IP QoS Objectives For Broadcast Services

Bengt J. Olsson
Net Insight AB
Stockholm, Sweden

Abstract - *The rapidly progressing transition from using traditional TDM-based transport mechanisms, such as SDH/SONET, to use IP based transport for broadcast services provides broadcast operators with a possibility to use a converged network platform for all their services and hence save on infrastructure costs. However, in the trail of this transition, there is a renewed focus on transport quality of service (QoS) issues and related to this, a focus on how to follow up service level agreements (SLAs) given for IP connectivity QoS. This article will discuss these questions with respect to the broadcasters' needs, while an emphasis will be given to the effects and handling of packet delay variation (PDV) since this is the least known QoS parameter and a good estimator for the general connection performance.*

INTRODUCTION, AND A LOOK IN THE MIRROR

In the days of TDM-based transport, such as SDH/SONET/ PDH and alike, transport QoS was rarely discussed and SLAs were more or less written to only specify the availability of a service in terms of events that could interrupt the service, such as fiber breaks, power outages and so on. TDM transport properties not only fitted the broadcast services well, but also provided a guaranteed QoS that was easy to interpret: either the service worked 100%, or the service was completely unavailable.

With the introduction of a new service landscape involving a much higher content of data networking, both for society as a whole and also for broadcasters that now rely on concepts such as non-linear production, use of intranets etc., TDM-based networking no longer provides the optimal networking platform. Consequently, service providers have migrated to IP-based transport technologies. While IP-based transport provides numerous benefits in the general perspective, the transition is not without its challenges for all services. Typically, IP networks were designed to effectively accommodate for *elastic* services by means of using *statistical multiplexing* to provide a high resource (link) utilization, which is at the expense of a varying delay in router buffers. By their nature, much of the services of a broadcaster are *non-elastic* to their nature, such as video and audio services. In TDM systems they were served by *deterministic multiplexing*, i.e. they were served by a channel that provided the resources that the service needed. A number of techniques are being introduced in IP networks to better handle non-elastic traffic, however it is still difficult to provide the TDM-like, guaranteed QoS that was found when using SDH.

As an example, compare the timing performance for an IP connection vs. an SDH circuit. An SDH circuit that is properly synchronized will have a negligible jitter and wander for all broadcast services. On the other hand, if an SDH network is not properly synchronized the circuits may be subject to a wander that is due to what is called "pointer justification wander." The amplitude of this wander is approximately 0.15 µs. Even this minute pointer adjustment jitter or wander has historically been known to cause problems in video transport by introducing artifacts in color hues. Now, compare this jitter to the PDV of typical IP connections. For reasons that will be clearer later, a PDV around 150 µs could be a typical value for a reasonably well managed connection. This is three orders of magnitudes higher jitter than in the SDH case! The pointer justification jitter is hardly visible in Figure 1 as compared to the IP connection jitter.

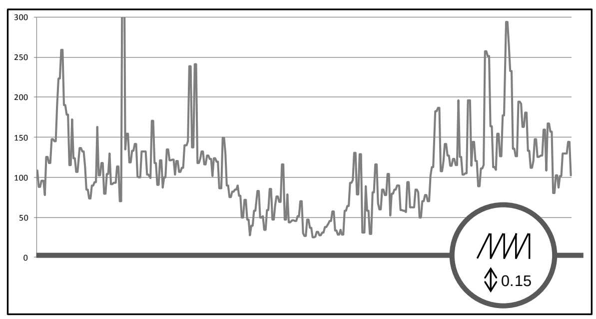

Figure 1. Network jitter (in µs) on an IP connection compared to the pointer justification jitter of an SDH connection (expanded)

But the services sent over the IP or SDH networks have the same requirements, seen from the user of these services. Video-over-IP adapters generally mitigate the high degree of packet jitter with what is called "jitter buffers." But do these suffice to provide studio quality of the video with respect to timing when subject to the network jitter? What packet jitter levels are compatible with high-quality video? What type of PDV should be expected from typical IP connections? These are questions that will be discussed in this paper.

IP QOS PARAMETERS

When discussing IP QoS there are typically three parameters that are discussed:

- Packet Loss Ratio (PLR)
- Packet Delay (PD)
- Packet Delay Variation (PDV)

These parameters directly affect the QoS of a service in different ways. Of these three parameters, PLR and PD are most well known and their influences on services are relatively easy to understand. However, the reasons for, and effect of, PDV is less known. Therefore, the emphasis in this paper will be to describe PDV issues. To note, there are also other parameters that have less influence on the quality or only affect the availability of the services, but these will not be discussed in this paper.

Packet Loss Ratio (PLR):

It is quite easy to understand the effects of packet loss: user data is lost and this in turn affects the video or audio as artifacts, macro blocking, audible clicks, etc. The PLR of an IP connection is probably the most important quality parameter for broadcast services and can be specified to meet acceptable data loss criteria. One such criterion proposed in the ITU-T recommendation Y.1541 [1] is that a video service should not be subject to more than "one hit per day." Assuming random packet loss, a service bandwidth in the order of 100 Mbps and the healing effects of using Forward Error Correction (FEC), it can be shown that the "one hit per day" criteria, that is one uncorrectable error per day, is met by the 10^{-5} PLR that is specified for QoS classes 6 & 7 in Table 3 of [1]. A higher PLR will give "hits" more often and vice versa.

It is possible to use more effective packet loss reducing techniques, such as hitless 1+1 merge, where the same service is sent over two geographically diverse paths and then merged at the destination. Packets are enumerated and only if the same numbered packets from both paths are lost a service packet loss occurs. This mechanism can be compared to FEC. Having two paths for the same data gives a 100% bandwidth overhead, vs. FEC that typically has a 5-25% bandwidth overhead. On the other hand hitless 1+1 merge offers stronger loss recovery, and also offers protection against packet loss bursts or plain path faults. It should be noted though that hitless 1+1 merge only have the packet loss reduction property when both legs of the 1+1 connection is working. The FEC mechanism usually gives larger delays than hitless 1+1 merge, unless the latter is used in for example large ring configurations where the differential delay can be substantial.

