Heterogeneous Networks for Audio and Video

Using IEEE 802.1 Audio Video Bridging

Michael D. Johas Teener
Office of the CTO, Broadcom Corporation
San Jose, CA, USA

Andre N. Fredette, Ph.D.
Extreme Networks
Research Triangle Park, NC, USA

Christian Boiger
Deggendorf University of Applied Sciences
Deggendorf, Germany

Philippe Klein, Ph.D.
Broadband Technology Group, Broadcom Corporation
Airport City, Israel

Craig Gunther, David Olsen
Corporate Technology Group, Harman International
Sandy, UT, USA

Kevin Stanton
Intel Corporation
Hillsboro, OR, USA

Abstract—The IEEE 802.1 Audio Video Bridging Task Group has created a series of IEEE standards that specify methods used to provide the appropriate quality of service for audio and video streams in a heterogeneous network. This paper describes the requirements for such a network and summarizes the methods described in these standards and how they are used in some example higher layer protocols.

Keywords—networks; Ethernet; audio/video streaming; standards; consumer electronics

I. Introduction

Computer networking has traditionally been optimized for "best effort delivery", and that has worked extremely well in the past and will continue to do so in the future for many uses. It is not, however, always good enough when a network is being used to replace the kind of point-to-point connections used for audio and video transmission and other time-sensitive applications.

There have been a number of successful projects to build networks and interconnects appropriate for audio/video delivery[1], but none have succeeded in getting wide market adoption, and none are useful in a heterogeneous network consisting of different layer 2 technologies bridged together. This paper describes the first fully standardized and comprehensive architecture for a bridged, multi-technology audio/video network that is forward compatible with existing standard best effort networks.

1. Best effort

So what it "best effort delivery"? According to Wikipedia (that font of all that is true in the Internet Age), best effort delivery means that it "does not provide any guarantees that data is delivered or that a user is given a guaranteed quality of service level or a certain priority"

Hmm … what is "best" about that?

In practice, it really means "transfer data as quickly as possible". So, in this case best means quickest, and that works! In many, many cases, "best effort" really is best:

- in lightly loaded networks
- where average delay is the primary metric
- if we can't, or don't want to, or it's too much trouble to differentiate between different types of traffic that have different time sensitivities

On the other hand, "best effort" is not best when the *time* is the important metric

2. Audio/video networks: time-sensitivity

"Time-sensitive" in the context of a network has two meanings:

- Data must be delivered within a certain window, typically before a specified maximum delay.
- Connected devices need to have a common sense of wall clock time for synchronization, coordination, phase locking, etc.

Both bounded delay and a well-known time are required in time-sensitive networks, such as those used for live audio and video streaming (and other applications such as control and sensor networks). Even home networks need those attributes whenever multiple devices coordiate to render a particular audio or video stream (think how bad it would be if the various speakers in a stereo or 7.1 presentation were not tightly coupled).

3. Requirements for audio/video applications

The timing-specific requirements for a professional live audio and video network include:

- 2 ms maximum delay. The maximum delay between a musician doing "something" and hearing that same "something" is 10 ms while the transit time of sound from monitor speakers to the musician, plus DSP delays,

---

1. Some examples include IEEE 1394 (commercialized as "FireWire") which is a successful A/V and mass storage interconnect in a relatively narrow market; and CobraNet which is a proprietary audio distribution network based on Ethernet components.

plus mixer delays, plus more DSP delays uses up 8ms so the network gets 2ms for the musician-to-monitor path.

- 1 μs maximum synchronization error. For speaker arrays the maximum synchronization error between speakers must be less than 10 μs and, of course, the designers want (and can use) better: down to 1 μs.

Control and sensor networks have different (and even more stringent requirements), while home networks are typically more relaxed ... although the spectrum of applications in homes ranges all the way up to something similar to "professional".

4. Standardizing a heterogeneous time sensitive network

In 2005, the IEEE 802.1 Working Group created the Audio Video Bridging Task Group (AVB TG) with responsibilities "for developing standards that enable time-sensitive applications over IEEE 802 networks". The IEEE 802.1 WG was the appropriate organization since it is responsible for bridging (including Ethernet "switches") between LANS and interoperability between networks of differing layer 2 technologies.

Given the requirements outlined above, the AVB TG had these goals:

- Provide a network-wide precision clock reference
- Limit network delays to a well-known (and hopefully small) value
- Keep non-time-sensitive traffic from messing things up

Four projects were started to meet these goals:

1. IEEE 802.1AS, the Generalized Precision Time Protocol (gPTP) [2], a layer-2 profile of IEEE 1588 Precision Time Protocol (PTP) [1] with extensions to support different layer 2 network technologies that are based on the IEEE 802 architecture;
2. IEEE 802.1Qav, "Forwarding and Queuing of Time-sensitive Streams" (FQTSS), a specification for a credit-based shaper;
3. IEEE 802.1Qat "Stream Reservation Protocol", registration and reservation of time-sensitive streams (both 802.1Qav and Qat were folded into the overall IEEE 802.1 specification in 2011 [3]); and
4. IEEE 802.1BA "AVB Systems" [4], an overall system architecture.

Together, these define common QoS services for time-sensitive streams and mapping between different layer 2 technologies. They also enable a common endpoint interface for QoS regardless of the particular layer 2 technologies used in the path followed by a stream, effectively defining an "API" or toolkit for QoS-related services for ALL layer 2 technologies.

While the AVB standards were still being developed, the group noted that there was a specification gap between what endpoint applications needed and the services provided by AVB. There needed to be a way to specify how existing applications based on IP (IETF-defined) architecture or IEEE 1394 could take advantage of the new specifications. This gap-filling has been done partially by work done within the IETF AVT group (see [7]) and partially by the IEEE 1722 and 1722.1 Working Groups which have defined streaming formats and management protocols that can enable end-to-end interoperability of professional A/V systems.

Finally, there was a need to ensure interoperability of components that nominally follow the AVB standards. This is not the charter of IEEE or IETF standards groups, so a separate organization, the AVnu Alliance [8], was formed with the specific charter to develop compliance and interoperability tests.

