Timing Challenges for Media Production and Transport over the Wide Area Network

Bengt J. Olsson
Net Insight AB
Stockholm, Sweden
bengt.j.olsson@netinsight.net

Abstract – *The transition from synchronous transport methods to IP for production quality video and audio services implies a deterioration of the timing quality of the signals, since IP is a completely asynchronous transport format. IEEE 1588 (PTP) has been introduced to convey precise timing over an IP network. In the Studio LAN, reaching the required quality level of timing transport with PTP is manageable, but doing it in the WAN is very challenging. In practice the only wide area solution today is to use GPS controlled Grand Master (GM) clocks at each remote location to provide synchronization with sufficient quality. This scheme has some disadvantages in that it depends on that a GPS signal is always available, which may not be the case, for several reasons. The industry also needs solutions for PTP transport over the terrestrial WANs. We will describe the challenges with PTP over the WAN, but also show how a synchronous overlay on top of the IP network can provide sufficient PTP quality over the WAN for professional Audio and Video services without resorting to GPS controlled GMs at each end point. Real operational cases using this functionality will be described.*

INTRODUCTION

The broadcast industry is presently in a transition phase with respect to its media production methods. In the core of this transition broadcasters are abandoning many traditional broadcaster specific work flows and instead adopting work flows which to a higher degree are based on Information Technology (IT) methods and technologies. These work flows are more software oriented and utilize COTS hardware, such as standard computers and IP/Ethernet network switches wherever possible. This development also naturally implies a shift to an all IP network infrastructure for media production, since IP is the base for IT networking.

IP networking provides many advantages for the broadcasters in terms of flexibility, abundance of equipment choice, standardized interfaces, unified cabling and network services, compared to the more rigid and specialized SDI/coax cabling methods that traditionally have been used for broadcaster Audio/Video (AV) services.

The adoption of IP for transport also enables the possibility of a more distributed work flow which is very important for the broadcasters. It has been used for a long time in non-linear production, but for live production the limited network capabilities has been a prohibiting factor for distributing production. With IP transport this is changing, and we can already see this manifested in the increasing interest for Remote or At Home production. Live events often have a high content value and distributed production is key for evolving live production and capitalize on its value.

So, while we can see that the transition to IP transport and networking has the potential to bring much more value to the broadcaster than just to provide a larger choice of equipment and a better cabling infrastructure, it comes not without challenges. There are aspects of IP that are not fully aligned with the requirements of professional Audio and Video requirements.

Especially the demands for precise timing and synchronization are challenging. While SDI was designed for broadcast services, IP was originally designed for another category of services that are more elastic than broadcast services to their nature, such as file transfers. IP is fundamentally an asynchronous network protocol with no inherent concept of timing. Its design is optimized for link utilization by having services waiting for transmission until link resources are free. This fits elastic services but is really not a good a fit for deterministic and non-elastic services like audio and video. The generic quality properties of AV transport over IP are described in an earlier NAB conference contribution “IP QoS Objectives for Broadcast Services” [1].

IEEE 1588, also termed “Precision Time Protocol” (PTP) introduces the concept of precision timing into IP. PTP is a two-way time transfer protocol that aims to provide real time over the IP network, accurate to the microsecond level. With PTP broadcasters get a tool to synchronize their AV services over the IP network. But PTP in itself is a service that is subject to the transport properties of the IP network. Meaning that the IP network will have to support the transport of the PTP protocol data to a sufficient degree of quality. This may be manageable in the LAN environment but is generally very difficult to manage in the WAN environment.

In this paper we will briefly describe the basics of PTP and its use for AV services and discuss its applicability in the WAN as well as practical deployment methods for using PTP to support professional AV services.

PTP BASICS

Only the key aspects of PTP, in a generic manner, will be discussed here. For more thorough and precise information on PTP, please refer to [2] and its references.

As mentioned, PTP is a two-way time transfer protocol. This means that the protocol communicates a set of time timestamps between sites A and B and applies arithmetic to these to get a measure of the offset of the clocks that are producing the timestamps at A and B respectively. If one of the clocks is known to be correct, the calculated offset can be used to synchronize the other clock.

A PTP clock has, like any clock, two fundamental properties

- A phase
- A rate of change of phase

The phase is related to the actual moment of time in the used time scale. The rate of change of the phase is the frequency of the clock. The two properties have very different use for broadcasters. The phase, or time, of two broadcast signals may be used for example to synchronize the signals in time for editing purposes. The rate of change of phase, its frequency, is important for the correct play-out of signals, to not over- or under-flow play-out buffers and to keep the transported signal within the AV stability specifications. Note that the use cases for time and frequency are very different and hence their applicability.