Even stronger loss recovery can be acquired using retransmission techniques. For general data, TCP is usually used. But for the streaming services that broadcasters use, UDP based algorithms are typically deployed in order to avoid the congestion-control algorithms used in TCP. The PLR reduction gain is exponential with respect to the number of round-trip delays that the service is subject to. Hence, at the cost of many round-trip times in delay, a very strong PLR suppression can be obtained. For unmanaged networks, or for transmission over the Internet, this may be a very usable technique. For managed networks typically the other two mechanisms to mitigate packet loss, FEC or hitless 1+1 merge, would be used.

Loss recovery mechanism	Typical enhance-ment	Overhead	Typical delays
FEC	10^{-5} -> 10^{-9}	5-25%	FEC matrix (10-100 ms)
Hitless 1+1 merge	10^{-5} ->10^{-10}	100%	Differential delay (1-10 ms)
Re-transmission	10^{-2} -> 10^{-6} (N = 3)	~PLR (low)	N x RTD (> 100 ms)

Table 1. Loss recovery mechanisms with example characteristics

Table 1. provides example characteristics for the described packet loss reducing techniques. A resulting PLR of about 10^{-10} is needed to fulfill the "one hit per day" criteria for an uncompressed HD-SDI service. As seen in the table this would be reachable by hitless 1+1 merge and nearly reachable with FEC at a real PLR of 10^{-5}. An interesting prospect would be to combine FEC with hitless 1+1 merge such that FEC operates independently at each leg of the hitless 1+1 merge connection, and recover packet losses to a certain degree before the hitless merge, which in its turn recover the residual packet loss. In a simple model of the recovery capabilities of FEC and hitless 1+1 merge (for random losses) would be that if the PLR can be expressed as 10^{-N}, the PLR after FEC would be $10^{-(2N-1)}$. For hitless 1+1 merge the resulting PLR would be 10^{-2N}. Hence the combined effect after FEC and hitless 1+1 merge would be $10^{-(4N-2)}$. That is, combining FEC and hitless 1+1 merge would make it possible to reach a PLR of 10^{-10} with an input PLR of 10^{-3}. This PLR is in line with may typical service provider SLAs, while a PLR of 10^{-5} is not.

Most packet loss calculations assumes that the packet loss is random, since otherwise the models becomes very complex. In reality however, packet losses mostly occurs in bursts. A parameter that is pertinent for packet loss issues is the "packet burst loss size" (PBLS). This parameter is not standardized but it has a large impact on packet loss behavior. FEC recovery calculations demand that packet loss distribution is more or less random. This is because a packet loss burst, where consecutive packets are lost, may not be possible to recover. In order to recover a packet loss burst, the size of the burst must be smaller than, or equal to, the column width of a FEC matrix. (This is true for 1-dimensional FEC, however, 2-dimensional FEC has somewhat better burst tolerance). To see the impact on PBLS on the "hit" rate of a FEC protected connection, consider the following example: assume a 100 Mbps stream of 1500 bytes packets. A distribution could be constructed with the following properties: approximately each 250 seconds a burst of 21 packets are lost. Since the maximum FEC column width in SMPTE 2022-1 [2] is 20, this burst loss would not be possible to recover. This means that while this distribution has a PLR that is still better than 10^{-5}, it

would lead to a hit each 250 seconds, which is many orders of magnitude worse than the objectives in [1].
Therefore, it would be desirable to introduce a metric for PBLS, as well as the distance between burst losses in the standards for IP QoS objectives, and for service providers to monitor and manage their connections with respect to PBLS.

Packet Delay (PD):

PD does not affect the QoS of the service in the same sense as PLR or PDV. It introduces a static delay in the information transfer that may harm the end user service in different ways. Hence different applications will have different requirements on the end-to-end PD. [1] specifies 100 and 400 ms for QoS classes 6 and 7 respectively. However, this should be considered as rather arbitrary recommendations, instead the service context should decide the correct PD requirement.

However, it is interesting to note that there is a relation from PD to both PLR and PDV of a connection. This is because, that beside the pure physical transmission delay in the medium (fiber, air for radio links etc.), the two biggest factors that decides the PD are:

- Size of the FEC matrix
- Size of the jitter buffer

Thus, by controlling the PLR and PDV of a connection, for example by applying traffic shaping in the routers along the connection, it is possible to decrease the sizes of the FEC matrices and jitter buffers which in its turn decreases the PD of the same connection.

Packet Delay Variation (PDV):

PDV is shortly defined as the variation in arrival time for the packets of a stream. There are several associated parameters to characterize this delay variation, which will be discussed in this paper later. As PDV is a less known quality parameter that affects the timing integrity of the transported service, the rest of this paper will discuss timing and synchronization issues and their relation to PDV.

TIMING AND SYNCHRONIZATION FOR BROADCAST SERVICES

Depending on the service, the importance of PDV effects can vary. By using play-out buffer techniques it has been possible to recover time to a level suitable for simpler video and audio applications. However, in order to preserve the timing properties of high-end services for contribution or production, often referred to as "studio quality," or in order to provide explicit time and synchronization services (e.g. IEEE1588, TToIP, G.703 2.048 MHz sync, 10 MHz sync etc.), more sophisticated time recovery mechanisms as well as tighter control of the network induced jitter and wander are needed.

Simple jitter buffer-based timing recovery is used in a variety of decoders, often with proprietary variations to cope with various aspects of the play-out quality. For example, these variations may involve non-correct recovery of the PCR clock in a transport stream in order to be able to change channels more quickly and use error concealment techniques to hide the manipulations in the play-out buffer.