5. Technology outline

The rest of this paper will discuss the technology and specifications mentioned in this introduction, starting with the time synchronization services defined by IEEE 802.1AS and continuing on to the stream reservation and traffic shaping parts of IEEE 802.1Q, and finishing with a discussion of the integration of the various layer 2 network technologies and the IEEE 1722-based higher layers for AVB systems.

II. Time synchronization: IEEE 802.1AS - gPTP

1. Motivation for Network Media Synchronization:

Time, as a fundamental unit of physics, is critically important any time audio or video are rendered, because humans perceive media through our ears and eyes, and our brains integrate these into what is (hopefully) a pleasant experience. We summarize this requirement as proper Media Synchronization. The rule of thumb for media synchronization is that all audio channels must be within 5-20 μs of each other (and stationary), and that video can lead audio by as much as 25 ms but video may lag behind audio by only 15 μs--this is due to the way human brains are wired to perceive late audio as normal, but early audio as unnatural.

Historically, audio and video rendering were confined to a single device (like a TV) or perhaps a set of tightly-coupled systems in an entertainment center. Progress eventually demanded that media be moved or streamed over a network, but to maintain proper Media Synchronization, the audio and video were unpackaged, synchronized, and rendered by a single device or a set of tightly-coupled devices connected with dedicated wires. Again, progress demands that we remove such limitations--users are increasingly demanding that audio and video be untethered from the entertainment center and other media devices--why can't I place my audio devices and video device(s) wherever I want, and use the network to distribute and synchronize the resulting rendering? The chief challenge is maintaining Media Synchronization which, for a good experience, requires orders of magnitude better synchronization than was possible prior to the advent of the IEEE 1588 Precision

Time Protocol and more specifically a profile defined for audio/video applications in IEEE 802.1AS-2011 or gPTP.

2. The gPTP (IEEE 802.1AS-2011) Protocol

gPTP first determines the best source of time in the Local Area Network. In a home, the best source of time is usually a device that isn't coming and going all of the time, and may be configured administratively, e.g., a home gateway. gPTP supports such prioritization but even without such configuration, the protocol will select exactly one device to be the clock master. It turns out that the source of gPTP time need not be the source of all or even any media streams, since the notion of a Presentation Time abstracts Network Time from Media Time, as described later in Section VI and in new improvements in RTP [7]. In the end, gPTP creates a clock tree from the Grand Master through all paths of the LAN that support gPTP, e.g., bridges (and even routers, but that is left for a future publication). In fact, if a legacy hub or buffered repeater is detected, it is automatically designated as "outside the AVB cloud", meaning that time information is not reliable, and reservation parameters cannot be assured (more on that later) as shown in figure 1.

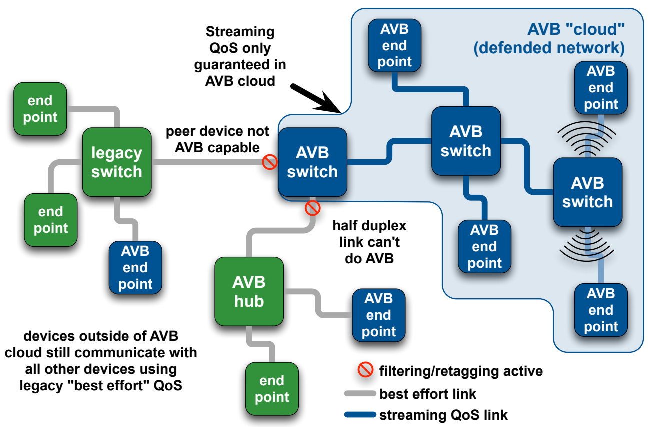

Fig. 1: AVB services cloud

It is important to note that gPTP, as an 802.1 standard, is defined for many different transports increasingly found in the home, including Ethernet, Wireless (commonly, Wi-Fi), MoCA, and G.hn. Thus any of these LAN technologies may be used in any combination, and still maintain accurate time. Each of these standards and industry specifications include a description of how they plug into the gPTP architecture. It bears special mention that the measurement utilized by gPTP for Wi-Fi links is defined as the Timing Measurement capability in IEEE 802.11v-2011.

Once the clock tree is established, the Grand Master periodically sends its time to the next device(s) downstream from it, but rather than relying on software to act in real-time to measure the transmission time, reception time, etc., gPTP defines hardware timestamping mechanisms for each LAN technology that propagates the time with very little degradation from hop to hop. In fact, bridges / APs measure the duration during which each timing packet is held within that device, and even the length/delay of the link between them and the upstream port. This yields extremely accurate time, on the order of a few hundreds of nanoseconds per hop, worst case. With such accurate time, streams may start quickly after system boot rather than wait for minutes for the time reference to "stabilize".

Once the Grand Master's time is known accurately by a talker and a listener, they can use the shared gPTP time as a reference for their media clock. And other sets of talkers and listeners can also use the same gPTP time reference to communicate their media clock—with no requirement that any of the media streams coming from the talkers be synchronized to each other.

III. The Stream Reservation Protocol

1. Introduction

The Stream Reservation Protocol (SRP), as defined in clause 35 of IEEE 802.1Q-2011, is one of the core protocols required for Audio Video Bridging (AVB) Systems. At the highest level, SRP is designed to allow the sources of AVB content (Talkers) to advertise that content (Streams) across a network of AVB Bridges, and users of the AVB content (Listeners) to register to receive the streams through AVB Bridges. SRP is a powerful tool that gives AVB end stations the ability to automatically configure a bridged network to deliver AV content without the need for network administration. In addition SRP is equally able to adjust to engineered networks such as those configured with multiple VLAN segments.

(a) SRP Benefits

In order to appreciate the importance of SRP in AVB systems it is helpful to understand the benefits it offers. Working in concert with FQTSS and gPTP, SRP performs the following functions:

- Allows Talkers to advertise Streams and Listeners to discover and register for Streams;
- Establishes a path through a bridged network between a Talker and one or more Listeners;
- Provides guaranteed bandwidth for AVB Streams;
- Guarantees an upper bound on latency;
- Discovers and reports the worst case end-to-end latency between the Talker and each of its Listeners;
- Automatically configures VLANs between Talkers and Listeners across the bridged network, or automatically adjusts to engineered VLAN networks;
- Reports failure reason and location when a path between the Talker and a Listener cannot be supported;
- Supports emergency priority streams such as 911 telephone calls, and fire and safety announcements;
- Provides a single bandwidth reservation protocol across multiple media types (e.g, wired, wireless and MoCA);

- Supports multiple classes of traffic with different latency targets; and
- Protects best effort traffic from starvation by limiting AVB traffic.

