# Heterogeneous Networks for Audio and Video: Using IEEE 802.1 Audio Video Bridging

*This paper presents the AVB standards that provide the appropriate quality of service for audio and video streams in a heterogeneous network.*

By Michael D. Johas Teener, Andre N. Fredette, Christian Boiger, Philippe Klein, Craig Gunther, David Olsen, and Kevin Stanton

**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 (QoS) for audio/video (A/V) 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** | Audio/video (A/V) streaming; consumer electronics; Ethernet; networks; standards

## 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/video (A/V) transmission and other time-sensitive applications.

There have been a number of successful projects to build networks and interconnects appropriate for A/V 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, multitechnology A/V network that is forward compatible with existing standard best effort networks.

### A. Best Effort

So what is "best effort delivery"? According to Wikipedia (that fount 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 (QoS) level or a certain priority."

Hmm ... what is "best" about that?

In practice, it really means "transfer data as quickly as possible to maximize throughput." 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 cannot, or do not want to, or it is too much trouble to differentiate between different types of traffic that have different time sensitivities.

On the other hand, this type of behavior can create congestion in the network, and provides no guarantees on

---

Manuscript received February 22, 2013; revised May 22, 2013; accepted May 23, 2013. Date of publication September 11, 2013; date of current version October 16, 2013.

**M. D. Johas Teener** is with the Office of the CTO, Broadcom Corporation, San Jose, CA 95051 USA (e-mail: mikejt@broadcom.com).
**A. N. Fredette** is with Extreme Networks, Research Triangle Park, NC 27709 USA.
**C. Boiger** is with the Deggendorf University of Applied Sciences, Deggendorf 94469, Germany.
**P. Klein** is with the Broadband Technology Group, Broadcom Corporation, Airport City 70100, Israel.
**C. Gunther** and **D. Olsen** are with the Corporate Technology Group, Harman International, Sandy, UT 84107 USA.
**K. Stanton** is with Intel Corporation, Hillsboro, OR 97124 USA.

Digital Object Identifier: 10.1109/JPROC.2013.2275160

[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.

latency or even packet delivery, so "best effort" is not best when the *time* is the important metric.

### B. A/V 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 A/V streaming (and other applications such as control and sensor networks). Even home networks need these attributes whenever multiple devices coordinate 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).

### C. Requirements for Audio/Video Applications

The timing-specific requirements for a professional live A/V 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 digital signal processing (DSP) delays, plus mixer delays, plus more DSP delays uses up 8 ms so the network gets 2 ms for the musician-to-monitor path;
- 1-$\mu$s maximum synchronization error; for speaker arrays, the maximum synchronization error between speakers must be less than 10 $\mu$s and, of course, the designers want (and can use) better: down to 1 $\mu$s.

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."

### D. Standardizing a Heterogeneous Time-Sensitive Network

In 2005, the IEEE 802.1 Working Group (WG) 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 local area networks (LANs) and interoperability between networks of differing layer-2 technologies. As the scope of the AVB TG grew beyond time-sensitive A/V streams to include time-sensitive control streams (such as those used in automotive or industrial automation applications), the TG was renamed the Time-Sensitive Networking (TSN) TG. The TSN TG is continuing the work on audio/video networking described in this paper as well as the new technologies needed for automotive and industrial control systems.

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 Std. 802.1AS, generalized Precision Time Protocol (gPTP) [2]: a very tightly defined 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 Std. 802.1Qav, Forwarding and Queuing of Time-Sensitive Streams (FQTSS): a specification for a credit-based shaper (CBS);
3) IEEE Std. 802.1Qat, Stream Reservation Protocol (SRP): registration and reservation of time-sensitive streams (both 802.1Qav and Qat were folded into the overall IEEE 802.1Q specification in 2011 [3]);
4) IEEE Std. 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 application programming interface (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 (Internet Engineering Task Force (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 WGs which have defined streaming formats and management protocols that can enable end-to-end interoperability of professional A/V systems. The various specifications and standards fit together in a protocol stack between the application-level control and management services and the switching, media access, and link services of a LAN, as shown in Fig. 1.

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

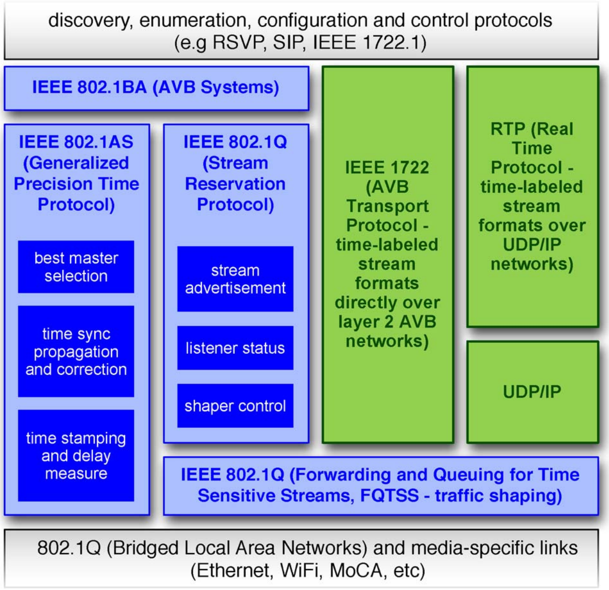

**Fig. 1.** ***AVB protocol stack.***

with the specific charter to develop compliance and interoperability tests.