Synchronized or Syntonized?

When two clocks have the same time and frequency they are *synchronized*. It the two clocks have the same frequency but not necessarily the same time, they are *syntonized*. Note that synchronized clocks are also syntonized (but the opposite is not true).

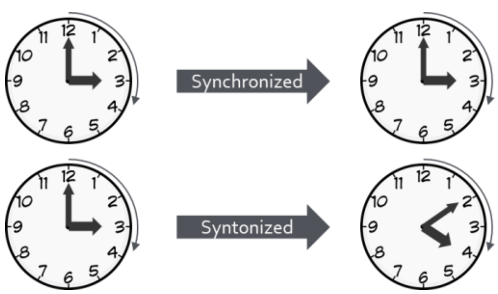

FIGURE 1: SYNCHRONIZED VS SYNTONIZED CLOCKS. BOTH PROGRESS AT THE SAME RATE AS THEIR MASTERS BUT A SYNTONIZED CLOCK MAY HAVE A FIXED OFFSET VS THE MASTER

For some applications synchronized clocks are needed at site A and B, but for many applications only syntonized clocks are needed. This will be important since it is much harder to achieve synchronization than syntonization (which is difficult enough over the WAN).

Another important aspect of clock syntonization is that while it does not produce a global absolute time, it produces a *locally common time* at the remote end. That is, the syntonized clock can be used to synchronize services that are related with each other. For example, video and audio services that emerges from one location and are transported, possibly over separate paths, to another location where they should be synchronized (think "lip-sync"). The syntonized clock handles this since absolute time is not needed, only a local common time. The timing model of MPEG-TS has used this syntonization principle for a long time as will be discussed later.

PTP FOR MEDIA APPLICATIONS

The synchronized "wall clock" paradigm

The basic method for synchronizing services transported over distances using PTP is to *add a constant time offset* to the signal at the egress point. The offset is from the time the signal enters the ingress of the media transport device. This constant offset is called the *link offset*, which is a central term for this discussion. The rationale behind this method is that if a signal enters ingress at Media Device A at time T_A and exits at the egress of Media Device B at time $T_B = T_A + LO$ where LO is the link offset all audio or video samples will be subject to the same constant delay LO independent of how they are delayed in the processing and/or transport. The slack between the link offset and the processing/ transport delay is handled by the egress play-out buffer.

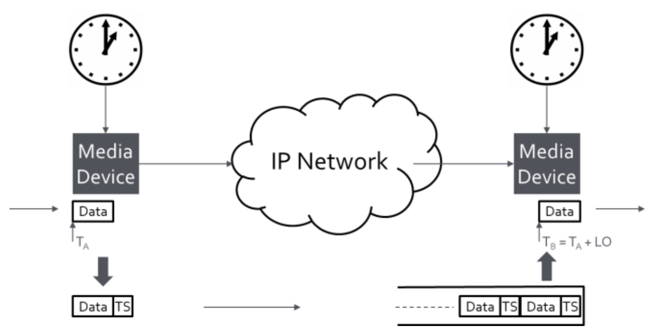

FIGURE 2: THE SYNCHRONIZED "WALL CLOCK" PRINCIPLE

This is a perfectly valid model for a timing transparent transport, but it relies on one important condition: The clocks that provide the times T_A and T_B at each side of the connection must be synchronized or syntonized. We call this the synchronized "wall clock" model, since in principle it doesn't matter what type of clock is used as long as it presents the same time at A and B. This fact provides some flexibility in providing practical solutions for different use cases.

Simplified Media Transport model

Figure 3 below depicts the essential elements of a modern media transport model, using IP for transport. The model is relevant for standards such as SMPTE 2110 and AES67 and similar implementations.

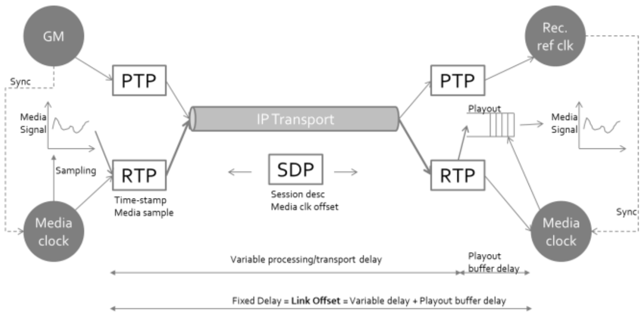

FIGURE 3: GENERIC MEDIA OVER IP TRANSPORT MODEL

The media clocks, that are used for the actual time stamping of the media signal, are synchronized to the PTP reference clocks at each end.