The discovered latency can be reported by Listeners through higher-layer protocols and used, in conjunction with gPTP and the transport protocol, to synchronize the playback of multiple Streams and/or multiple Listeners.

As this list of features shows, SRP offers many benefits beyond the simple establishment of a stream between a Talker and a Listener.

In addition, the IEEE 802.1 TSN Task Group is continuing to work to enhance the capabilities of standard networking for applications such as Automotive and Industrial control. SRP will likely be used for configuring many of these new capabilities.

(b) SRP Applications

SRP can be used in many different applications including Consumer Electronics, Professional Audio/Video, Automotive and Industrial Control. Here the benefits of SRP for Consumer Electronics applications are examined in more detail. The Consumer Electronics (CE) environment is unique it that it is often built with a variety of network types including wired, wireless, coax, power line, and others. In addition it is not uncommon for the network topology and available devices to change from moment to moment.

This constantly changing heterogeneous topology is easily handled by SRP. Since SRP was designed from the beginning to work across multiple network types it can easily establish a reservation with a Talker on a MoCA network, which then transitions through a wired Ethernet segment and on to a Listener connected via a wireless AP.

Existing Listeners can easily establish a stream with a Talker that recently powered on, or just joined the secured wireless network. In a similar way a portable speaker system and/or video display could temporarily be installed and instantly play a movie, even if it is a wireless device. Or, you just bought the newest A/V device from your local CE store, plugged it in, and it was immediately available for streaming to/from all the other existing equipment in your A/V system. All this is possible as a result of the flexibility of SRP, and you don't have to call your resident network expert to get your system running.

Integrated support for emergency services, like a 911 telephone call, is another benefit of using SRP in the home. Thankfully emergency phone calls do not occur very often and it would be unfortunate if a home network always had to reserve a set amount of bandwidth for something that, hopefully, never happens. With SRP there is no need to pre-reserve any bandwidth. In the unfortunate event that an emergency situation occurs, the SRP based network will instantly force other nonemergency reservations off the network so the 911 call can be placed.

2. SRP Technical Overview

This section presents an overview of how SRP is implemented and how it provides the functions described in SRP Benefits. SRP is based on the Multiple Stream Registration Protocol (MSRP) and the Multiple VLAN Registration Protocol (MVRP). MSRP and MVRP in turn are based on the Multiple Registration Protocol (MRP). MSRP additionally works with Forwarding and Queuing for Time-Sensitive Streams (FQTSS) to manage resources, and with the Generalized Precision Time Protocol (gPTP) to discover the SRP Domain.

(a) SRP Operation

The details of MRP are not covered here, but from a high level, MRP defines the rules and procedures to allow applications, such as MSRP and MVRP, to advertise (or withdraw) necessary information across a network and to act on that information in each bridge.

MSRP uses four types of messages including Domain, Talker Advertise, Talker Failed and Listener.

For AVB to work correctly, it must be supported and configured correctly end-to-end. SRP establishes domain boundaries using Domain messages from MSRP and state from gPTP. By exchanging and comparing Domain messages, MSRP determines whether MSRP is operational between the local and peer nodes on a link, and whether the SR class to priority mapping is configured consistently. Similarly, gPTP maintains a variable for each link called asCapable. If `asCapable` is true, gPTP has determined that gPTP is operational between the local and peer nodes on the link. If both the MSRP Domain and asCapable checks succeed, the port is considered to be part of the SRP domain, and streams are allowed to be established over the port. Otherwise, the port is marked as an SRP domain boundary port and streams are not allowed. In addition, any non-AVB traffic that enters through an SRP domain boundary port using AVB priorities will be mapped by the bridge to a non-AVB priority, thus protecting AVB traffic from interference by all other traffic.

Talkers advertise streams by sending Talker messages, and listeners subscribe to streams by sending Listener messages. As illustrated in the following diagram, Talker messages are flooded over the ports on which SRP is enabled, while Listener messages are forwarded only back to the source of the Talker.

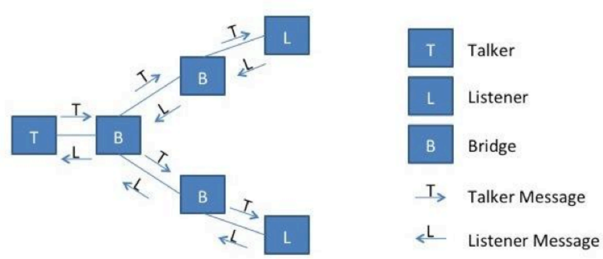

Fig. 2: Stream reservation

Talker Advertise messages contain the following information necessary to make a reservation:

- Stream ID (48-bit MAC address associated with the Talker plus a 16-bit ID)
- Stream DA
- VLAN ID
- Priority (determines traffic class)
- Rank (Emergency or non-emergency)
- Traffic Specification (TSpec)
  - Max Frame Size
  - Maximum number of frames per class measurement interval
- Accumulated Latency

The TSpec and Traffic Class are used to determine the bandwidth required for the stream. As Bridges forward the Talker Advertise messages across the network, they evaluate several factors to determine whether a reservation can be successfully made. These factors include (among other things) whether sufficient bandwidth exists on each port, whether sufficient resources exist on the Bridge, and whether the port is part of the SRP domain. It is important to note that this is an evaluation of whether it is possible to make the reservation, and the resources are not reserved until a Listener message is received as described below. As each node forwards the Talker message, it updates the Accumulated Latency field in the message with the worst case latency for the given hop. A discussion on how the worst case latency is calculated and guaranteed is discussed in Section IV on FQTSS. When the Talker message arrives at a prospective Listener, the accumulated Latency field carries the end-to-end worst case latency for the stream from the Talker to the Listener.