While this may be a sufficient and even a good strategy for consumer types of products, a production facility should not be designed by relaxing timing requirements and then use error concealment strategies to cover the problem in the transport. This is a bad network design strategy. A simple but pertinent example will show that. Consider time base correctors, or "frame stores", which are used to adapt the frequency of incoming video frames to the frame frequency governed by the local studio clock. Such frame stores will produce a frame drop or a frame repeat every now and then. Depending on the broadcaster’s policy, this may or may not be suitable, and might seem quite harmless. But now, consider that the broadcaster wants to contribute 4K produced material by transporting 4 independent 3G-SDI signals that should be combined at the receive end. Then the drop or duplication of frames in the frame store will cause big problems. All four signals will probably drop frames at the same rate but with arbitrary offsets versus each other's. This will be seen in the resulting 4K image as a cycling offset in time of the quadrants of the picture, which is very difficult to correct or conceal.

A better strategy is then to have designed the network or connection for proper timing and synchronization transport. The possibilities to design this and the extent to which PDV will become a problem will be discussed below.

"INBAND" OR "OUTBAND" SYNCHRONIZATION

There are two different options available to convey synchronization between the ingress of a network and the egress. The first option is to let the network itself carry the timing information, together with the data as an intrinsic embedded clock or using some other means to transport synchronization from ingress to egress within the network.

The other option is to use a "common clock" architecture where both ingress and egress are synchronized to a common clock source, outside the network, which in most cases would be a GPS clock.

The advantage of the second option is that all PDV issues with respect to timing and synchronization simply disappear! Jitter buffers must still be used since network jitter must still be absorbed, but the jitter buffer will not be used to recover the timing.

While the second option looks very appealing with respect to the timing characteristics it provides, it is seldom a practical solution for broadcasters. It may be difficult to apply a GPS clock to all ingress and egress points of the network for several reasons: cost, right-of-way for cabling to antennas, maintenance of extra equipment, etc. GPS is not infallible either and many factors may affect its performance, such as bad weather and intentional or un-intentional jamming or spoofing of the GPS signal.

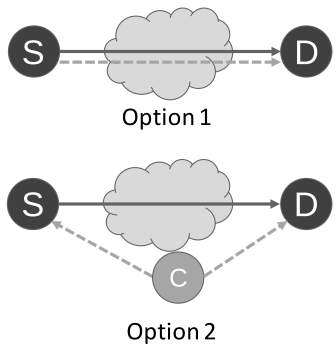

Figure 2. Two basic options for network synchronization.

Another caveat with option two is that it is not only necessary that the ingress and egress transport nodes are synchronized, but the ingress signal of an IP video adapter must in general be synchronized with the common clock as well. For example, in order to clock out an SDI frame from an IP video adapter using a 10 MHz GPS reference requires that the actual camera that produced the ingress signal be also synchronized to the GPS, otherwise the egress SDI signal will slip frames.

PDV ORIGINS

PDV is inherently due to the asynchronous nature of packet transport. This in turn is manifested as a number of effects:

- Statistical multiplexing in routers and switches (queuing)
- "Head of line" blocking in output queues
- Variability in router table lookup times for each packet

These are the effects attributed to PDV generation in [1]. In this model, in-elastic traffic (real-time, video, audio, etc.) is sent in a priority class with strict priority over the default forwarding class, hence the second effect, "head of line" blocking. This means that each high-priority packet that encounters a low-priority packet that is under transmission through the interface has to wait for it to complete the transmission before the high priority packet transmission can commence. While queuing effects in general dominate the PDV, head of line blocking can contribute significantly to the total PDV over a multi-hop connection.

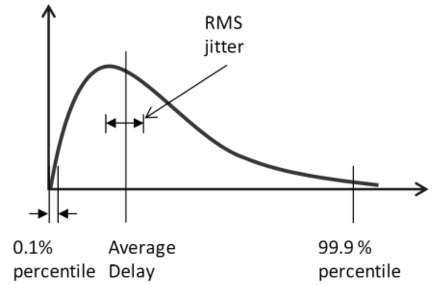

Figure 3. Example of PDV distribution

The 3rd effect however can be considered as more or less negligible today. In [1] this variability term is chosen to 3 µs per router, which is a very defensive value looking at today's silicon based hardware look-ups that are much faster. It must be noted though that the very general scope of [1] forces its authors to consider a broader range of equipment and network/link types than would be used by broadcasters. Business access routers connected by E1 links will have very different PDV properties than 1/10 GbE switches used by broadcasters.

Beside these three effects that are more or less purely statistical, there are other, more systematic effects that may contribute to PDV:

- "Beating" patterns in service traffic
- Varying network load
- Re-routes in the network

These effects produce more slowly varying PDV’s, which in many respects is more difficult to handle than what is caused by the faster statistical PDV. For example, a beating pattern may be produced by having a number of not synchronized media streams traversing the same link. At most times their packets will be spread in time, but inevitably their arrival will coincide in time to produce longer packet trains that affect the end-to-end delay for packets belonging to certain streams and hence their PDV. In [3] the author shows that just three similar constant bit rate (CBR) video streams over a 100 Mbps Ethernet connection produces a beating pattern, which makes it very hard to recover the videos to studio quality in terms of timing. The more video streams, the larger the effect of the beating between streams will be. A way to avoid this beating problem is to use a synchronous scheduling technique over IP, where video streams are synchronously multiplexed onto a common single IP bearer and optionally switched synchronously within the network to reach different end destinations [12].

Varying network loads have the effect to shift the center of gravity of the PDV distribution. This is very problematic for time recovery circuitry that is based on jitter buffer fill levels, since these low-pass filter the fill level, which essentially means to track the center of gravity of the PDV distribution. When the mean of the PDV distribution moves, this is manifested as a wander in the time recovery. This effect, if not considered, will kill any attempt to acquire accurate synchronization, and especially for phase synchronization, to recover absolute time accurately. As will be covered later, a much better strategy is to pre-select the minimum latency packets (the packets within the 0.1 percentile in Figure 3 to the left) since the distribution of these packets are much less sensitive to load variations.

Re-routes create such large PDV changes and must be handled by additional means to conceal the resulting phase jumps from reaching the time recovery circuitry in the first place, if accurate synchronization is required.