## E. 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

## A. Motivation for Network Media Synchronization

Time, as a fundamental unit of physics, is critically important when audio or video is 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 $\mu$s 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 $\mu$s; this is due to the way human brains are wired to perceive late audio as normal, but early audio as unnatural.

Historically, A/V rendering was 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 A/V was 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 A/V 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?

## B. Time Synchronization in Networks

Unfortunately, networks used for information technology use, such as web browsing and e-mail were only designed to carry as much information as possible as reliably as possible. Speed was important, efficiency was important, delay minimization was important, but maintaining synchronization was only a secondary concern (at best). All concept of "time" was lost in network specifications except for physical layers. Delays in buffers and queues were not communicated or measured.

So, how do we do it?

The key is the measurement of delay which can be done via a packet exchange such as done in the Network Time Protocol (NTP) and shown in Fig. 2. A client can then use a time value transmitted by a server just by adding the delay.

The delay measurement requires the use of "timestamps": the value of some timer when a particular event occurs, and the more precise and deterministic the event, the better. By communicating the value of timestamps at well-known events, we can correlate network actions with actual time.

There are two primary difficulties in computing the delay.

- First, the timing of the events may not be well controlled. For NTP, this is usually a kernel/driver software event as close to the hardware as possible, but usually with accuracy that, at best, is in the tens of microseconds. To get better results, we use

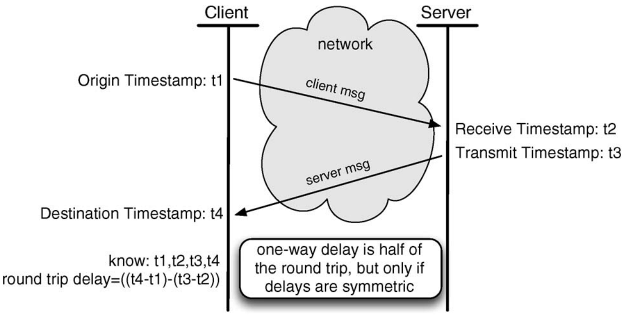

**Fig. 2.** ***Delay measurement.***

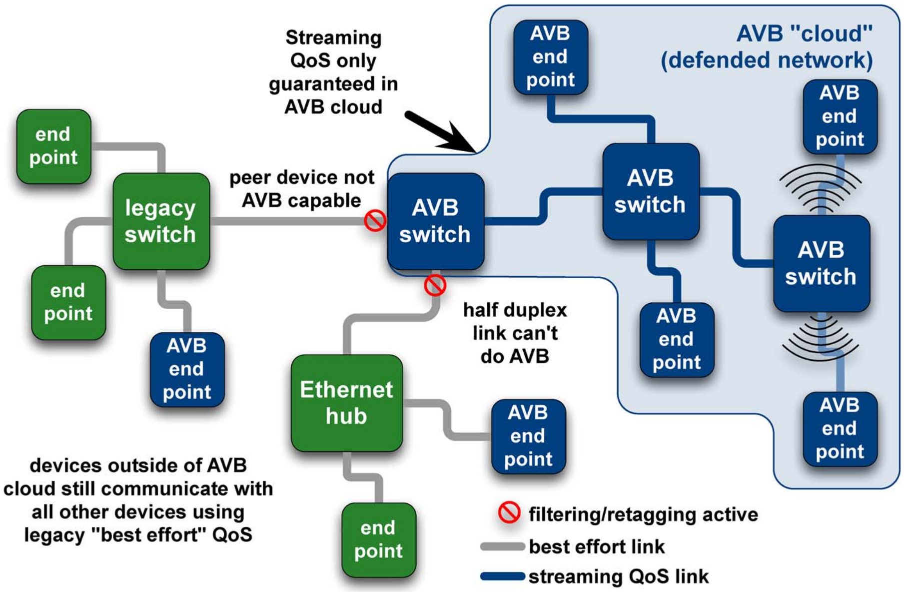

**Fig. 3.** ***AVB services cloud.***

a physical layer event. In Ethernet and WiFi, that is the start of packet, with accuracy on the order of tens of nanoseconds. IEEE 1588 [also known as the Precision Time Protocol (PTP)] standardized the use of physical layer timestamps to compute network delays and define synchronization events.

- Second, and even more important, we need some way to handle the variability of the delays through the network. In NTP and early versions of 1588, this was done by taking many, many samples and applying various types of advanced filters to reduce the effect of the continually changing delays. In the current version of 1588 (IEEE Std 1588-2008, the basis of 802.1AS) the concept of a "transparent clock" was introduced, which dynamically compensated for the network buffering delays. See Section II-C1.

### C. The gPTP (IEEE 802.1AS-2011) Protocol

gPTP first determines the best source of time in the LAN, called the grand master (GM). In a home, the best source of time is usually a device that is not 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 in Section VI and in new improvements in RTP [7]. In the end, gPTP creates a clock tree from the GM through all paths of the LAN that support gPTP, e.g., bridges (and even routers, where the router acts as an IEEE 1588 boundary clock). 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 Fig. 3.