According to [3] the media clocks shall be synchronized to the PTP time scale and share the PTP epoch of January 1, 1970, 00:00:00 TAI. Time stamps are provided with the RTP clock, that is synchronized with the media clock. The RTP clock can be seen as a media clock that wraps at 2^{32} ticks due to the finite 32 bit size of the RTP time stamp field. The RTP clock may have an offset vs the media clock that should be signaled to the peer. (Differs somewhat between specifications, [3] requires this offset to be zero).

The Session Description Protocol (SDP) [4] is used to communicate the capabilities and attributes used by the communication peers.

Note that, following the "wall clock" paradigm above, the egress end does not necessarily have to derive its PTP reference clock from the ingress PTP clock that is transported over the network. It could use a local PTP reference, *as long as the two PTP reference clocks are synchronized, or syntonized* depending on the use case. The PTP clocks could be synchronized via a GPS or GNSS reference or syntonized via a common frequency source. This will be seen in the section about deployment models below.

Compare MPEG-TS timing model

This timing model for media transport is not entirely new. A similar approach has been used in the MPEG-TS system for many years. The MPEG-TS timing model can be said to be the syntonization version of the synchronized wall clock principle, i.e. it transfers a clock, the System Time Clock (STC) using the Program Clock Reference (PCR) timestamp. The transferred STC has a correct frequency but an arbitrary static phase offset with respect to the time of the ingress STC. The Presentation Time Stamp (PTS) takes the role of the Link Offset in this model. Since the objective of MPEG-TS is to provide a media signal with timing integrity and synchronize contained program and elementary streams, only a local common time is needed, which is what syntonization of the STC clocks provide. This has been a very successful paradigm for many years.

PTP DEPLOYMENT MODELS

Depending on the use case PTP time could be provided in different ways. The two typical examples are

Studio to Studio interconnect

In this use case the objective is to connect two studios, possibly to cooperate in the live production like one "virtual" studio, or at least provide a consistent time scale for time codes etc. In this case it is conceivable that each studio can bear the cost of a GPS controlled Grandmaster Clock, since this cost would be divided upon many devices inside the studio. The GM in this case could be considered as an upgraded "house clock" of the studio, providing "genlock" to PTP based media equipment within the studio.

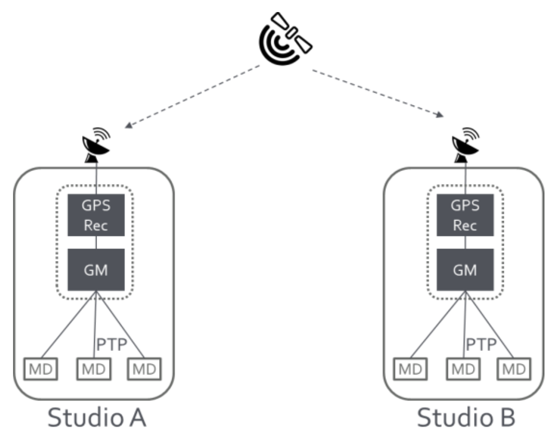

FIGURE 4: STUDIO TO STUDIO WITH GPS CONTROLLED GMS

Each studio would then have synchronized "wall clocks" to handle play-out and synchronization of media streams. Moreover, the variable delay and asymmetries in the connecting network are non-issues since network-carried PTP is not used. High PDV (packet delay variation) in the network will translate to a larger link offset and higher delay but will not otherwise affect the transport quality. (Note however that a high PDV is often associated with high PLR (packet loss ratio) and vice versa, affecting the quality indirectly). The connectivity objectives will focus on keeping the packet drop low and the overall delay low.

In this deployment case the SDP attribute descriptor line that indicates which Grand Master the devices should use would be

`a=ts-refclk:ptp=/traceable/`

to indicate that the reference clock is traceable to GPS.

A potential issue in this case is the reliability of the GPS controlled GM. GPS signals may be disturbed, either unintentionally by weather conditions, cable problems etc., but also intentionally by jamming or spoofing sources. In many cases there are also regulatory requirements that the time source shall be GPS independent. Hence it is recommended to use PTP transported over the WAN as a backup time source. A very good solution for this is to use a frequency aided PTP over WAN transport as described later in this paper.

Studio to Remote interconnect

In the Studio to Remote use case it is, for many different reasons, not viable to synchronize the Remote site by a GPS controlled GM at that site. If the remote site is a single camera for example, it may not have free sight to the GPS satellites and it is not otherwise feasible to have a separate GPS/GM equipment just to drive the timing of that site.