BROADCAST SERVICES TIMING SPECIFICATIONS

Timing specifications for video and audio signals, as well as for pure synchronization or time transport signals, are needed in order to be able to recreate the original signal at the egress to be within a desired quality. What this means will, as discussed before, depend on the usage context and it is not always easy to know exactly what is needed. However, the most well specified requirements are those termed as "studio quality" and are also those that should be regarded as target specifications for a professional media network. Figure 4 below depicts these requirements in the time domain for the most common broadcast services and will be discussed in detail below.

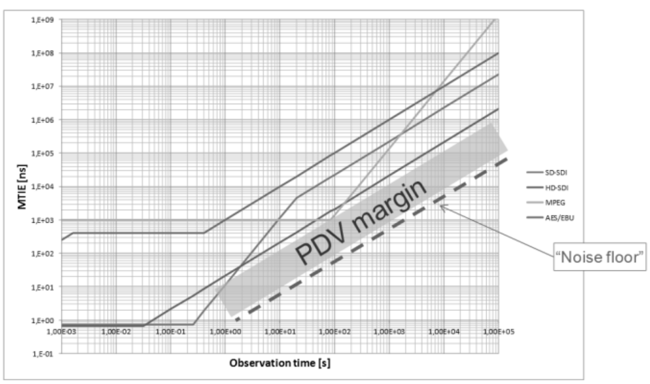

Figure 4. Time domain jitter and wander specifications, for video/audio services, compatible with studio quality requirements. (Data from [4])

The general appearance of these specifications is that to the left, depicting very short time periods, the function is a constant, which represents a fast jitter with high frequency. When it comes to time recovery, this region is fairly easy to handle since time recovery circuits, that in almost all cases use some form of PLL (phase locked loop) mechanism, can easily filter out this high-frequency "noise".

On the other end of the spectrum, towards very long time periods, it can be seen that more time deviation is allowed. This region does in general not pose a practical problem either for time recovery, since in a limiting sense this case would effectively be equivalent to run the complete system at a somewhat different frequency, and in some respect the choice of reference frequency is arbitrary and does not affect the quality of media transport. There is a very important exception from this statement and it relates to phase or time synchronization where there is a limit on the absolute time deviation and hence the frequency and phase must be absolutely tracked, and no long time deviation is allowed. This also makes absolute time synchronization the most difficult task to make work satisfactorily over a packet network.

Seen in the Figure 4 is also a dashed line that schematically represents what could be called a "noise floor" of the time recovery circuitry, which depends on physical factors such as the quality of the used oscillators, ambient temperature variations, etc. In between this noise floor and the specification curves there is what could be called an "allowance" for PDV. What also can be seen is that the most sensitive region is in the mid section, with time periods reaching from parts of seconds up to hundred of seconds. It is here where it is the largest risk to violate the studio quality standards. This is because the allowance for timing jitter is small and the jitter or wander frequencies in this region are low enough to make it hard to suppress by low-pass filtering techniques.

TIME RECOVERY TECHNIQUES

A general model of a time recovery circuit is presented in the figure below (from [5]). It consists of a few blocks that will be discussed below. From the left there is the "packet timing clock" or "master" clock that is to be recovered on the right side. The clock data ("tick") is encapsulated into a packet stream. The network will add noise to this clock stream in the form of PDV. The packet stream arrives at the destination and is processed.

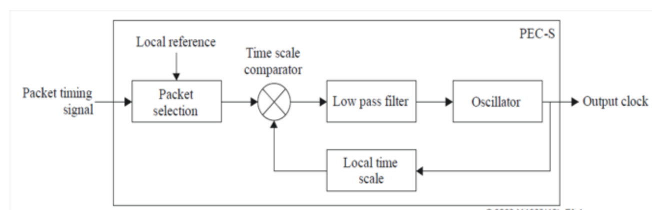

Figure 5. Functional model of a packet-based equipment clock (From [5])

The first step in the recovery process, and one that is not always implemented, typically in jitter buffer fill level-based implementations, is to pre-select the packets that are most relevant for clock recovery. This is a very important step since we have an a priori knowledge that some packets are more relevant than others for this purpose. These are the packets *with lowest delays*. If from a sample of timing packets, a "block", the packet with the lowest delay is selected and all others are discarded, the selection will have lower variability than if all timing packets were used. (See Figure 6 below). A way to see this is to imagine that of all packets in the block, there might be one or more packets that traverse the network without being delayed by either queues or blocking lower priority packets. In principle these packets would have the same delay through the network with more or less zero variability. On the other hand you need to have a sufficient amount of timing packets for the time recovery circuit to discipline its clock, so there is a trade-off. However, if this technique is not used it is more or less impossible to exclude wander depending on, for example, load variation in the network.

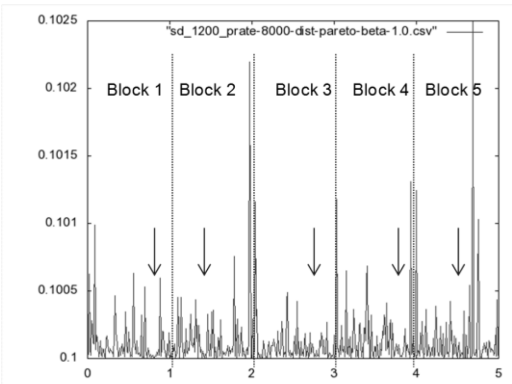

Figure 6. Pre-selection of timing packets with lowest delay

The second step is the time recovery circuit itself. As described before, this is essentially a digital phase lock loop that low-pass filters the timing information in order to achieve a stable output clock. The actual implementation of the digital PLL can vary and the performance is subject to the craftsmanship of the designers. But since the general objective is to low-pass filter the timing information, its function can be discussed in general terms.