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, WiFi), Media over Coax Alliance (MoCA), HomePlug, and G.996x, as well as infrastructure technologies such as the Ethernet passive optical network (EPON).[2] Thus, any of these LAN technologies may be used in any combination, and still maintain accurate time. Each of these standards and industry specifications includes a description of how it plugs into the gPTP architecture as a "media-dependent" leg, as shown in Fig. 4.

The primary components of gPTP are:

- time distribution;
- link delay measurement;
- best master clock selection.

*1) Time Distribution:* gPTP, which contains a 1588 profile, requires that a time source, known as a GM periodically transmit time information as a "sync" message that contains the time of day (gPTP uses TAI[3] by default) along with a timestamp of when the message was actually sent.

---

[2] See http://en.wikipedia.org/wiki/Ethernet_passive_optical_network.

[3] International Atomic Time (French "Temps atomique international"), which has the advantage over UTC (universal or "Greenwich" time) of not having leap seconds or other discontinuities. See http://en.wikipedia.org/wiki/International_Atomic_Time.

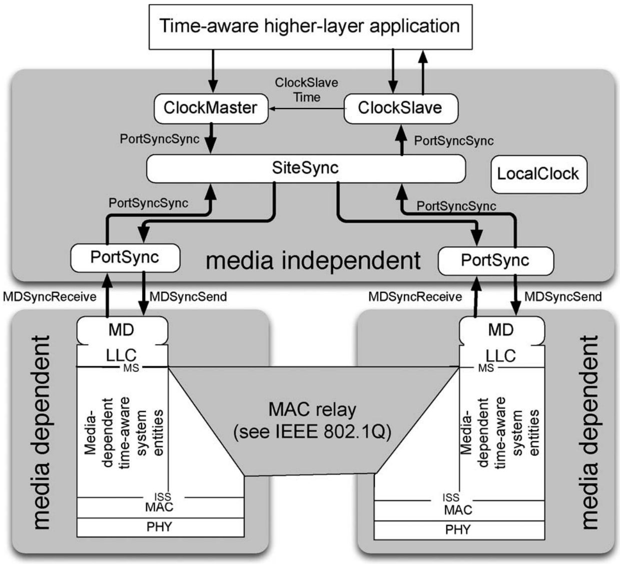

**Fig. 4.** ***802.1AS architecture.***

Since sending the time that a packet is transmitted within that same packet is somewhat difficult, gPTP specifies 1588 "two-step processing" where the termination timestamp is sent in a later packet, the "follow up," as the "precise origin timestamp." This is illustrated in Fig. 5.

Just sending a precise sync sending time is not enough; we still have to correct for delays introduced by intermediate systems such as bridges, switches, and wireless access points. This is done using the 1588 concept of a "transparent clock" where a "correction field" in the follow up message is incremented by the upstream delay and the residence time ($t3 - t2$), as shown in Fig. 6. The correction field plus the precise origin timestamp plus the upstream delay is the correct time.

*2) Link Delay Measurement:* Although the procedures defined above compensate for all the intermediate system delays, there are still the delays due to signal propagation time: the "link delay." This value may be known as an inherent characteristic of the underlying network technology, but is frequently a slowly changing value due to physical parameters. For instance, Ethernet propagation time is dominated by cable length, but is also affected by temperature. gPTP measures link delays, in this case using the delay measurement procedure illustrated in Fig. 7.

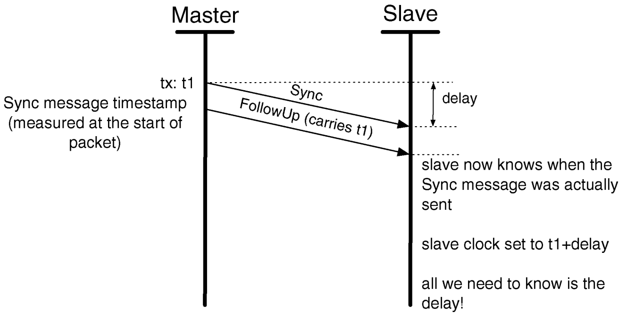

**Fig. 5.** ***Two-step sync message.***

Path-delay measurements for other link types are similarly dependent on the abilities and specifications of those technologies. For instance, the path delay for WiFi links is defined as part of the timing measurement capability in IEEE 802.11v-2011.

*3) Best Master Clock Selection:* A gPTP system only has one active clock source (what 1588 calls a "timing domain"). The selection of this source uses a version of the 1588 "best master clock algorithm," which has been improved to support much faster switchover in case there is a failure of the current GM clock.

In gPTP, all GM-capable devices announce their capabilities to all their immediate neighbors, and intermediate devices (bridges, switches, etc.) retransmit to all their neighbors only announcement of the "best" capability that they receive. If a GM-capable device receives a "better" capability message, then it stops transmitting its own announcements. Soon, there is only one device sending announcements, and that is the "best" GM.

This "best" quality comprised (in decreasing order of importance):

- priority (configurable);
- multiple "clock quality" fields;
- media access control (MAC) address.

If the GM dies and announce messages cease, all GM-capable devices announce and a new GM is quickly chosen.

*4) gPTP Summary:* The gPTP protocol yields extremely accurate time, on the order of a few tens of nanoseconds per hop, worst case, and frequently much better. With such accurate time, streams may start quickly after system boot rather than wait for minutes for the time reference to "stabilize."

