White Paper

Synopsys
Predictable Success

Ethernet Quality of Service

New IEEE Specifications Driving a New Generation of Network Products

June 2010

Author

Lokesh Kabra
Senior R&D Manager, Synopsys Inc

John A. Swanson
Senior Staff Marketing Manager, Synopsys Inc

Abstract

As designers look to their next-generation network designs, they are faced with a set of new challenges when developing products that incorporate the common Ethernet interface. Not only is Ethernet continuing to grow and be driven into faster and faster network communications, as seen by the move to 40G and 100G standards, but it is also being updated to support streaming audio and video applications over Local Area Networks (LANs). In 2008 and 2009, we saw several IEEE specification updates either ratified or advanced to solid draft versions that are targeted at improving networking systems' Quality of Service (QoS). To maintain Ethernet as a dominate and long-lasting network interface, incorporating these updates into new consumer products will be critical to achieve the throughput and meet the demands of the consumer. The work being done by the 802.1 Audio Video Bridging (AVB) work groups not only applies to Ethernet but to all 802-compliant networks, and will give consumers a range of options to support streaming audio and video in their home and automotive networks.

Audio Video Bridging

Most of today's entertainment media is available in digital form, and soon perhaps all will be. Applications supporting streaming audio and video have resorted to proprietary formats, making support of the network more complex. Yet it's the audio video streaming, along with interactive applications such as multi-player and interactive gaming, which has driven the enhancements made to the bridged LANs required to make these types of online activities possible. These improvements have led to the standards that not only address consumer requirements, but will also be used in professional audio video setups in mixed wired and wireless networks. Prior to the development of the new IEEE Audio Video standards, there was no interoperability among bridges that supported audio and video streaming, let alone any standard way of setting up or maintaining such a service. To succeed in moving from the tech-savvy customer to the mass market, these standards are required and the work being done by the IEEE appears to be fully supported by companies with existing solutions as well as plans for new products supporting audio video streaming in 802-compliant networks.

The IEEE updates that will drive a new generation of home networking products include the new or updated specifications related to Audio Video Bridging (AVB) including: IEEE 802.1-AS, version D6.x and IEEE 802.1-Qav, which was released in January 2010 as IEEE 802.1Qav-2009 for Audio Video (AV) traffic; IEEE 802.1Qat, version D6.0 as well as the IEEE 1588-2008 PTP specification approved in 2008, which is required as a hardware implementation to support IEEE 802.1-AS in Ethernet devices.

Other key updates related to AVB in an 802-compliant LAN are in early draft versions including: IEEE802.1BA version D1.1 which defines AVB systems; IEEE 1722 Version D2.2 which defines the AVL L2 transport layer protocol including the AVB packet structure; and IEEE 1733 version D6.0 defining the Layer 3 transport for AVB.

Why are these specifications critical to future home and LAN applications? Why will they drive a new generation of networking consumer products including Ethernet products? As a simple example, consider the evolution of different configurations in a home audio video system. The first type of configurations is what is called a single source single listener. This configuration is shown in figure 1.

Figure 1: Single source single listener

In this simple configuration, no AVB is required as the television can only see the single stream provided and separates the audio and video. A more typical set-up today, known as single source multiple listeners, is shown in figure 2.

Figure 2: Single source multiple listeners

In this configuration the speakers, or listeners, have digital settings to set their location; each speaker sees the entire stream and extracts the appropriate data for its location. Historically, Ethernet has not been used in this type of application, but with the new AVB standards the door is open for Ethernet devices to expand into this market in home, automotive and other applications. Why Ethernet for this type of application? Simply, it is so a standard cable can be used for all connections, removing the miss-mash of cables used today.

In a more complex set-up, becoming more and more common in today's home networks, is what is called multiple source multiple listeners. This configuration is shown in figure 3.

Figure 3: Multiple sources multiple listeners

Now here there are multiple audio and video streams. Users are able to select from a programmable "decoder", which streams media data to different components so consumers can watch and listen to audio and video anyplace in their home network. In this scenario, enabling efficient and reliable connection of the network components is the goal of the audio video bridging working groups.

The new IEEE 802.1 specifications enable this synchronized home network to support all of the various types of digital media being provided by external network sources as well as internal network streaming. Consider a typical digital home. It receives Internet access from a land-line or cable provider, with TV possibly coming from the same cable provider or perhaps from a satellite provider. Then you add in satellite radio, wireless internet, etc. – the possibilities are endless. Consumers have a range of choices and want to build their digital home to synchronize and enjoy this new age of media. Figure 4 shows today's tech-savvy consumer's digital home, and most likely the home of tomorrow's mass-market consumer.