A low-pass filter is characterized by a frequency region that lets the signal through unchanged and a frequency region where the signal is attenuated. In this case the low-frequency component, the wander, is let through unchanged in order to follow the stable long term behavior of the clock, and the high-frequency part, the jitter or noise, is filtered out. The frequency in which filtering sets in is called the "cutoff" frequency. In SDH-based networks, this frequency is typically 10 Hz. In time recovery circuits used for IP connections, this cutoff frequency is much lower, typically in the 10 mHz region. (Compare this difference to the example in the beginning where the "jitter" of a typical IP connection were three orders of magnitude higher than for an SDH connection with pointer adjustments, then it can be seen that the lower cutoff frequency makes sense.)

By understanding the basics of clock recovery it is possible to estimate what PDV limitation is required to acquire the clock to a certain accuracy. Consider this example: A video signal should be recovered to a timing accuracy of 1 µs. It is subject to a PDV with a dominating 1 Hz jitter component. Now, assume that the clock signal is filtered by a first order filter with a cutoff frequency at 10 mHz With a slope of 20 dB/decade, this filter would provide about 100 times suppression at 1 Hz, hence the allowable PDV would be 100 µs. This is, as mentioned, a very simplified example to get some order of magnitude understanding. In reality, the complete jitter/wander spectrum of the PDV distribution must be understood, the suppression given in the example is for sinusoidal signals, etc. But at least it provides some insight to what is possible to achieve in terms of jitter/wander suppression.

SHORT STANDARDIZATION SURVEY

There are a number of standardization efforts within the field of IP QoS, three of these will be surveyed in this chapter.

ITU-T specification Y.1541 "Internet protocol aspects – Quality of service and network performance" [1] could be considered as the main reference on this subject and a very good starting point. First approved in 2002, it was last updated as late as of 2011. As already mentioned, this recommendation has such a large scope that it has not been possible to cover the need of the broadcast community in detail, but it has a pertinent description of objectives for broadcasters in its Table 3 that specifies "provisional objectives" for broadcast types of services (with a focus on IPTV services). Provisional in the sense that the they may be altered if new knowledge is put forward to better describe the QoS objectives for broadcasters. Table 3 provides the following objectives for the discussed QoS parameters.

- PLR < 10^{-5}
- PD < 100/400 ms
- PDV < 50 ms, 10^{-5} quantile

As mentioned before, the PLR objective supports the "one hit per day" objective, and is further motivated in Appendix VIII of the Recommendation, on IPTV applications. The PD is hard to give a clear objective as discussed, it will be very much dependent on the context.

However, the PDV objective of 50 ms (compares to about 40 ms at the more commonly used 10^{-3} quantile) is not good enough for higher-quality broadcast services, as can be seen from the discussion above. ITU recognizes this and opens up for recommending a lower value of this parameter.

Another important specification is the Metro Ethernet Forum, MEF CE 2.0 specifications [6]. It does not explicitly mention broadcast services; rather it gives a few service classes and has geographical contexts for each class. For example the highest service class, named "High" with Metro and Regional scopes provides the following specifications:

- PLR < 10^{-4}
- PD < 20 (Metro) / 75 (Regional) ms
- PDV < 5 (Metro) / 10 (Regional) ms, 10^{-3} quantile

Even though the PDV requirements are higher than for the Y.1541 standard, the PLR requirements are relaxed and thus not fully compliant with the high requirements to provide full studio quality for broadcasters. This has not been an explicit goal for this standard body, which has a high focus on the enterprise and mobile backhaul segments.

When it comes to PDV, both the above specifications provide specifications at the PDV distribution width, given as 99.9 or 99.999 percentile, meaning that this amount of all packets will have an arrival time distribution within the limits. The question is, how relevant is this for timing recovery, if only part of the distribution is used for clock recovery? Obviously not very much. As mentioned above, it is really the distribution of the fastest packet that is used for clock recovery, that is the most interesting. The width of the total distribution has its value in that it could be used to design the jitter buffer sizes such that no, or few, packets are lost, but for the timing properties of the carried signal it does not convey much information.

However, a third very interesting standard, that considers pre-selection of timing packets, is the G.8261.1 recommendation [7] that also provides network limits on the PDV. It is also directly aimed at clock recovery and synchronization. This standard targets the mobile backhaul networks' needs for frequency synchronization. (It also explicitly excludes applicability to phase, i.e. time, synchronization that is for further study.) Very simplified, the specification reads as follows:

- *More than 1% of the timing packets shall have a delay that is less than 150 micro-seconds above the minimum delay*

Hence, if the 1% "fastest" time-stamp carrying packets are selected, these should have a delay distribution that is within 150 µs. From the previous example with the low-pass jitter filter, it can be seen that it may be possible to provide a clock with microsecond accuracy, given this PDV distribution. Again, the properties of the jitter distribution, also when using this pre-selection method, is critical for the outcome. There may be low-frequency components in the delay distribution that is not possible to suppress sufficiently to reach the single microseconds accuracy also when having only 150 µs delay spread. But in general this is a step in the right direction for timing specifications.

IP NETWORK BEHAVIOR

The preceding chapters have hopefully provided a slightly better understanding of the QoS requirements for broadcast services. Now it is time to take a look at the possibilities to fulfill these in real IP networks. First a theoretical overview, including some calculated results on expected PDV that are given for a model network, and then performance of some real IP connections will be discussed. It will be evident that real IP connections may be of almost any quality and that it is important also for the broadcaster to monitor that their leased connections really provide the necessary QoS.

Theoretical aspects of IP connection PDV:

As mentioned before, [1] provides a calculation model for a multi-hop IP connection by considering the added effects of statistical multiplexing (queuing), head-of line blocking and router look-up times. A better model is provided in [8] that can be seen in the figure here:

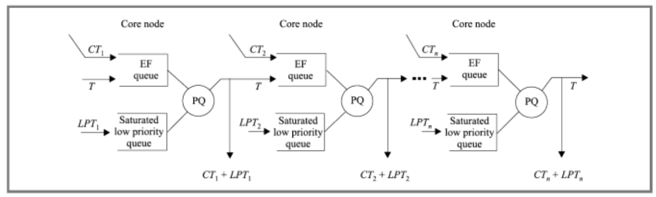

Figure 7. Model of a multi-link connection with a mixture of high and low priority traffic. (From [8])

The model comprises a set of identical cascaded nodes. In each node there are two queues, one for the high-priority traffic for which the PDV should be determined, and one for low-priority, back-ground, traffic. In the model the low-priority queue "saturates" the node, i.e. there will always be a low-priority packet to send when there is no high-priority packet available. The high-priority queue is modeled as an M/D/1 queue. This can be considered as a worst-case distribution for many multiplexed CBR streams.