Once the GM'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

### A. 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

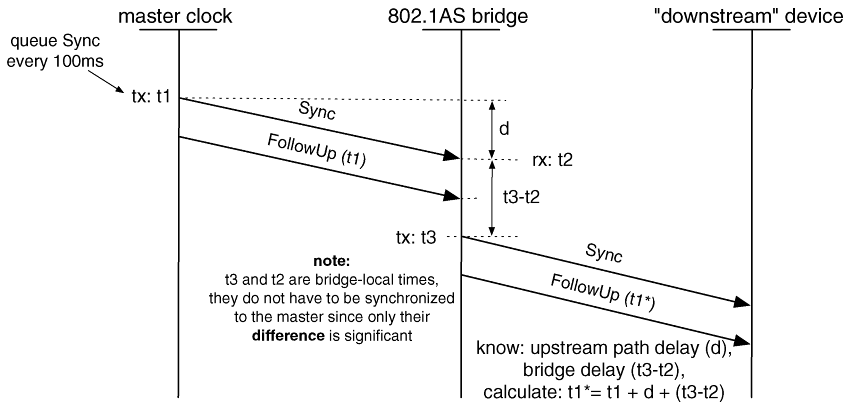

**note:**
t3 and t2 are bridge-local times,
they do not have to be synchronized
to the master since only their
**difference** is significant

know: upstream path delay (d),
bridge delay (t3-t2),
calculate: t1*= t1 + d + (t3-t2)

**Fig. 6.** ***Time correction in a bridge.***

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 virtual local area network (VLAN) segments.

*1) 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., 802.3 Ethernet, 802.11 wireless and MoCA);
- supports multiple classes of traffic with different latency targets;
- protects best effort traffic from starvation by limiting AVB traffic.

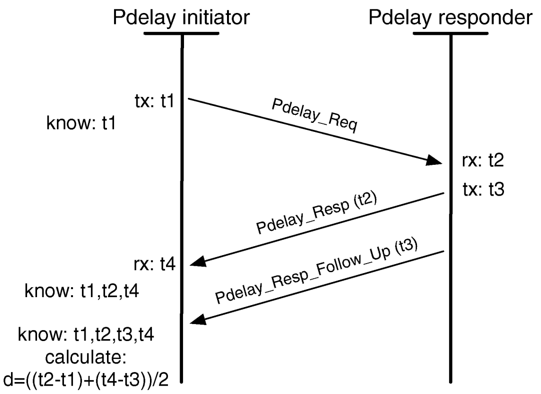

**Fig. 7.** ***Path-delay measurement in Ethernet.***

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 TG 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.

*2) SRP Applications:* SRP can be used in many different applications, including consumer electronics (CE), professional A/V, and automotive and industrial control. Here the benefits of SRP for CE applications are examined in more detail. The CE environment is unique in 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 establish a reservation with a talker on a MoCA network, which then transitions through a wired Ethernet segment and onto a listener connected via a wireless AP.

Existing listeners can 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 do not 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 prereserve 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.

### B. 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 FQTSS to manage resources, and with gPTP to discover the SRP domain.

1) *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.

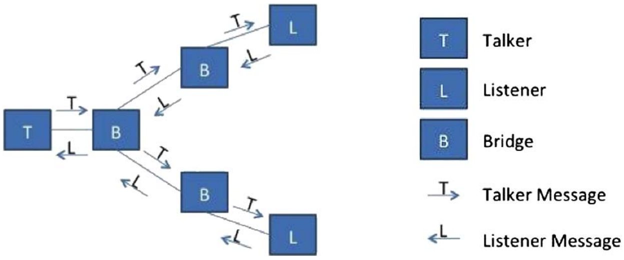

**Fig. 8.** ***Stream reservation.***

Talkers advertise streams by sending talker messages, and listeners subscribe to streams by sending listener messages. As illustrated in Fig. 8, 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.

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

- stream ID (48-b MAC address associated with the talker plus a 16-b ID);
- stream DA;
- VLAN ID;
- priority (determines traffic class);
- rank (emergency or nonemergency);
- 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 at least one listener that has requested the stream but the corresponding reservation could not 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.

*2) 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 nonemergency streams when all bandwidth is being used.

*3) 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.

*4) 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 do not use more bandwidth than is available in the network.
- The amount of bandwidth available to each SR class is determined by the configurable delta bandwidth parameter provided by FQTSS; and CBSs 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 also protects valid AVB streams from misbehaving 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 misbehaving 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, bridges should drop frames with AVB priorities received on AVB ports that do not have a reservation.
- Non-AVB traffic is allowed to use any unused bandwidth that has been reserved for a stream.

### C. The Future of SRP (Not Just for AV Anymore)

While AVB may have started as a solution for transporting A/V 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:

- seemless redundancy;

- preconfigured (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 TSN TG 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 10 to 100 Mb/s to 1 Gb, to 10 Gb, 40 Gb and beyond, AVB will continue to evolve as well.

## IV. TRAFFIC SHAPING

### A. Introduction

In order to ensure QoS, additional mechanisms besides the 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 environments in which A/V 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 A/V applications have very stringent 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 range 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 do applications require low latency, but so also does the network itself. Latency in a network directly corresponds to the memory requirements in bridges. This results from 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 A/V streams.