In this case PTP must be transferred over the network, “in-band” to the remote site.

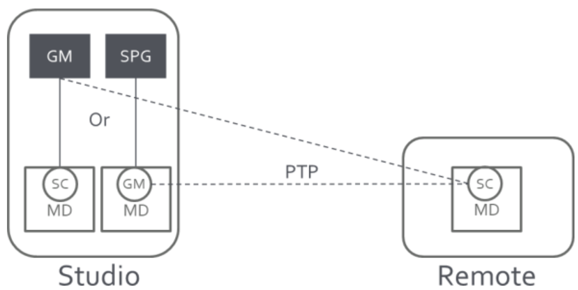

FIGURE 5: STUDIO TO REMOTE WITH “IN-BAND” PTP

The Slave Clock (SC) is configured to connect to the right GM which is attributed for example as
`a=ts-refclk:ptp=IEEE1588-2008:39-A7-94-FF-FE-07-CB-D0:0`
in the SDP message.

NETWORK IMPAIRMENT OF THE PTP CLOCK

A PTP data stream that is traversing a network, in a steady state (not subjected to re-routing etc.), will be subject to Packet Loss (PL) and Packet Delay Variations (PDV) impairments. These factors will influence the steady state performance of the transported PTP clock. In case of sudden topology changes that introduces re-routing in the network, there is also an impairment in form of changed Packet Delay (PD), which can be seen as the minimum delay of a packet from source to destination. This impairment generally affects the asymmetry of the two-way transfer delay and hence changes the offset calculation.

We will first concentrate on the steady state behavior which is mainly affected by the PDV properties of the connection. Given that the PLR is not unusually high, packet loss of PTP packets are seldom a problem, it is the PDV that hits the recovered clock performance.

PDV is inherently due to the asynchronous nature of packet transport. This in turn is manifests itself mainly as[1]

- Statistical multiplexing in routers and switches (queuing)
- "Head of line" blocking in output queues

There may also be other effects contributing to the PDV, such as queue policies in routers and switches, traffic shaping etc. All these effects will contribute to a PDV of the PTP data packets that may aggravate the recovery of clocks, especially in the WAN case.

The AES67 standard recommends an Expedited Forwarding (EF) traffic class for forwarding of PTP data packets. The EF class is characterized by using a “strict priority” scheduling algorithm, meaning that when a PTP packet arrives to the output buffer it should be scheduled to be transmitted on the output link as soon as possible. This minimizes the waiting time through the router, but timing packets will still be subject to some statistical multiplexing depending on the load of EF traffic.

Also the effect of “head of line” blocking is still prevalent for EF traffic. This means that even a high priority packet will have to wait for a lower priority packet that has already started to be transmitted on the egress interface. Depending on the load on the router or switch, this effect only can severely impact the PTP packet traffic if the connection contains multiple hops. Consider that a 1500B packet consumes more than 12 μs to be transmitted on a 1 Gbps link. Thus, a PTP packet could be delayed by 0 – 120 μs in a 10-hop connection with 1 Gbps links even for EF traffic, if there is a high load of large packets in the network. If “jumbo frames” are considered, the head of line blocking waiting times could be even longer.

Simple model of the delay variation for a PTP WAN connection

To provide a “correct” model for determining the delay variation properties of a PTP WAN connection and the following clock recovery is of course immensely difficult. There are so many factors affecting the connection properties in a real network, so every model will be more or less flawed. Also, how the delay variation affects the clock recovery will depend on how the clock recovery is designed and implemented. Here a very simple model will be used in order to understand the challenges that PTP WAN transport is subject to.

Model description

First step is to model the delay variation in the connection itself. For this we are going to base the model on the only to us known specification of packet delay variation for PTP transport, the ITU-T 8261.1 Recommendation [5,6]. This

[1] Refer to [1] for a more comprehensive discussion on causes and effects of PDV on media streams.

Recommendation targets frequency synchronization for mobile backhaul purposes over a 10-hop network denoted HRM-1, Hypothetical Reference Model 1. The HRM-1 network consists of 10 switch/router hops out of which 3 is over 10 GbE links and 7 over 1 GbE links. The HRM-1 model uses no on-path support for PTP, such as Boundary or Transparent clocks (BC/TC) on the switches/routers, which would be unlikely to have in a service provider network.