All packets are considered to be of 1500B in size. The three assumptions, M/D/1 distribution, saturating background traffic and all packets being of maximum size, provide a sort of worst-case scenario.

The model handles head-of line blocking by low-priority packets by modeling the M/D/1 queue to be served by a server with "vacations" that stops serving the queue while the lower-priority packets are served. Also by using the fact that a M/D/1 tail distribution can be very well approximated with an exponential distribution, the convolution over N hops becomes Erlang-N distributed. This finally results in a model that can be executed in a spreadsheet to provide the following example graph:

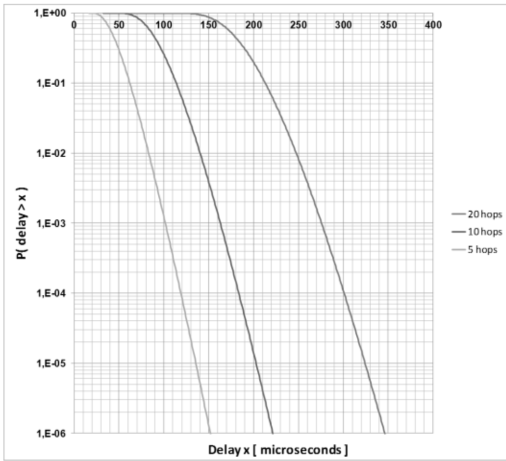

Figure 8. Calculated delay distribution from the model above, with a 30% load of real-time (high-priority) and 70% of elastic (low-priority) traffic. The curves represents the PDV for the real-time traffic over 5, 10, 20 hops over Gigabit Ethernet links.

The figure describes a case with 5, 10 and 20 hops over GbE links, with 30% of the traffic being high priority and 70% low priority traffic. It can be seen that the 99.9 percentile distribution width in the 10 hop case is approximately 160 µs. Roughly the same proportions of the delay variation can be attributed to queuing and head-of line blocking respectively in this case.

The model scales linearly with the link speed such that of 100 Mbps links where used instead, the corresponding 99.9% distribution width would become 1.6 ms instead, showing the importance of using as fast an infrastructure as possible to keep jitter levels down.

Finally it should be noted that theoretical models like the one above are very far from representing actual physical networks with their more complicated patterns of cross traffic, and that have largely varying implementations of classification and policing mechanisms, of queues, of scheduling mechanisms, and with processes such as traffic re-routing and load balancing among others that all will affect the latency, often to higher values than what the simulation indicates. But still the models are valuable to provide a measure of insight to delay processes within a network.

Probed performance of real IP connections:

In this chapter a number of measurements on real IP connections will be present and discussed. They are all monitored using precise probe functionality that are part of the Net Insight's Nimbra MSR equipment [12] and the parameter of interest is the PDV. Three or four measures are presented depending on the software version:

- Peak-to-peak PDV
- 99.9 percentile PDV
- RMS PDV
- (0.1 percentile)

The peak-to-peak PDV measures the maximum delay variation between packets in sample of 100.000 packets. 99.9 percentile provides the interval of PDV that contains 99.9% of the sample. RMS PDV measures the RMS value of the packet-to-packet jitter. The 0.1 percentile is a new probe that measures the distribution width of the 0.1% fastest packet in a sample. This distribution is very important since it will closely relate to the quality of the clock or synchronization recovery as discussed earlier. The simplest and most practical usage of the counters could be to for example:

- Use the 99.9 percentile value to decide the size of the jitter buffer (for example, set jitter buffer size to twice this value)
- Use the 0.1 percentile value to get an understanding of the timing quality of the IP connection

Other probes measures the PLR and also the PD of the link, but these will not be discussed below.

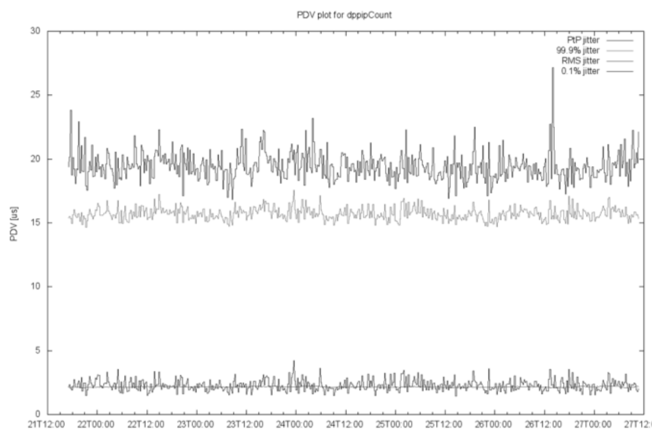

Figure 9. PDV measurement including 0.1 percentile width

The first example is of a rather lightly loaded nation-wide IP network consisting of approximately 10 hops over 10 GbE links. As can be seen the 99.9 percentile PDV is around 16 µs which agrees very well with the 160 µs calculated from the 10 times slower connection in the theoretical example. The agreement should however be considered as coincidental since we don't know more specific details about the traffic parameters for this 10 GbE connection. But it supports the general statement that a higher bitrate on the connection provides lower PDV. We can also see that the 0.1 percentile is in the single µs region which should make it easy to recover studio quality timing from this connection.