### B. 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 seven hops and of the second highest AVB priority (SR class B) no more than 50 ms over seven hops. More hops result in correspondingly longer delays. In order to achieve these goals, the CBS was standardized in IEEE Std. 802.1Qav-2010 FQTSS (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 significantly increases the maximum latency of this traffic and thereby also the memory demands in the bridges and end stations.

Additionally, long bursts 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 high priority stream frames and again the memory requirements in bridges.

Another task of the shaper is to enforce the bandwidth reservations. 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 amount.

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, it is 4 kHz. These frequencies are used for the bandwidth reservation. It is possible to use multiple of these frequencies and it is not required that a stream frame be sent in every transmission period. For example, 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 may be 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 the 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 the SR class A stream frames and one for the SR class B stream frames. The two AVB stream queues have the highest priority (the SR class A is above the 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 previously mentioned parameters. The idle slope, which defines the rate at which credit is accumulated, is defined as

$$\begin{aligned}\mathrm{idleSlope} &= \frac{\mathrm{reservedBytes}}{\mathrm{classMeasurementInterval}} \\ &= \mathrm{reservedBandwidth}. \end{aligned} \tag{1}$$

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

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

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, the credit is accumulated with the idle slope until zero credit is reached.
- If there is an AVB stream frame in the queue that cannot be transmitted because a non-AVB stream frame is in transmission, the credit is accumulated with the idle slope. In this case, the credit is not limited to zero.

An example of the credit, ingress, and egress of a bridge port is illustrated in Fig. 9. The colored packets are AVB stream frames. Each color represents one AVB stream. The white frame represents an interfering non-AVB stream frame (i.e., the best effort frame).

The CBS 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 the positive credit that allows for a limited burst of stream frames to catch up.

Thus, the CBS allows for a converged network with best effort and AVB stream traffic in one network with controlled small latency.