The importance of the Recommendation for our purposes is that the recommended PDV limits are actually considered as achievable in real networks. The HRM-1 PDV limits are specified as (somewhat simplified)

- *More than 1% of the timing packets, during a 200 second measurement window, shall have a delay that is less than 150 micro-seconds above the minimum delay.*

Thus the Recommendation is based on the clock recovery model of [7], where the clock pre-selects the "fastest" packets only, that ideally gives a low variability. In this case there shall be enough timing packets with a variability less than 150 μs. Importantly, a newer addendum (2015) to the Recommendation adds the following sentence: "For instance, some measurements on HRM-1 networks show that FPP (n, W, δ) $\geq$ 1% is respected when considering δ = 75 μs". Meaning that we could for our model assume a 75 μs spread of the 1% fastest PTP packets instead of 150 μs.

Given the HRM-1 network model (using the 75 μs cluster interval) and a PTP Sync message rate of 8 packets per second (according to the default value specified by [8]), our timing transfer model will have the following properties:

- 1% of the PTP packets are "eligible" for clock recovery
  - For PTP sync packet rate R = 8 pkts/s => 8 x 200 / 100 = 16 pkts / 200 s
- These "eligible" pkts are assumed equally spaced in time
  - 16 pkts / 200 s => Δt = 12.5 s
- Each packet has a random delay within the cluster interval of 75 μs
  - D = Floor delay + Rand[0..75 μs]
- The clock recovery circuit is approximated by a low pass filter using the exponential averager [9]
  - $\alpha = \Delta t / \tau$ where τ is the time constant of the low pass filter

This is of course a strong simplification of the true behavior of a network connection, but is motivated by that any model will be subject to large imperfections vs reality, and that we only need a simple model to draw general conclusions. The largest simplification is that the "eligible" timing packets are equally spaced in time, which is quite un-physical, but makes it easier to implement the model. In reality these packets would be more randomly spread in time, and hence make it harder for the clock recovery circuit, so this makes the model somewhat optimistic. However, the model does not lose its generality because of this and it can be implemented using a spreadsheet, which is done here.

The parameters to vary in the model are

- PTP sync packet rate R
- Clock recovery time constant τ

An example simulation of the PDV and resulting time recovery error is given in the figure below

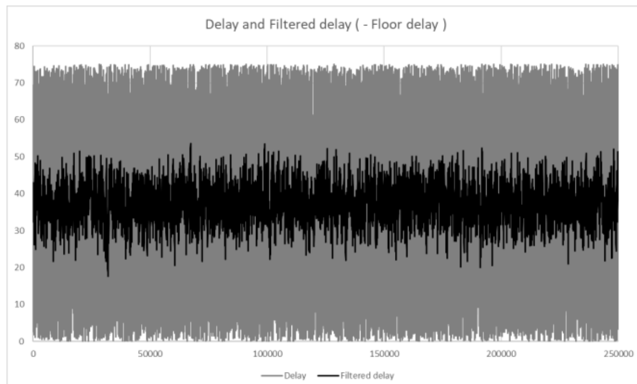

FIGURE 6: SIMULATION WITH R = 8 AND τ = 120 s

The simulation exhibits a peak PDV of 75 μs as expected and a recovered time with a peak deviation of more than 32 μs. This is not surprising since the clock recovery circuitry must track a packet clock with large deviation and with data points spread by 12.5 seconds. One could argue that a "real" clock recovery circuit should use more data points than 1% to get a better statistic, but the reason the clock model of [5] use pre-selection of packages is to keep the variability of the time carried by the timing packets low, and then use a narrow band clock filter to further suppress the variability. In any case it is hard to imagine that any variation of the clock circuitry (except increasing τ to unreasonable values) could get the time error down to region of 1 μs.

It can also be noted that the time constant τ of 120 s is far from the approximately 5 s that is implicitly assumed in a LAN environment [8]. A τ of 5 s in this model would in principle give no suppression of the PDV time deviation. That is one of the conclusions for PTP transport over the WAN. Packet time must be averaged over much longer times than in the LAN, given the higher PDV.

Increasing the PTP sync messages rate R to 128 pkts/s which is the maximum allowed in [8] gives the following clock recovery in this model:

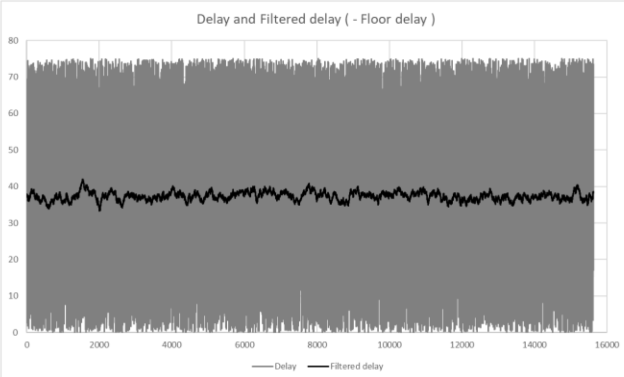

FIGURE 7: SIMULATION WITH R = 128 AND τ = 120 s

Now the recovered time deviation is better as expected. Peak deviation is now about 6 µs with a standard deviation of about 1 µs. This is still far from the peak deviation of 1 µs that is required in the LAN environment. It is also not evident that all equipment will handle this high rate of PTP sync messages.