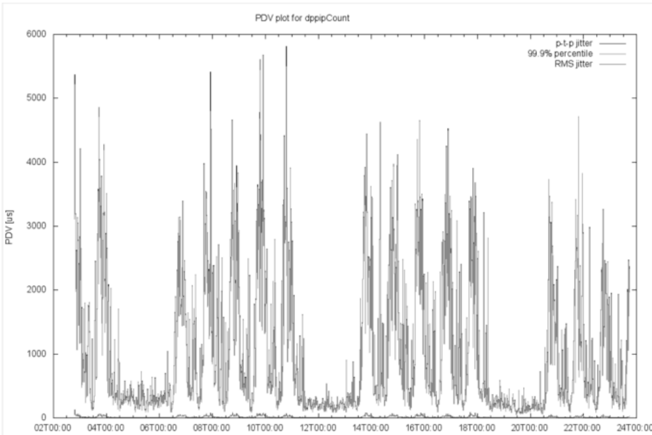

Figure 10. Connection over heavily loaded IP/MPLS links

Figure 10 shows the PDV of a high-capacity connection traversing an IP/MPLS network. The capacity transported here is ~997 Mbps and hence saturating at least the 1 GbE links of the connection. A load dependent jitter distribution with a maximum PDV close to 5 ms at peak hours can be seen, and a more normal PDV at 200-300 µs at weekends.

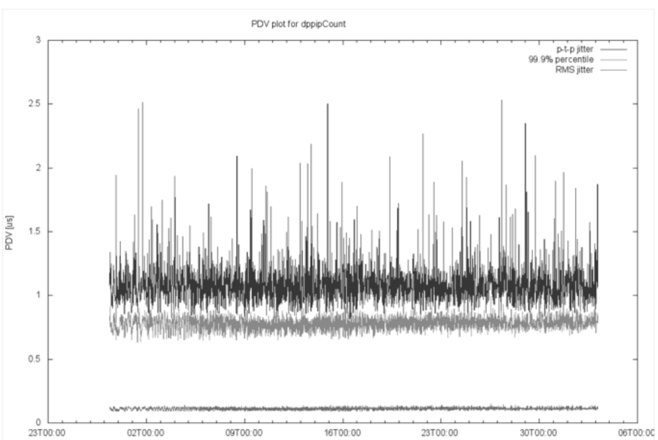

Figure 11. PDV performance for radio-link hop

Figure 11 shows an example of extremely good PDV performance. This is over an Ethernet radio-link for DTT distribution. The transported capacity is around 200 Mbps and there is no disturbing traffic (except for low volume radio-link management traffic). This link is also used for providing absolute time to the SFN network using Time Transfer over IP (TToIP) [12] functionality. The link has a PDV in the single µs range over the several weeks measurement period.

As can be seen, the three examples exhibit very different performances. This is a typical feature of IP networking and also to be expected since different service providers manage services differently with respect to QoS policies, type of infrastructure used, etc. Providing the required QoS for a larger number of demanding services in for example a IP/MPLS network, especially if these are of occasional use type, is a difficult task for a service provider. Broadcasters need to be aware of this and investigate how their services are provided and, when in service, probe the connections and follow up the SLAs.

IP QoS SLAs

As stated in the beginning of this paper, traditionally SLAs used to more or less describe the availability of the particular service with respect to longer or shorter interruptions due to fiber cuts, equipment failure, etc. With the introduction of packet-based networking there are some new quality issues that must be quantified and specified in the SLAs. Some of these parameters, such as PLR, PD and PDV, has been discussed within this paper and are the most important to quantify in an SLA. There are other parameters as well, but these are of less importance and out of scope for this paper.

Furthermore, the way many service providers specify their SLAs may not be appropriate for broadcast services. PLR specifications are in general not a problem, but the specification of the time dependent parameters PD and PDV are often not very useful for broadcasters. For example many service providers provide an average target for the delay of a leased line. But depending on the jitter level a real application may have to use a larger play-out buffer in order not to have buffer tail-drop, offsetting the specified delay with the jitter buffer delay. An alternative would be to specify the delay similarly to how PDV is defined, i.e. to specify for example the 99.9% delay, (or 99.999% to be consistent with a PLR of 10^{-5}) a delay that 99.9% of the packet complies to. For short overall delay links this difference could be significant. This also have the advantage that it relates the PD and PDV in the respect that PD and PDV are both given from the same delay distribution of packets, differing only by a fixed offset (the minimum of the delay distribution).

The PDV objective is usually even more confusing. Sometimes it is given as a single number without stating what part of the PDV distribution that is targeted, or how it is measured. Sometimes it is stated as "average PDV (or jitter)" without interpretation. This could be interpreted as the RMS of the jitter distribution, or the arithmetic mean of the absolute inter-packet delay variation or something similar. In more detailed specifications also the maximum jitter could be described as for example "not exceeding X ms for more than 0.1% of a calendar month". This could be interpreted as the 99.9 percentile measured over a month. The problem with having such a long measurement period is that there could be long contiguous periods with very high PDV, rendering the service very bad for broadcast services, while still fulfilling the SLA. In a packet based network it is not uncommon that the during steady state conditions the network behaves good with respect to PDV, but at times when services are provisioned or re-configured, temporary congestions occur that affects the PDV and hence broadcast services. Today's typical SLA definitions then very much favors the service providers.

All this indicates the needs for enhanced SLA specifications with stricter definitions of the SLA parameters that also takes into account performance on shorter time scales than the monthly averages that is common today.

A complete SLA can be very complex, both to define, but also for customers to understand, making it difficult to follow up and hence learn and improve from it. Service providers within the telecom sector have worked with this subject for a long time and found a way to consolidate performance data in a condensed and uniform way. Within the ITU-T G.826 recommendation [9] a model is used where performance events are divided into two classes:

- Anomalies
- Defects

Anomalies would represent "the slightest deviations from ideal behavior" that, in an IP networking context, could be a lost packet for example. (The same methodology may be used for services where for example an SDI line CRC error would be an anomaly.) These are not network faults in any way, but they affect the performance of the transport. Likewise, defects that are in some respects more serious deviations from ideal behavior, for example, a detected loss of signal (LOS) or "link down," which would be the equal Ethernet defect. Defects can be defined in many ways, for example Y.1731 [10] does not recognize the LOS defect but instead defines a similar loss of continuity (LOC) defect that is based on dropped CCM frames. A way to capture these performance events in a unified manner is to define the concepts of:

- Errored seconds (ES)
- Severely errored seconds (SES)
- Unavailable seconds (UAS)
- Unavailable time (UAT)

Hence a second with at least an anomaly would render an ES. A second with a defect, or where the density of anomalies is above a defined threshold, is marked as an SES. (As a side note, Y.1731 does not seem to embrace the concept of "anomalies", it only considers "defects", and as a consequence only SES, not ES are measured.)