If any device on the path from Talker to Listener determines that the stream cannot be supported, it changes the type of the message from Talker Advertise to Talker Failed, and adds the following additional information to the message:

- Failure Information
  - Bridge ID where the failure occurred.
  - Reservation Failure Code to identify the reason for the failure.

The Failure information allows a control system or administrator to pinpoint the exact location of the failure in the network, the reason for the failure, and fix it.

Listeners indicate that they want to receive a stream by sending a Listener message. The Listener communicates the status of the Stream by sending either a Listener Ready if it received a Talker Advertise or a Listener Asking Failed if it received a Talker Failed. Bridges use a third type of message called the Listener Ready Failed message to indicate that both Listener Ready and Failed messages have been received on two or more ports.

Reservations are made as the Listener messages are propagated back toward the Talker. When Bridges receive a Listener Ready (or Ready Failed) message for a valid stream on a given port, they make a reservation on that port by updating the bandwidth on the FQTSS shaper for the queue associated with the traffic class, updating available bandwidth for the given port, and adding the port to the forwarding entry for the stream VLAN ID/DA to allow the stream to flow. They then propagate that Listener message toward the Talker. When the Talker receives a Listener Ready (or Ready Failed) message, it may begin transmitting. If the Talker receives a Listener Asking Failed it knows that there is one or more Listeners that have requested the stream, however no reservations could be created.

Both Listener and Talker must use MVRP to join the VLAN indicated in the Talker Advertise message prior to sending the Listener Ready or starting stream transmission, respectively. Tagged packets are needed for AVB traffic to communicate class priority, and MVRP enables the AVB end points to automatically configure the necessary VLANs on the AVB bridges.

As one might imagine, there is also a procedure for withdrawing streams and reservations, but that is not covered here.

(b) Emergency Streams

A key feature of SRP is support for Emergency streams. In general, bandwidth is used by streams on a first come first served basis. However, as mentioned earlier, it may be necessary to transport an emergency stream across the network. SRP uses Stream Rank to allow emergency streams to preempt non-emergency streams when all bandwidth is being used.

(c) Automatic Network Configuration

While it may be possible to statically engineer a network for A/V content, configuration of VLANs, priority to queue mappings, and engineering bandwidth requirements is cumbersome and error prone. SRP does all of this automatically, and when there is an error, it identifies exactly what it is and where it occurred.

(d) SRP Protection and other Features

The use of SRP and FQTSS provides protection for both AVB traffic and non-AVB traffic in a number of ways.

- The SRP domain detection mechanism ensures that if a stream has a valid reservation, AVB is supported end-to-end.
- Frames received on SRP domain boundary ports are prevented from interfering with AVB traffic.
- AVB Talkers are required to make reservations prior to transmitting; therefore, they don't use more bandwidth than is available in the network.

- The amount of bandwidth available to each SR class is determined by the configurable deltaBandwidth parameter provided by FQTSS; and credit-based shapers are used on AVB queues to limit the bandwidth to no more than what is reserved. Because this upper limit is placed on AVB stream traffic, the remaining bandwidth is reserved for non-AVB traffic.
- This shaping is also protects valid AVB streams from mis-behaving Talkers. If a Talker transmits at a rate higher than allowed by its reservation, the shaper on the first bridge will limit the traffic, and therefore limit the damage a mis-behaving Talker can do to the rest of the network.
- By managing the forwarding entries for AVB traffic, SRP limits transmission of that traffic to ports that have valid reservations.
- While not explicitly required, it is highly recommended that bridges drop frames with AVB priorities received on AVB ports that don't have a reservation.
- Non-AVB traffic is allowed to use any unused bandwidth that has been reserved for a stream.

3. The Future of SRP (Not just for AV anymore)

While AVB may have started as a solution for transporting audio and video over data networks, it has been recognized that the capabilities provided by AVB help to solve the general problem of running time-sensitive applications over networks. As such, AVB is being applied to Automotive, Industrial control and other problems spaces, and new features are being evaluated.

The following information describes some of the enhancements that may appear in the Gen 2 release of SRP. Be aware that none of the features discussed here are guaranteed to be implemented in the next generation of SRP.

- Redundancy and failover support;
- Pre-configured (static) reservations;
- Configuration of various traffic shapers;
- Reduced latency based on packet preemption;
- Standard-based support for configuring SR class priority and default VLAN ID;
- Integration with Layer 3 (IP protocol) support;
- Configuration of Ingress Policing;
- Dynamic changes to bandwidth and latency;
- Report worst case latency assuming no additional reservations allowed;
- Configurable worst case latency in a bridge which will be used to restrict reservations;
- Link aggregation;
- Multiple Talkers per stream;
- Expanded support for Energy Efficient Ethernet.

Obviously the intent is for SRP to add functionality as the protocol continues to evolve. What that functionality might be is currently under discussion.

The desire of the AVB Task Group is for all the AVB protocols to continue to provide more and more capabilities over time. Some AVB detractors have used this as an argument to say that "AVB is not ready yet". Obviously this is misleading since there are products in the market today which illustrate that AVB has successfully delivered on its first generation promises. Just as wired Ethernet speeds are continuing to evolve from 10Mbps to 100Mbps to Gigabit, to 10 Gig, 40 Gig and beyond – AVB will also continue to evolve as well.

IV. Traffic Shaping

1. Introduction

In order to ensure quality of service additional mechanisms besides the stream reservation protocol (SRP) are necessary. IEEE Std. 802.1Q-2005 only described the strict priority transmission selection algorithm for the prioritization of frames. This mechanism follows the basic idea that highest priority traffic goes first. Such a concept works well as long as there is only a small amount of high priority traffic and no need to fulfill hard latency guarantees. This mechanism does not provide a deterministic low latency; hence the number of interfering higher and same priority frames is not limited.

This type of prioritization scheme does not fit to environments in which audio and video streams are the predominant type of traffic, i.e. occupy a big part of the bandwidth. In the past this problem was solved with big buffers in the end stations, which guaranteed, that enough samples are buffered. This solves the problem as long as the buffers in the devices (end stations and bridges) are big enough and the applications do not require low latency.