It would be possible to vary model parameters in many ways, for example changing the clock recovery time constant, or change the cluster interval size, but this is out of scope for this report. The general idea should be clear with these examples. As mentioned, it is relatively trivial to set up this model in a spreadsheet and change the parameters as wanted.

Transient network conditions

Until now it has been assumed that the connection is in the steady state and that PDV is a purely statistical effect. However, in a real network there are also node and link failures that affects the connection in form of service interruptions and re-routes. These will of course also impact the timing performance of the services. The most problematic aspect for PTP is the asymmetry changes such transients experience. Assume that the asymmetry for a connection has been calibrated to provide zero offset between two synchronized clocks during the steady state of the connection. Now a re-route will in most cases severely offset this synchronization and the clocks needs to be re-calibrated. Given that it is hard or not at all possible to calibrate the clocks (in principle GPS clocks are needed for the calibration) this becomes a serious problem with small remedy possibilities when needing absolute time synchronization.

What about the frequency stability?

SMPTE ST 2059 [8] does not specify frequency stability requirements, but the timing requirement of ±0.5 µs accuracy gives an implicit indication of this stability. To find out, examine this simple example: Consider a worst-case timing scenario. This would be when the clock from having a stable -0.5 µs time error (TE) suddenly (at time t = 0) jumps to a TE of +0.5 µs. This represents a phase step of 1 µs as seen in the figure below. The phase response to this step can be modelled by using a first order low pass filter with a time constant τ of 5 s corresponding to the Studio LAN settling time value of [8].

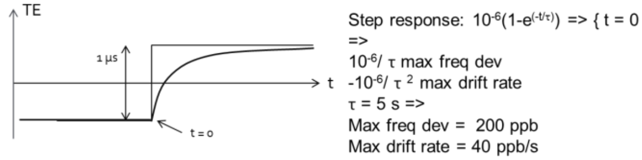

FIGURE 8: ESTIMATION OF FREQUENCY STABILITY

By looking at the 1st and 2nd derivatives of the phase response we can conclude that while not fully meeting the old SDI specs on frequency deviation and drift rate they should be close enough. These requirements should be manageable also in the WAN case. Even if the TE might be larger, the value of τ will, as we concluded earlier, also be larger and compensate more than enough for the larger TE.

USING A SYNCHRONOUS MULTIPLEXING OVERLAY OVER IP FOR SYNC TRANSPORT

The performance of timing transport can be significantly enhanced by using a synchronous multiplexing overlay over the IP network for transporting timing sensitive data. With this technique, the overlay equipment builds a synchronous overlay network on top of the IP network that utilizes synchronous multiplexing and switching of services. Nimbra Media Service Routers (MSRs) uses this technique. The MSRs builds overlay circuits that are encapsulated in IP datagrams. To provide synchronization each IP data packet includes a time stamp, which makes the density of timing information very high, without adding a significant overhead to the data transport.

Without loss of generality we can apply the same model and simulate for transport over such a synchronous overlay by just increasing the timing packet rate to the rates used for MSRs. The minimum packet rate of an MSR is 1000 pkts/s but depending on service bandwidth and choice of MTU, the packet rate can be as high as 8 million pkts/s (at 100 Gbps), each containing a time stamp. An example with 4000 pkts/s (typically corresponding to a service bandwidth of about 45 Mbps) would yield the following clock recovery in the model:

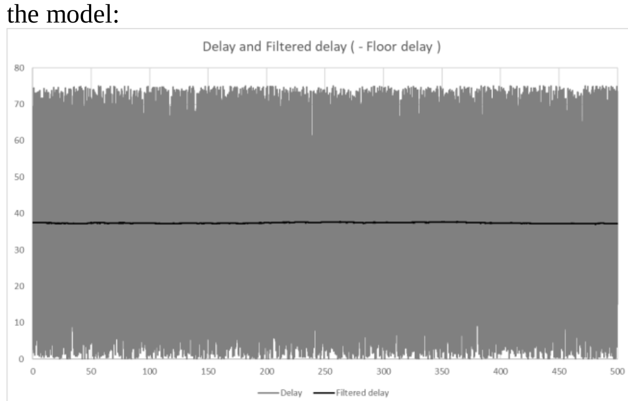

FIGURE 9: SIMULATION WITH R = 4000 AND τ = 120 s

The peak deviation of the recovered clock is now down to less than 0.6 µs with a standard deviation of less than 170 ns. Due to the large density of timing packets we can now suppress PDV down into the sub-microsecond level using a reasonable time constant PDV suppression filter.