In a simplified diagram the process looks as follows in Figure 12.

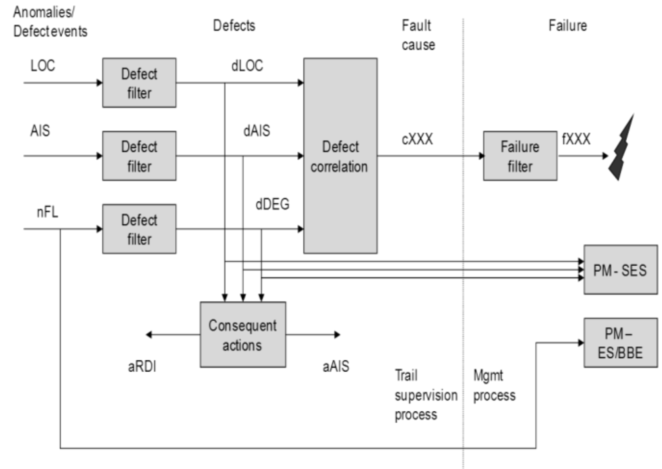

Figure 12. Simplified fault and performance management model

This model describes how performance events are qualified to either defects/faults or just contributes to the performance statistics via the ES/SES counters.

UAS/UAT is important for the SLA since they give the basis for availability calculations. UAS is defined as: getting 10 SES in a row starts a period of UAT where the first UAS is the first SES in the sequence. Likewise UAT ceases when a sequence of 10 non-SES seconds emerges, and UAT then ceases with the first non-SES in this sequence. (See [9] for a more detailed description.)

In this way it possible to describe the SLA in simple terms where the complexity is hidden in the actual definitions of the anomalies and defects. It also provides a clear definition of UAT to be used in availability specifications.

Then there is a simple matter for the broadcaster to follow up the SLA by just counting the ES, SES, and UAS/UAT. A convenient and much used way to do this is to collect this performance data in 15 min / 24 h "bins" such that a table is provided that lists the number of, ES/SES/UAS etc. for each quarter of the day, and for each day [11].

This methodology provides a simple interface to the SLA-related performance measurements that are easy to use in a business context. The difficult part is to agree on what performance events should trigger ES or SES. For packet drops this is easy, one or more packet drops triggers an ES while if the number of packet drops during a second in relation to the nominal packet rate is above a certain threshold, SES should be declared. For other parameters, like PDV, this is trickier. We saw earlier that the typical SLAs of today handles performance objectives for PD and especially PDV in a way that is not good enough for broadcasters. To be able to define performance events also for PDV that could render seconds as ES or SES could be very useful and the SLAs would be more uniform. On the other hand, some timing objectives, that are related to wander for example, needs longer measurement periods. So further studies of how to specify PD/PDV and the related timing performance are needed. Until good definitions are at hand, it may be more beneficial to just display the probe values, consolidated in some form, during each 15 min / 24 h period.

SOME FINAL REMARKS ON THE STATE OF IP QOS

How network congestion in IP networks (which is ultimately what packet loss and delay variation is all about) is handled differs very much between different service providers, and also for connections within the service providers network. The requirements for professional media services are very different from the requirements for normal data services that most service providers are used to. In order to reach full production or studio quality, much stricter requirements must be set on the IP connection QoS than what is common today. It is possible to engineer connections with the requested QoS, but it becomes an operational challenge for service providers when the services become many, and even more so if services are not static, but provided for occasional use. Use of an IP technology with support for synchronous scheduling and switching of services will make it easier to provide the needed QoS. Broadcasters can also help service providers by providing the right requirements for production or studio quality connectivity, monitor their leased connections and feedback their measurements to the service providers. This requires that vendors provide relevant probing functionality in their equipment.

REFERENCES

[1] ITU-T Y.1541 (12/2011), "Internet protocol aspects – Quality of service and network performance"

[2] SMPTE ST 2022-1:2007, "Forward Error Correction for Real-Time Video/Audio Transport Over IP Networks"

[3] Geoffrey M. Garner, Felix Feng, "Delay Variation Simulation Results for Transport of Time-Sensitive Traffic over Conventional Ethernet" IEEE 802.3 ResE SG 2005-07-18

[4] Geoffrey M. Garner, “End-to-End Jitter and Wander Requirements for ResE Applications", IEEE 802.3 ResE SG 2005-05-16

[5] ITU-T G.8263, "Timing characteristics of packet-based equipment clocks"

[6] MEF 23.1, "Implementation Agreement Carrier Ethernet Class of Service – Phase 2", January 2012

[7] ITU-T G.8261.1 (02/2012), "Packet delay variation network limits applicable to packet-based methods (Frequency synchronization)"

[8] Sylwester Kaczmarek and Marcin Narloch, "Methods for evaluation packet delay distribution of flows using Expedited Forwarding PHB", Journal of Telecommunications and Information Technology, 2/2004

[9] ITU-T G.826 (12/2002), "End-to-end error performance parameters and objectives for international, constant bit-rate digital paths and connections"

[10] ITU-T Y.1731 (05/2006), "OAM functions and mechanisms for Ethernet based networks"

[11] ITU-T G.7710 (02/2012), "Common equipment management function requirements"

[12] Refer to *www.netinsight.net* for more information on Net Insight's media networking technology

Original article published in NAB Broadcast Engineering Conference Proceedings 2014. This version is slightly updated.

Stockholm, May 2014.
bengt.j.olsson@netinsight.net