But many audio and video applications have very high requirements regarding latency (i.e. very low latency) and as the latency of the network is only one part of the total latency, it needs to be in the rage of few milliseconds. In any case the worst case latency needs to be known in order to know how many bytes a device needs to buffer to allow a reliable playback.

Not only applications require low latency, but also the network itself. Latency in a network is also a measure of the memory requirements in bridges. This results of the simple fact, that a frame which is not in transmission has to be stored somewhere (accumulating latency). As the memory in bridges is limited, it is necessary to transmit traffic without undue delay through the network. This especially applies to bandwidth intensive applications like audio and video streams.

2. Credit Based Shaper

It is the goal of AVB to delay traffic of the highest AVB priority (SR class A) no more than 2 ms over 7 hops and of the second highest AVB priority (SR class B) no more than 50 ms over 7 hops. More hops result in corresponding longer delays. In order to achieve these goals the Credit Based Shaper (CBS) was standardized in IEEE Std. 802.1Qav-2010 "Forwarding and Queuing of Time Sensitive Streams" (later merged into the overall IEEE Std. 802.1Q-2011).

The CBS spaces out the high priority AVB stream frames as far as possible. For this the shaper uses the information about the reserved amount of bandwidth for AVB streams, which is calculated by SRP. The spaced out traffic prevents the formation of long bursts of high priority traffic, which typically arise in traffic environments with high bandwidth streams.

These bursts are responsible for significant QoS reductions of lower priority traffic classes in such traffic environments, as they completely block the transmission of the lower priority traffic for the transmission time of the high priority burst. This strongly increases the maximum latency of this traffic and thereby also the memory demands in the bridges.

On the other hand long bursts also increase the interference time between high priority stream frames from different streams (which arrive from different ports) inside a bridge. This interference increases the maximum latency of this high priority stream frames and again the memory requirements in bridges.

Another task of the shaper is to enforce the bandwidth reservation. Hence the shaping is performed on a per stream per class basis in the talker and on a per class per port basis in the bridges. This enforces on the one hand that every AVB stream is limited to its reserved bandwidth in the talker, and on the other hand that the overall AVB stream bandwidth of each port (in talker and bridges) is limited to the reserved one.

AVB stream frames are sent with a specific frequency. For SR class A the minimum packet frequency is 8 kHz and for SR class B 4 kHz. These frequencies are used for the bandwidth reservation. It is possible to use multiple of this frequencies and it is not required that a stream frame is sent in every transmission period, i.e. if a stream with an 8 kHz packet frequency is reserved it is also allowed to send less than 8000 stream frames in a second (e.g. necessary for rate adaptive codecs). The unused bandwidth is not lost and is used for best effort traffic (i.e. non AVB stream traffic).

These frequencies also define the observation interval in which the reserved bandwidth can be measured if there is no interference with non AVB stream traffic. Hence this interval is also called class measurement interval.

On the basis of the reserved amount of bandwidth and the class measurement interval it is possible to calculate two parameters which define the accumulation and reduction rate for the credit.

The shaper algorithm is similar to the leaky bucket algorithm. AVB stream frames are sorted in two queues, one for SR class A stream frames and one for SR class B. The two AVB stream queues have the highest priority (SR class A is above SR class B).

Frames of a specific SR class are only transmitted as long as there is positive or zero credit for this class. When the credit of a class is negative no frame of this AVB queue is transmitted, even though AVB stream frames have the highest priority.

The calculation of the credit is based on the two already mentioned parameters. The idle slope, which defines the rate with which credit is accumulated, is defined as:

$$\begin{aligned} idleSlope &= reservedBytes / classMeasurementInterval \\ &= reservedBandwidth \end{aligned} \tag{1}$$

The send slope defines the rate with which the credit is reduced and can be calculated as:

$$sendSlope = idleSlope - portTransmitRate \tag{2}$$

The credit is calculated according to the following rules:

- If there is positive credit but no AVB stream frame to transmit, the credit is set to zero.
- During the transmission of an AVB stream frame the credit is reduced with the send slope
- If the credit is negative and no AVB stream frame is in transmission, credit is accumulated with the idle slope until zero credit is reached.
- If there is an AVB stream frame in the queue but cannot be transmitted as a non AVB stream frame is in transmission, credit is accumulated with the idle slope. In this case the credit accumulation is not limited to zero, also positive credit can be accumulated.

An example of the credit, ingress and egress of a bridge port is illustrated in figure 3. The colored packets are AVB stream frames (SR class A). Each color represents one AVB stream. The white frame represents an interfering non AVB stream frame (i.e. Best Effort frame). For simplification the ingress of the non AVB stream frame is not shown in the figure.

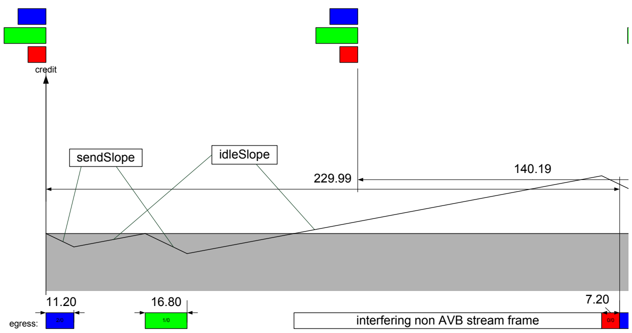

Fig. 3: Credit-based shaper

The Credit Based Shaper spaces out the frames based on the idleSlope and sendSlope. Interfering traffic which blocks the transmission of an AVB stream frame leads to an accumulation of positive credit which allows for a limited burst of stream frames to catch up.

Thus the Credit Based Shaper allows for a converged network with Best Effort and Reserved Traffic (AVB stream traffic) in one network with controlled small latency.

3. Future work in traffic shaping

To achieve even lower latencies in a network, as it is required for control applications in automotive and industrial networks, a new standardization project (IEEE P802.1Qbv) was started in 2012. This project introduces a new type of traffic, the so called Scheduled Traffic.

In order to reduce the latency significantly (compared to the current AVB traffic), it is necessary to reduce the interference between frames with the highest priority, as well as the interference between traffic from lower priority classes with the highest priority class. This can be realized with time aware traffic scheduling.