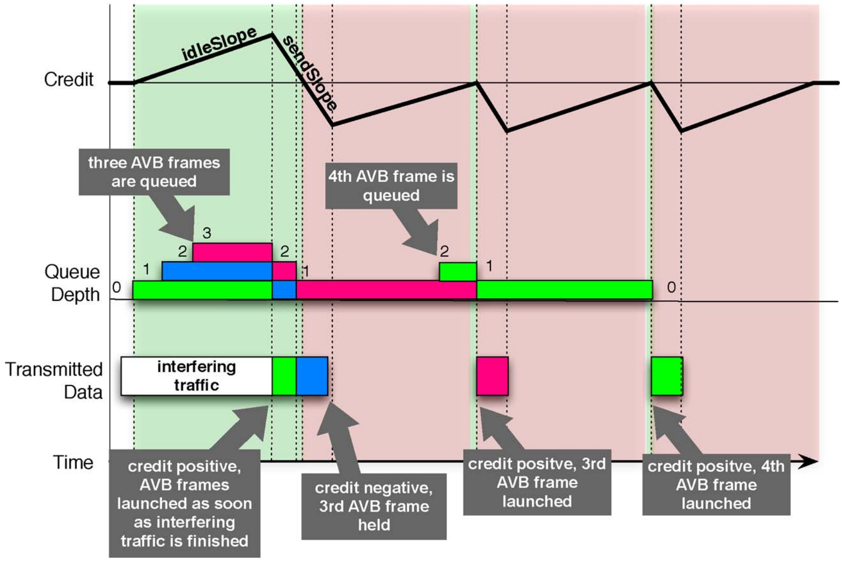

three AVB frames are queued

4th AVB frame is queued

credit positive, AVB frames launched as soon as interfering traffic is finished

credit negative, 3rd AVB frame held

credit positve, 3rd AVB frame launched

credit positve, 4th AVB frame launched

**Fig. 9.** ***Credit-based shaper.***

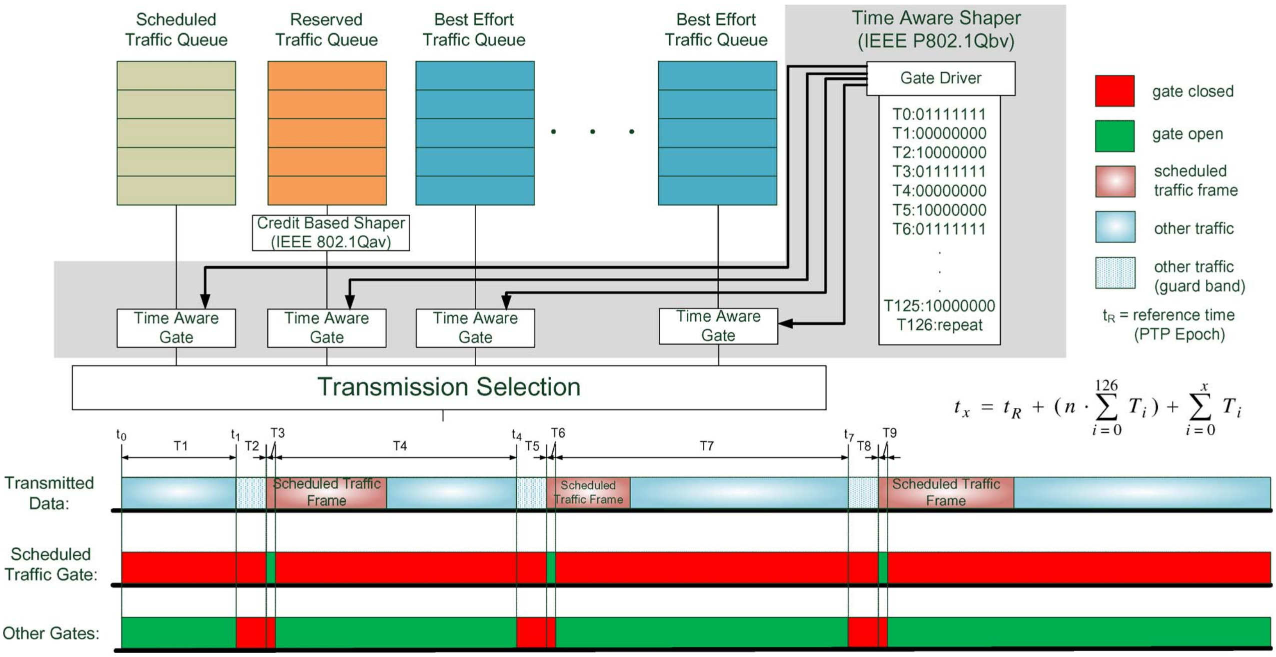

**Fig. 10.** ***Time-aware shaper example.***

### C. Future Work in Traffic Shaping

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

In order to significantly reduce the latency (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 the time-aware traffic scheduling.

The scheduling is done in bridges and end stations with the time-aware shaper (TAS). The TAS allows for forwarding of frames based on the time-based connection and disconnection of the queues within the transmission selection function of an egress port.

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 ($t0$). To make this happen, 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 maintain the schedule and is a precondition for a very small latency.

Hence, it is possible to achieve minimum latency and delivery variation for scheduled traffic. For example, in gigabit Ethernet networks, it is possible to reach latencies on the order of few microseconds per hop. An example of a TAS is shown in Fig. 10. The gates connect and disconnect the queues such that no stream frame of a queue is transmitted during the gate closed state.

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 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 the cut-through mode. The TAS, however, guarantees that the port is idle and, thus, the frames can be forwarded in the cut-through mode after the destination is known.

The CBS and the TAS make it possible to build a converged network with best effort traffic, reserved traffic (e.g., A/V 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 TG, IEEE 802.3 formed the 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.

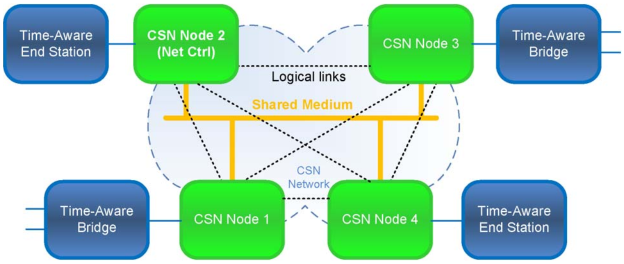

**Fig. 11.** ***Example of a CSN backbone in an AVB LAN.***

## V. INTEGRATION OF DIFFERENT L2 TECHNOLOGIES

Several standards and industry bodies have defined a variety of networking protocols for the home network, and today's home networks comprise an interconnection of heterogeneous technologies, transporting Ethernet frames over a variety of media. The more recent orthogonal frequency-division multiplexing (OFDM)-based home network technologies, MoCA for coax, HomePlug AV/IEEE 1901 for powerline, and WiFi/IEEE 802.11 for wireless networks share common characteristics generically called coordinated shared networks (CSNs).

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.

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 changes in the media or signal propagation environment. The multicast/broadcast transmission characteristics are the worst case common characteristics of all of 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, as illustrated in Fig. 11.

### A. Time Synchronization: gPTP

Time information is propagated across a CSN in a way dependent on the accuracy of the native time synchronization facilities of the CSN.

- If the CSN node local clocks are fully synchronized to the network clock reference with an accuracy that complies with the standard requirements [Fig. 12(a)], the CSN nodes do not need to implement the path-delay mechanism. Instead, they can treat the path delay as part of the residence time of the CSN treated as a 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 [Fig. 12(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 used by Ethernet. Sync messages are time stamped with the CSN clock at the edges of the CSN nodes.