The synchronous properties of the MSR overlay is also the foundation for the Time Transfer function that the MSR features. The MSR Time Transfer function provides absolute phase in form of time synchronized 1 PPS and 10 MHz signals. It uses a similar two-way time transfer protocol as used in PTP but utilizes the synchronous properties as described above to keep the absolute phase error low also over the WAN. It is widely in use to provide the phase synchronization used in DTT and DAB networks that uses the Single Frequency Network (SFN) principle where TV or Radio transmitter phases are synchronized to the microsecond level over IP networks. By building redundant overlay links, the synchronous overlay network can also handle node and link failures in the underlay network. However, handling asymmetries is still a challenge given the tight specifications on synchronized time.

EXAMPLES OF PTP WAN SYNCHRONIZATION USE CASES SUPPORTED BY THE MSR

- Sports commentary case

During the 41st World Junior Ice Hockey Championships 2017, a Nordic broadcaster used Ravenna/AES67 based commentary units to report from the games. The connection for 8 bi-directional audio channels was a 15 Mbps Nimbra MSR/ETS (Ethernet Transport Service) channel between Montreal and their facility. In this case the audio equipment supported the needed link offset of about 40 ms (2048 audio sample times at 48 kHz), which is compatible with the 6000 km great circle distance that would give ~30 ms in cable delay. It worked flawlessly according to the broadcaster.

- AES67 contribution and distribution network

This case is for an European broadcaster that requires a GPS independent contribution and SFN broadcast TV/audio network, using AES67/Ravenna for audio transport. They use a setup described in [10] from which figure 10 below is taken. The "Media Router" in this figure is actually a Nimbra MSR that distributes a highly accurate 2.048 MHz clock throughout the complete overlay network. The 2.048 MHz sync signal is locked to the GM and is used to provide for an accurate clock syntonization of the PTP Master and Slave clocks. The PTP equipment also uses techniques to initially synchronize the time, resulting in a very stable PTP slave clock time over the WAN, that is used both for audio contribution as well as for SFN synchronization of DVB-T2/DAB transmitter sites. This network is up and running since the summer of 2017.

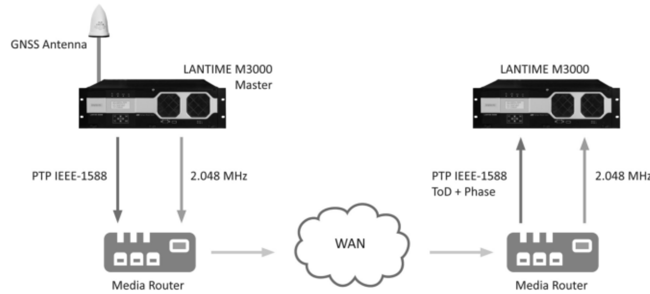

FIGURE 10: MSR ENHANCED TIME SYNCHRONIZATION OVER A WAN. FROM MEINBERG WHITE PAPER [10]

This setup may be generalized to an almost ideal solution for secure inter-studio synchronization: By using a GPS controlled GM on each site, combined with frequency synchronization aided PTP over WAN as backup the following properties can be achieved:

- GPS clock accuracy during normal conditions
- Calibration of the network asymmetry offset for the PTP over WAN backup solution
- Surveillance of the slave site GPS time by the frequency aided PTP over WAN time

It should be noted that a clock which is locked to a network frequency source is not sensitive to network asymmetries.

DISCUSSION

Synchronization options

From the discussion on synchronization and syntonization we can see three main clocking scenarios:

1. Frequency synchronization

Only synchronizing frequency, no phase information at all. This is what was used in legacy media transport, when no phase information of the clock was available. For PTP, its main use is in strengthening the PTP synchronization over a WAN like in the example above. It can also be used to provide frequency synchronization of legacy synchronization equipment (like SPGs). Note that frequency synchronization does not require PTP and is a one-way technique with no asymmetry considerations.