The scheduling is done in bridges and end stations with the Time Aware Shaper (TAS). The TAS allows for a time based forwarding of frames. This is achieved with the time based connection and disconnection of the queues from the transmission selection.

With this mechanism it is possible to guarantee that the port of a bridge or end station is idle at a defined point in time (t0). For that all queues get disconnected from the transmission selection at a specific time interval before t0, so that the port is idle at t0. Thereby it is possible to schedule the transmission of the Scheduled Traffic frames at these points of time. This guarantees the immediate forwarding of the frames as the port is idle and as a result a very small latency and delivery variation. A small delivery variation is an important factor to keep the schedule and therefore also a precondition for a very small latency.

Hence it is possible to achieve minimum latency and delivery variation for Scheduled Traffic, e.g. in Gigabit Ethernet networks it is possible to reach latencies in the order of few microseconds per hop. An example of a Time Aware shaper in figure 4. The gates connect and disconnect the queues such that no stream frame of a queue is transmitted during the gate closed state.

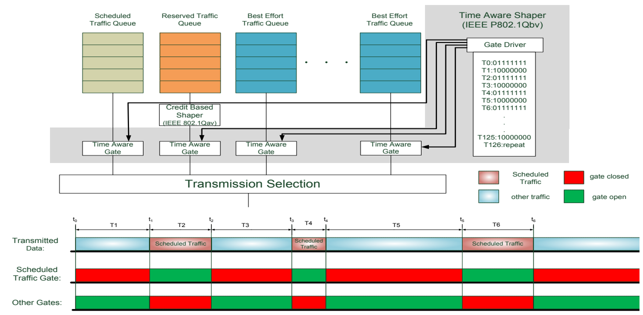

Fig. 4: Time aware shaper example

Further latency improvements are possible with the combination of this mechanism and cut-through switching (i.e. starting the forwarding process after the destination is known and not after the whole frame has been received). In the general case cut-through switching has only marginal or no advantages compared to store and forward switching. As long as the port is not idle and the queue empty, the frame ends up in a queue, even if the bridge is operating in cut-through mode. But the Time Aware Shaper guarantees that the port is idle and thus the frames can be forwarded in cut-through mode after the destination is known. Hence it is possible to use the benefit of cut-through switching.

The CBS and TAS make it possible to build a converged network with Best Effort Traffic, Reserved Traffic (e.g. audio/video streams) and Scheduled Traffic (e.g. industrial/automotive control). Further mechanisms to improve the convergence of these traffic classes are currently under investigation.

Besides the mechanisms defined and investigated in the IEEE 802.1 Time Sensitive Networking Task Group, IEEE 802.3 formed a “Distinguished Minimum Latency Traffic in a Converged Traffic Environment Study Group”. IEEE 802.3 defines the “lower layers” (Ethernet MAC and PHYs). The new study group studies further improvements for network convergence and latency on the “lower” layers.

V. Integration of different L2 technologies

Several standards and industry bodies have defined a variety of networking protocols over the home network and today’s home networks an interconnection of heterogeneous technologies, transporting Ethernet frames over a variety of medium. The more recent OFDM based home network technologies, MoCA for coax, HomePlug AV/IEEE 1901 for powerline and partially Wi-Fi/IEEE 802.11 for wireless networks share common characteristics generically called Coordinated Shared Network (CSN).

A CSN is a contention-free, time-division multiplexed-access network, supporting reserved bandwidth based on priority or flow. One of the nodes of the CSN acts as the

Network Coordinator (NC) node, granting transmission opportunities to the other nodes of the network. The NC node also acts as the bandwidth resource manager of the network.

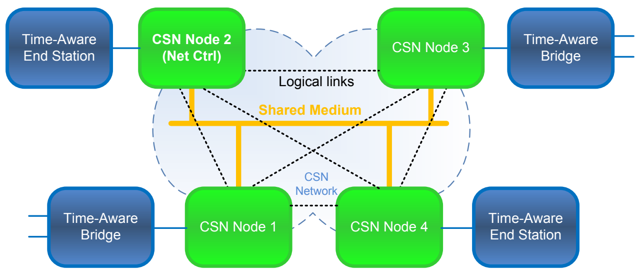

Fig. 5: Example of CSN Backbone in an AVB LAN

CSNs support both unicast transmission for node-to-node transmission and multicast/ broadcast transmission for one-node-to-other/all-nodes transmission. Each node-to-node link has its own bandwidth characteristics which could change over time due to the periodic ranging of the link. The multicast/broadcast transmission characteristics are the lowest common characteristics of multiple/all the links of the network.

A CSN network is physically a shared network, in that a CSN node has a single physical port connected to the half-duplex medium, but is also a logically fully-connected one-hop mesh network, in that every node could transmit to every other node using its own profile over the shared medium.

1. Time Synchronization - gPTP

The way time synchronization messages are propagated across a CSN is dependent of the accuracy of the time synchronization between CSN nodes provided by the CSN native mechanism.

For CSN in which the CSN node local clocks are fully synchronized to the network clock reference with an accuracy that complies with the standard requirements (figure 6-a), the CSN nodes do not need not implement the path delay mechanism but rather treat the path delay as part of the residence time of the distributed system: the Sync message is time-stamped at the edge of the CSN network by the ingress and egress nodes and the reported path delay is the residence time of the message within the whole CSN.

In the opposite case (figure 6-b), each CSN node is treated as an independent bridge with its own free running clock. The path delay across the CSN is the sum of the residence times of both the ingress and egress nodes and the CSN link delay between these two nodes. The path delay measurement either uses a native method (if the CSN features a native mechanism that provides an accurate path delay measurement), or the Pdelay protocol. Sync messages are time-stamped with the CSN clock at the edges of the CSN nodes.