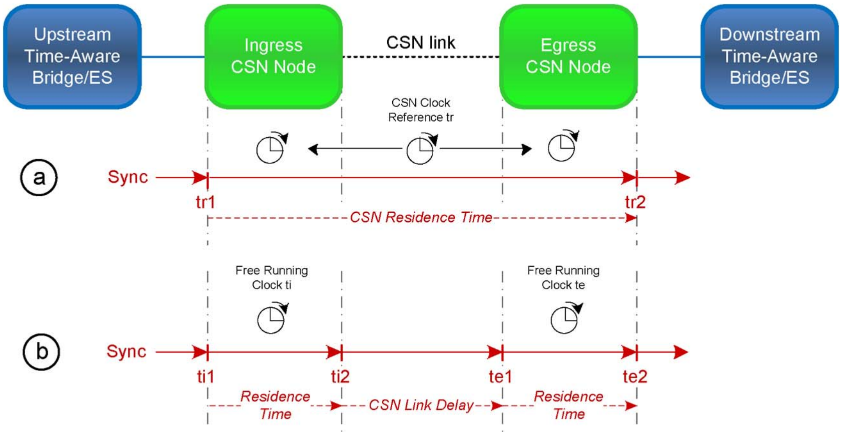

**Fig. 12.** ***IEEE 802.1AS Sync message propagation over CSN.***

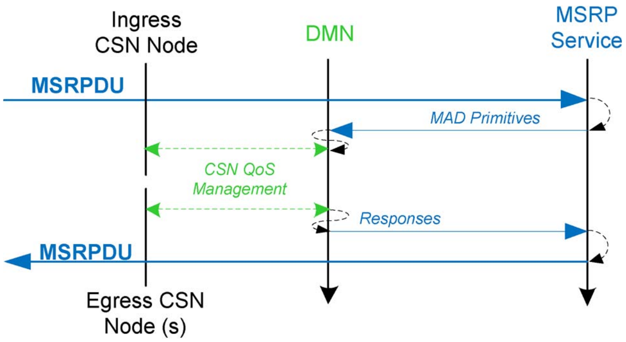

**Fig. 13.** ***MSRPDU flow over CSN.***

### B. 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 is used for an IEEE 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.

An MSRP-aware CSN node sends the MSRPDUs received on its non-CSN interface to the DMN over the CSN. The DMN delivers MSRPDUs, along with information about the originating interface, to the MSRP service, as illustrated in Fig. 13. 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 invokes 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.

### C. 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.

### D. 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 being 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 network 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 TGs 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 A/V TRANSPORT PROTOCOL

### A. Introduction

The Audio/Video Transport Protocol (AVTP) is defined by IEEE Std. 1722-2011, Layer-2 Transport Protocol for Time-Sensitive Applications in a Bridged Local Area Network [5], and was designed specifically to take advantage of the new capabilities added to 802 networking by the 802.1 AVB TG. When AVB was nearing completion, there was no A/V protocol that was directly suitable for use on AVB networks so AVTP was created.[4]

### B. 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 A/V formats where possible;
- maintain A/V coherence regardless of network topology;
- support multiple media clocks;
- wire replacement.

Design decisions for AVTP came from the above goals. AVTP was never designed to transport A/V across the country; instead, it was optimized for individual-venue-sized installations where a venue could be anything from a small concert theater or playhouse up to a stadium or large outdoor amphitheater.

By keeping AVTP simple and reusing existing well-known A/V formats, it was possible to maximize multivendor interoperability. It was critical to the success of this technology to keep it simple enough to be used by a garage band, and yet be flexible enough to fill the needs for a large concert hall.

---

[4] Similar work is currently underway in the IETF and elsewhere to bring the benefits of AVB to the Real-Time Protocol (RTP). For more information, see [7].

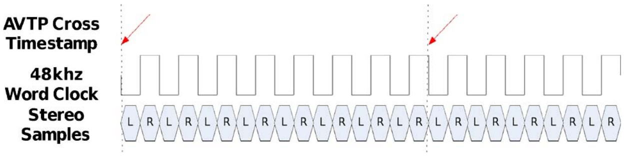

**Fig. 14.** ***AVTP sample timestamps.***

### C. AVTP Basic Concepts

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

*1) Data Formatting:* A fundamental requirement for transporting media is standardizing media formats in a packet. Since so much work has previously gone into this problem, there is no need to reinvent the wheel, so AVTP makes use of the IEC 61883 A/V formats, which have been used for years on IEEE 1394 (Firewire). IEC 61883 defines a rich set of media content encodings including simple mono or multichannel audio, encrypted surround sound, low-resolution raw video, high-bandwidth compressed video streams, and even musical instrument digital interface (www.midi.org) data.

*2) Media Clock Reconstruction:* In order to maintain real-time performance, it is critical that the source and sink of A/V data maintain synchronized media clocks. This eliminates the need for sample rate conversion and greatly reduces 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 48 kHz and 44.1 kHz 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, as illustrated in Fig. 14.

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.

*3) Presentation Time/Latency Normalization:* Another key concept of AVTP is the "presentation time": this normalizes 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 A/V samples to be simultaneously presented to media interfaces regardless of the number of network hops between the source and the sink.

AVTP has a default presentation time of 2 ms which allows most networks to operate with real-time performance and without unreasonably limiting network topology. The presentation time offset can, however, be adjusted to accommodate either extremely low latency or unusual network topologies. If network latency lower than 2 ms 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.

### D. 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 A/V 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, by itself, resolve the lip sync issue, it does create a coherent system that can then be used to time align multiple A/V sources. AVTP compensates for the different delays between a source and various listeners, so it is a simple calculation to combine those delays with the fixed codec delays to achieve tight lip sync.

### E. 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 a 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 WG.

## VII. A MANAGEMENT PROTOCOL FOR AVB DEVICES: AVDECC

### A. Introduction

The Audio/Video Discovery, Enumeration, Connection Management, and Control (AVDECC) standard [6] 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 A/V systems where every component is from a single vendor. AVDECC enables multivendor system to work together seamlessly.

### B. AVDECC

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

- discovery;

- entity model;
- connection management;
- enumeration and control.

