivering ever higher numbers of streams simultaneously by providing a predictable, repeatable measurement rather than relying on subjective human observations. Use of the MDI further provides a network margin indication that warns system operators of impending operational issues with enough advance notice to allow corrective action before observed video is impaired.

IneoQuest Technologies, Inc.

170 Forbes Boulevard – Mansfield, MA 02048

Toll Free: +1 866 464 4636

Fax: 508 339 4727

IneoQuest UK Office (44) (0) 1865 784322

www.ineoquest.com

WHTP-MDIT-1-0716