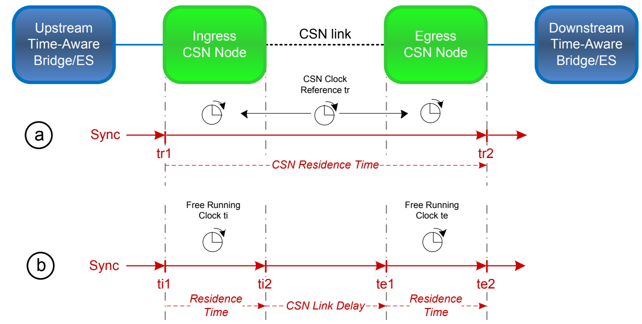

Fig. 6: IEEE 802.AS Sync Message Propagation over CSN

2. Bandwidth Reservation – MSRP

From the bandwidth reservation stand point a CSN network is modeled as a Bridge. Each node-to-node link is equivalent to a Bridge's path from an ingress port to an egress port. The MSRP Service for CSN is the same MSRP Service that manages 802.1 Bridge.

The CSN provides a single entity called the Designated MSRP Node (DMN) which communicates with the MSRP Service to manage the CSN specific bandwidth resources for the MSRP streams.

Depending on the CSN technology, the DMN might correspond to a static node or dynamically migrate between nodes during normal operation. Over time the DMN dynamically constructs its database by handling the MSRP Declarations generated by the nodes of the CSN. If the DMN migrates, the new DMN reconstructs the database by asking the nodes to re-declare their MSRP attributes.

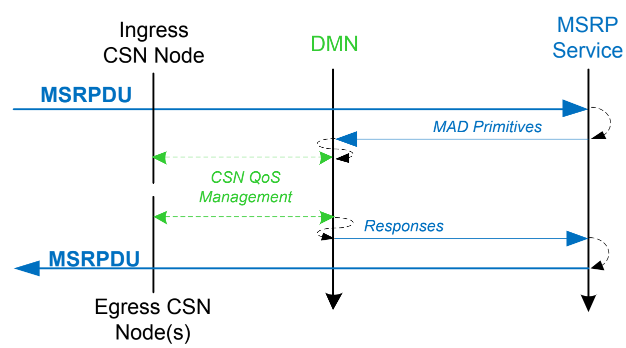

Fig. 7: MSR PDU Flow over CSN

A MSRP-aware CSN node send the MSR PDUs received on its non-CSN interface to the DMN over the CSN. The DMN delivers MSR PDUs, along with information about the originating interface, to the MSRP Service. Upon invocation by the MSRP Service, the DMN translates the MSRP MAD primitives and the MSRP TSpec parameters into CSN Specific QoS primitives and parameters and invoke these primitives to query reserve or relinquish CSN bandwidth resources. After the DMN completes the CSN QoS transactions, the DMN behaves as an MSRP application on a Bridge and propagates (MAP) and distributes (MAD) MSRP attributes.

3. Traffic Shaping

The CSN network is a contention free network in which transmission opportunities on the shared half duplex medium are centrally scheduled by the network coordinator. The NC scheduler shapes AVB streams according to their Tspec parameters.

4. Future L2 Technologies

New developments for AV services are focused on improving the user experience through more resilient network and optimized networking coverage of the home.

A significant effort is currently made to standardize stream bridging protocols supporting multipath to optimize the available bandwidth offered by the whole network topology and provide path redundancy for selected services.

AV services will also take advantage of the converged home networks which better integrates and manages the heterogeneous medium of the network. An important development in this regard is the newly formed IEEE 802.11ak and 802.1Qbz Task Groups aim to standardize the support of 802.1 bridging services over IEEE 802.11 Infrastructure networks and CSNs.

VI. A streaming format for AVB: the Audio Video Transport Protocol

1. Introduction

The Audio Video Transport Protocol (AVTP) is defined by IEEE Std. 1722-2011 and was designed specifically to take advantage of the the new capabilities added to 802 networking by the 802.1 AVB Task Group. When AVB was nearing completion there was no audio/video protocol that was directly suitable for use on AVB networks so AVTP was created

2. AVTP Goals

The AVTP protocol was designed to accomplish the following goals:

- Take advantage of AVB capabilities
- Lightweight protocol to maximize bandwidth usage
- Low Latency suitable for real time applications
- Reuse existing audio/video formats where possible
- Maintains audio/video coherence regardless of network topology
- Multiple media clocks
- Wire replacement

Design decisions for AVTP came from the above goals. AVTP was never designed to transport audio and video across the country the design has been optimize for individual venue sizes installations where a venue could be anything from a small concert or playhouse up to a stadium or large outdoor venue.

By keeping AVTP simple and reusing existing well-known audio/video formats it is possible to maximize interoperability between multiple vendors equipment. It is critical to the success of this technology to keep it simple enough to be used by a garage band yet flexible enough to fill the needs or a large concert hall.

3. AVTP Basic Concepts

There are several basic concepts that are required for any system to transport audio/video data across a network.

(a) Data Formatting

The most basic concept for transporting media is how the media is formatted in a packet. This is a basic concept for interoperability. Where is the data and what is the format? So much work has previously gone into this problem that there is no need to reinvent the wheel. AVTP make use of the IEC 61883 audio and video formats. These are the audio and video formats that have been used for years on IEEE 1394 (Firewire). IEC 61883 defines a rich set of formats including everything form simple mono audio to encrypted surround sound and low resolution raw video to high bandwidth compressed video streams.

(b) Media Clock Reconstruction

In order to maintain real time performance it is critical that the source and sink of audio/video data maintain synchronized media clocks. This eliminates the need for sample rate conversion and greatly reduced the amount of buffering required.

AVTP allows each stream to maintain a separate media clock. This means that a single AVB network can accommodate multiple clock rates. It is not only important that multiple clock rates such as 48Khz and 44.1Khz can be used together, but also to allow multiple streams that use the same nominal clock rate but are not synchronized to be used.

AVTP uses the wall clock defined by 802.1AS to create cross timestamps with designated media clock edges.

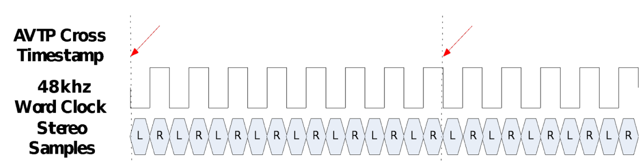

By transporting these cross timestamps along with the associated samples it is possible to precisely recreate the original media clock with the correct sample and clock alignment.