*1) 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.

*2) Entity Model:* In an A/V system, there is a need to not just discover a device but also to discover the 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 A/V device may be comprised of a network streaming port, other external ports or jacks, and internal ports. In order to intelligently manage an A/V 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 does not 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 A/V signals.

*3) Connection Management:* AVDECC connection management controls the making and breaking of connection between AVB stream sources and sinks.

*4) Enumeration and Control:* AVDECC Enumeration and Control Protocol allows AVB devices to be queried to understand their capabilities and use the capabilities. Many A/V devices that seem like single function devices are in fact multifunction devices. 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.

### C. Summary of AVDECC

By combining all the capabilities of AVDECC, a multivendor network A/V system can be managed from a single controller. All A/V 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 A/V industry.

## VIII. CONCLUSION

The package of described standards are the AVB standards—plus new layer-2 transport and control protocols—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 WiFi 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 [9], 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. ■

## Acknowledgment

Much of the material in this paper was derived from contributions made by the authors to the IEEE 802.1, 1722, and 1722.1 Working Groups as well as the AVnu Technical Working Group. The references list 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.

## 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] *IEEE Standard for Device Discovery, Connection Management and Control Protocol for IEEE 1722 Based Devices*, IEEE Std. 1722.1-2013.

[7] A. Williams, K. Gross, R. van Brandenburg, and H. Stokking, "RTP clock source signalling," May 2013. [Online]. Available: http://tools.ietf.org/html/draft-ietf-avtcore-clksrc

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

[9] M. Johas Teener, A. Huotari, Y. Kim, R. Kreifeldt, and K. Stanton, "No-excuses audio/video networking: The technology behind AVnu," AVnu Alliance White Paper, 2009.

## ABOUT THE AUTHORS

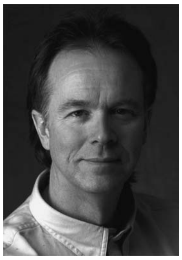

**Michael D. Johas Teener** received the B.S. degree from California Institute of Technology (Caltech), Pasadena, CA, USA, in 1971 and the M.S. degree from the University of California Los Angeles (UCLA), Los Angeles, CA, USA, in 1976.

He is currently a Senior Technical Director at Broadcom Corporation, San Jose, CA, USA, with major responsibilities for time-sensitive and high-performance/low-cost network technologies. In the past, he was the Plumbing Architect at Apple, Inc., where he was the Chief Architect of Apple's Firewire (IEEE 1394) technology and a primary contributor to both the development and standardization of multiple generations of 1394 technology. His early career was spent designing and architecting advanced radar and sonar equipment and early distributed digital PBX systems. He is currently the Chair of the IEEE 802.1 Time-Sensitive Networking Task Group, and the Co-Chair of the IEEE 1722.1 Working Group. He holds numerous patents related to telecom, FireWire, and time sensitive networking.

Mr. Johas Teener received the 2009 IEEE Standards Medallion for "exceptional leadership, clarity of vision, and perseverance in the effort to develop and standardize IEEE 802.1 Audio/Video Bridging."

**Andre N. Fredette** is a Senior Principal Engineer at Extreme Networks, Research Triangle Park, NC, USA. He has a long history in designing and developing networking infrastructure equipment and standards. He is the Lead Architect for AVB products at Extreme Networks, and has been active in the AVB/TSN Task Group in the IEEE 802.1 Working Group.

**Christian Boiger**, photograph and biography not available at the time of publication.

**Philippe Klein** received the Ph.D. degree in nuclear physics instrumentation from the Louis Pasteur University, Strasbourg, France.

He was a Research Fellow at the European Center for Nuclear Research (CERN), Geneva, Switzerland, and a Consulting Engineer at Digital Equipment (DEC). He is currently an Associate Technical Director with the Broadband Technology Group, Broadcom Corporation, Airport City, Israel. His areas of interest include security, power management, and broadband networks with an emphasis on hybrid and converged home networks.

**Craig Gunther** is a Senior Principal Engineer in the Corporate Technology Group, Harman International, Sandy, UT, USA. He is the editor for the Stream Reservation Protocol (SRP), clause 35 of IEEE Std 802.1Q-2011, and contributor to other Audio/Video Bridging (AVB) standard. He is currently working with the Time-Sensitive Networking Group of the IEEE 802.1 Working Group on the next generation of AVB. Prior to his work at Harman International, he was a Distinguished Member of Technical Staff with Lucent Technologies, and earlier with Ascend Communications, where he worked on high-density remote access servers that were instrumental to the Internet's growth during the late 1990s.

**David Olsen** is a member of the Corporate Technology Group within Harman International, Sandy, UT, USA, responsible for research into new network technologies. He is the Chair of the IEEE 1722a Audio Video Transport workgroup. He was the editor of IEEE Std. 1722-2011 and deeply involved in the development of the 802.1 AVB Standards.

**Kevin Stanton** has a Ph.D. degree in electrical and computer engineering.

He is a Principal Engineer at Intel Corporation, Hillsboro, OR, USA, with responsibility for accurate time synchronization technologies and its applications in high-volume devices. He was an active contributor to the IEEE 802.1 AVB standards, leading the expansion of accurate time to 802.11/WiFi technologies.