2. Local time synchronization

What we have called "clock syntonization". It provides a clock at the destination that has the right frequency but may have a fixed offset versus the original clock. The reason for the offset is most likely due to non-controlled asymmetries. If the fixed offset is sufficiently small, this option approaches the 3rd option, global clock synchronization. The main requirement in this option is to have a syntonized clock with high frequency stability (low jitter and wander).

With local time synchronization it is possible to synchronize streams that are related in terms of their origin and destination.

3. Global time synchronization

Provides an absolute correct time everywhere, to better than 1 µs accuracy. This option has the highest requirements, demanding both a high frequency stability and full control of network asymmetries end-to-end.

Use cases

From a timing point of view there are three "typical" use cases to consider, with increasing complexity. Most of these use cases will be handled by local clock synchronization (clock syntonization, option 2 above). Only the 3rd use case would need global clock synchronization.

1. Simple streaming of services (contribution)

Example of this is AES67 audio contribution. The objective is mainly to transport an audio stream, for example game commentary, from an event to a central studio. PTP time needs to be syntonized (option 1 or 2), not synchronized, at both ends to play-out the audio stream within frequency stability requirements. There must be some element of time synchronization also to set up the correct link offset, but time accuracy for this is not critical, and could be managed manually as well by adjusting the link offset until the stream is correctly positioned in the play-out buffer.

2. Streaming related media (remote production)

In this use case AV media is recorded on a remote site and processed in the central studio. On the remote site media data is timestamped using the WAN carried PTP clock. Since the remote cameras and audio devices must be gen-locked to the PTP clock there is again a requirement on the frequency stability of the PTP reference. But for synchronization of the media streams at the central studio, only local synchronization, or clock syntonization, is needed (option 2).

3. Phase-accurate synchronization between studios

This use case would be to extend a studio to a larger "virtual" studio over the WAN and use same work flow (at least with respect to timing) as in a single LAN studio. It would then require the same sub-microsecond accuracy also over the WAN, that is full the global time synchronization according to option 3 above. The setup used in Figure 10 would be ideal in this case, it combines the properties of a GPS controlled GM with the frequency synchronization aided PTP over WAN transfer in a very good and secure way as described in that use case.

Remote production falls under use cases 1 and 2. Straight PTP clock syntonization according to option 2 should be chosen in this case. For proper inter-studio synchronization, use case 3, the GPS GM with frequency aided PTP over WAN as backup is an ideal solution.

CONCLUSION

PTP faces difficult challenges if transported over the WAN. The main challenges are:

- Suppression of PDV over the WAN, both in the steady-state and when network transients occur
- Handling asymmetries affecting the absolute time offset that the PTP protocol computes

The first challenge affects both synchronization and syntonization which is why it is of greater importance. If the network handles this challenge, PTP syntonization or local time synchronization, can be supported. This sync scenario would support many uses cases, most notable remote production, in an efficient way. To meet the service requirements as specified today, either on-path support (BC/TC) or transport using a synchronous overlay is likely needed for the WAN use case.

The second challenge affects time synchronization only. Asymmetries makes it very difficult to provide absolute time synchronization to the sub-microsecond level for the typically available network connections. Using a synchronous overlay technique to provide frequency aided PTP over WAN transport is a way to achieve the wanted accuracy over a WAN, either as a primary solution, or as a backup solution when PPS controlled GMs are used to time each studio. This would be an ideal solution for secure inter-studio synchronization.

REFERENCES

[1] Olsson, B. J., "IP QoS Objectives for Broadcast Services", NAB 2014 conference contribution (may be retrieved at https://www.researchgate.net/publication/261476525_IP_QoS_Objectives_For_Broadcast_Services)

[2] https://en.wikipedia.org/wiki/Precision_Time_Protocol

[3] SMPTE ST 2110-10:2017, "Professional Media Over Managed IP Networks: System Timing and Definitions"

[4] Internet Engineering Task Force (IETF) RFC 4566 "SDP: Session Description Protocol"

[5] ITU-T G.8261.1/Y.1361.1 (02/2012), "Packet delay variation network limits applicable to packet-based methods (Frequency synchronization)"

[6] ITU-TG.8261.1/Y.1361.1 Amendment 1 (05/2014), "Amendment1: Revision to clause 8 on packet delay variation"

[7] ITU-TG.8263/Y.1363 (08/2017), "Timing characteristics of packet-based equipment clocks"

[8] SMPTE ST 2059-2:2015, "SMPTE Profile for Use of IEEE-1588 Precision Time Protocol in Professional Broadcast Applications"

[9] https://en.wikipedia.org/wiki/Exponential_smoothing

[10] https://blog.meinbergglobal.com/2017/11/30/select-your-source/