(c) Presentation Time/Latency normalization

Another key concept of AVTP is the presentation time. The presentation time is key to normalizing network latency and maintains sample coherence along multiple network paths. Presentation time is expressed as an offset that is added to the AVTP cross timestamps. The presentation time offset allows audio/video samples to be simultaneous presented to media interfaces regardless of the number of network hops between the source and sink.

AVTP has a default presentation time of 2 milliseconds. This default presentation time offset allows for most networks to operate with real time performance and without unreasonably limiting network topology. However the presentation time offset can be adjusted to accommodate either extremely low latency or unusual network topologies. If network latency lower than 2 milliseconds is desired the number of network hops can be limited to accomplish this. Likewise if a very large network is required a larger presentation time offset can be used to accomplish this.

4. Lip Sync

As you will notice lip sync was never listed in the AVTP goals. However lip sync always comes up in any discussion about audio and video delivery. AVTP was intended as a "wire replacement" with no consideration for lip sync. Lip sync is an extremely complex problem considering that codec delays are not fixed, video and audio codec typically have very different delays, and even room geometry and speaker placement relative to video screens can affect lip sync.

Even though AVTP does not "solve" the lip sync issue it does create a coherent system that can then be used to time align multiple audio and video sources. The network latency in AVTP can be fixed and presentation times aligned regardless of network topology it is possible to calculate the desired additional delays to achieve tight lip sync.

5. The Future of AVTP

The development of AVTP is ongoing and new and exciting features are on their way. One of the great strengths of AVTP is the ability for every stream to have an independent media clock. However there are environments that would prefer to have a shared media clock with multiple media sources using an identical media clock. AVTP is rapidly being adopted in specialized market such as the automotive market. These markets require specialized media formats that are not currently supported. These and other enhancements are currently in development in the IEEE 1722a workgroup.

VII. A management protocol for AVB devices: AVDECC

1. Introduction

The Audio Video Discovery, Enumeration, Connection Management and control (AVDECC) standard defined by IEEE P1722.1 creates a common language for managing AVB/AVTP nodes. A common language to manage AVB/AVTP nodes is a critical piece to allow creation of fully interoperable solution. There are very few networked audio/video systems where every component is from a single vendor. AVDECC enables multivendor system to work together seamlessly

2. AVDECC

AVDECC covers four main areas that are required to manage a streaming media system.

- Discovery
- Entity Model
- Connection Management
- Enumeration and Control

(a) Discovery

The first step with any network management system is to discover all devices on the network. AVDECC Discovery Protocol allows AVB devices to announce their availability on the network, announce they are departing from the network and discover specific or all devices on the network.

(b) Entity Model

In an audio video system there is a need to not just discover a device but also to discover that paths through and the capabilities of a device. The AVDECC Entity Model is used to describe the internal structure of an AVB device. An AVB audio/video device is comprised of network streaming port, other external ports or jacks, and internal ports. In order to intelligently manage an audio/video system a controller needs to be aware of and in control of all these paths. Simply routing audio from a networked media player to an amplifier doesn't solve the problem if the controller cannot then create the connection from the amplifier to the speaker. The AVDECC Entity model allows end-to-end routing of audio/video signals.

(c) Connection Management

AVDECC Connection Management controls the making and breaking of connection between AVB stream sources and sinks.

(d) Enumeration and Control

AVDECC Enumeration and Control Protocol allows AVB devices to be queried to understand their capabilities and use the capabilities. Many audio/video device that seem like single function devices are in fact multifunction. A modern TV cannot be understood by simply describing it as a TV. A TV may contain a video tuner, a video mixer, an audio mixer, an audio amplifier, speakers and a video monitor. For a networked controller to manage a multifunction device each capability must be understood and the controls for each need to be understood. AVDECC provides the ability to enumerate each of these separate capabilities and control these capabilities across a wide spectrum of devices.

3. Summary of AVDECC

By combining all the capabilities of AVDECC a multivendor network audio/video system can be managed from a single controller. All audio and video signals can be routed and each individual device can be controlled. AVDECC is the first management system of this type that has been designed from the ground up to support the audio/video industry.

VIII. Conclusion

The package of standards described are the AVB standards — plus a new layer-2 transport protocol — which are now deployed in the professional and commercial audio market over Ethernet LANs, delivering excellent quality of experience for both content creation and content delivery through accurate time synchronization and deterministic latency limits. The next exciting (and growing) application areas are automotive infotainment and home networks where LAN heterogeneity is an obvious requirement—where product capabilities naturally expand from wired Ethernet to Wi-Fi and other coordinated shared networks like MoCA, HomePlug/IEEE 1901, and HomeGrid/G.hn—all of which are supported by the AVB architecture and standards. With strong industry support through the AVnu Alliance, multiple certification programs for these and other markets are expected to ensure interoperability of devices that implement the AVB capabilities on a diversity of IEEE 802-compatible networks.

IX. Acknowledgment

Much of the material in this paper was derived from contributions made to the IEEE 802.1, 1722, and 1722.1 Working Groups as well as the AVnu Technical Working Group. The references section includes the major document sources, but the authors would also like to acknowledge the innumerable smaller contributions made by the other members of the working groups.

X. References

[1] "Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems", IEEE Std 1588-2009

[2] "IEEE Standard for local and metropolitan area networks—Timing and Synchronization for Time-Sensitive Applications in Bridged Local Area Networks", IEEE Std 802.1AS-2011

[3] "IEEE Standard for local and metropolitan area networks — Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks", IEEE Std 802.1Q-2011

[4] "IEEE Standard for Local and metropolitan area networks--Audio Video Bridging (AVB) Systems", IEEE Std 802.1BA-2011

[5] "IEEE Standard for Layer 2 Transport Protocol for Time Sensitive Applications in a Bridged Local Area Network", IEEE Std 1722-2011

[6] "Draft Standard for Device Discovery, Connection Management and Control Protocol for IEEE 1722 Based Devices", IEEE P1722.1 draft 21, August 2012

[7] http://tools.ietf.org/html/draft-ietf-avtcore-clksrc

[8] AVnu Alliance, http://www.avnu.org

[9] Johas Teener, M., Huotari, A., Kim, Y., Kreifeldt, R., and Stanton, K., "No-excuses Audio/Video Networking: the Technology Behind AVnu," AVnu Alliance White Paper, 2009