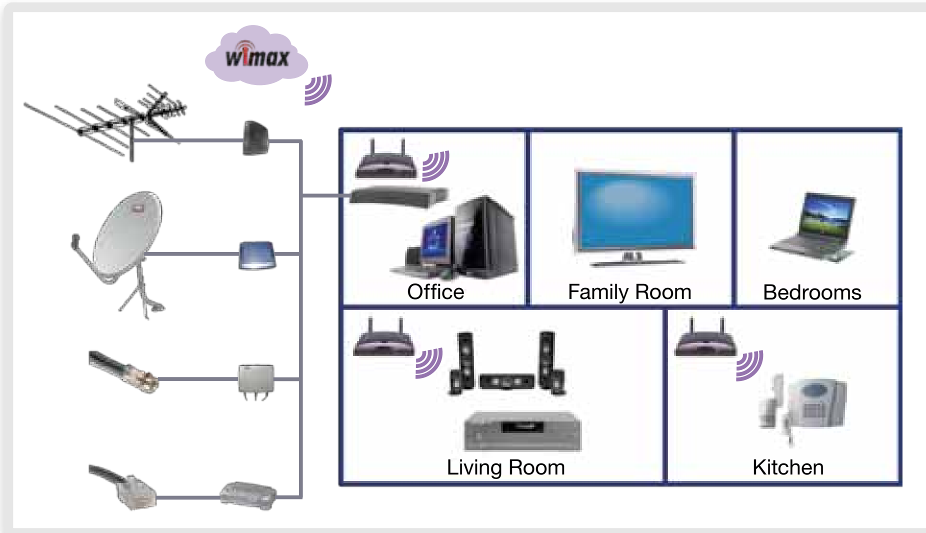

Figure 4: The digital home

When implemented, the new AVB-compliant network's AVB-compatible clouds will be able to connect to high-speed providers not supporting AVB as well as interoperate with non-AVB compliant networks and network components. While AVB traffic will not be able to pass through a non-AVB compliant bridge, non-AVB traffic will be able to pass through an AVB-compliant bridge. When put together, you can easily build an AVB LAN into your home and connect to the external internet as shown in figure 5.

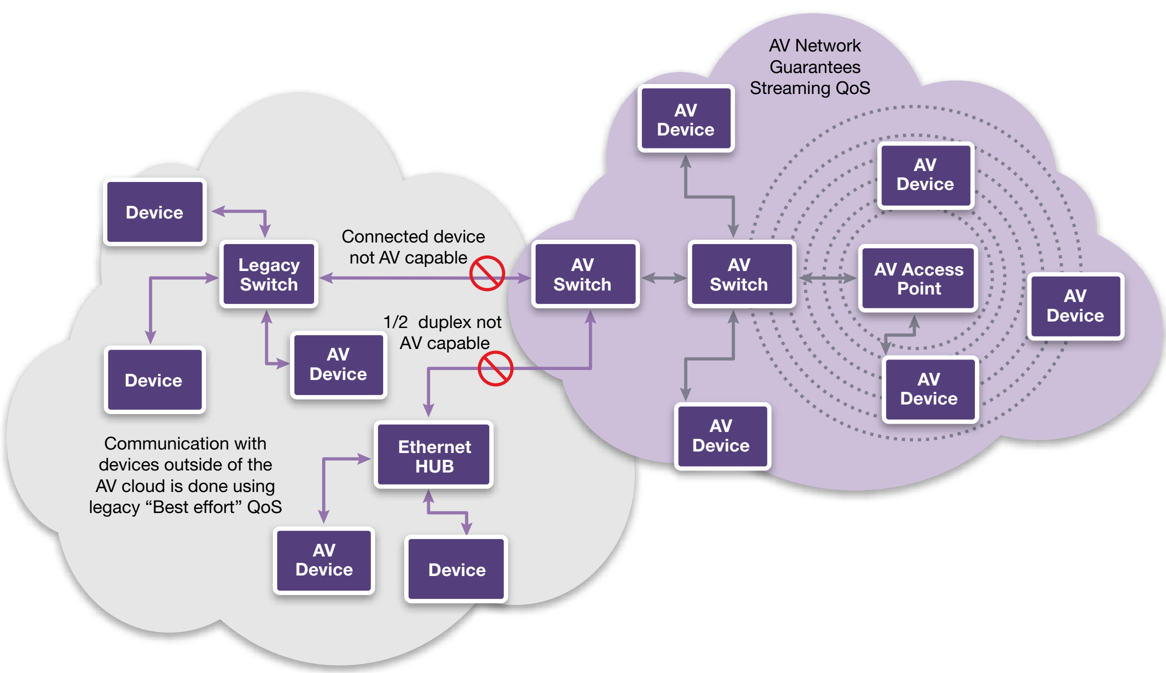

Figure 5: AVB network connected to legacy network

New IEEE Specifications Enabling Audio Video Bridging Today

The AVB updates to the IEEE 802.1 standards which are responsible for the higher level services in 802-compliant networks have advanced to implementable versions of the draft specifications in 2009. The IEEE 802.1, working groups developed updates to the specifications that are related to the architecture, link level security, inter-networking and network bridging, as well as Layer 2 management of 802-compliant networks. The AVB updates are being undertaken to support multimedia networking adding QoS features to support streaming traffic. In an AVB network, all of the components should support the AVB specifications and adhere to the following:

- No transparent or non-AVB devices are allowed in the AVB cloud
- AVB only supports full duplex links
- Bandwidth reservation is made for end-to-end connections for AV traffic before actual AVB data transfer
- AVB end stations negotiate the AVB capabilities using LLDP
- Maximum of 75% / 25% ratio for AVB / non-AVB traffic allowed at any AVB node
- The total number of hops or bridges allowed in an AVB network are limited by the latency requirements of the AVB data and traffic
- Traffic scheduling is required for isochronous traffic
- Latency delay for isochronous traffic within each intermediate node is limited to 250us at 100Mbps and 125us for 1Gbps links

The AVB specifications are being developed by several different working groups. These groups include the IEEE 802.1-AS working group, IEEE 802.1Qav working group and the IEEE 802.1Qas working group. In 2009 draft versions of these new specifications were released that are solid enough for companies to begin to build a new generation of products supporting real audio video broadcasting in a Local Area Network (LAN). In addition to the IEEE 802.1 AVB activity the IEEE 1588-2008 (also known as PTP Version 2) specification is an integral part of a quality audio video network solution and is referenced by the IEEE 802.1AS working group. Each of these IEEE 802.1 working groups as well as the released IEEE 1588-2008 specification is summarized in the following sub-sections. First let's take a more detailed look at how the IEEE 1588 PTP works as timing is one of the most critical components of a QOS networking component.

IEEE 1588-2008 Precision Time Protocol (PTP)

Although not part of the actual AVB working groups, the IEEE 1588 PTP standard is part of the solution enabling real audio video traffic in a network. The PTP protocol is used to establish a master-slave time relationship in a network, ensuring that all the devices on the network gets synchronized to the common "time" maintained by the master. Using this protocol, the master and slaves exchange messages that determine the offset between the clocks of the slaves and the master on the network. These exchanges reduce the impact of the clock drifting between the master and slaves on the network.

The PTP is transported over UDP/IP as per the original IEEE 1588 specifications published in 2002. The system or network is classified into master and slave nodes for distributing the timing/clock information. Figure 6 shows the process that PTP uses for synchronizing a slave node to a master node by exchanging PTP messages.

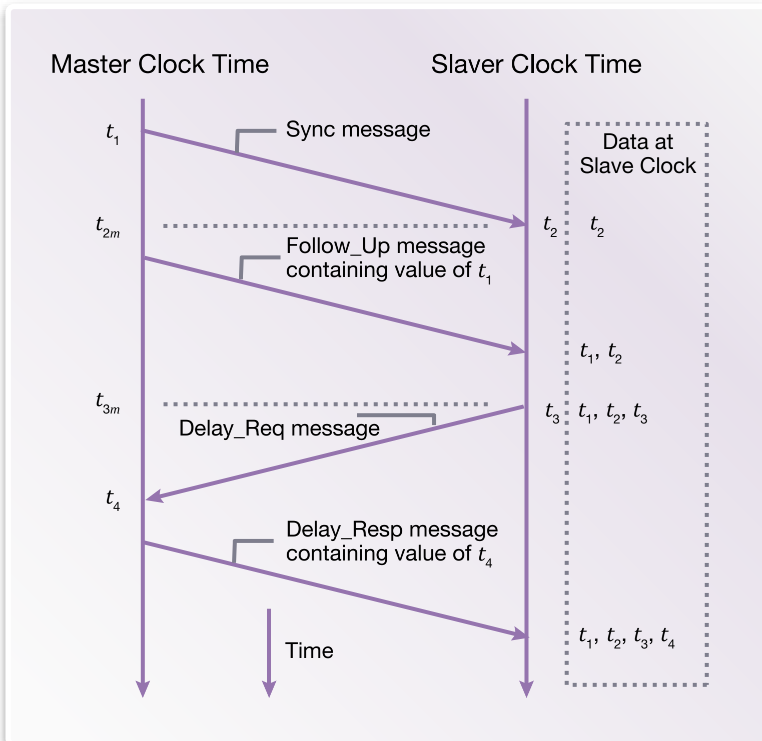

Figure 6: PTP synchronization between master and slave nodes

As shown in Figure 6, the PTP uses the following process:

1. The master broadcasts the PTP Sync messages to all its nodes. The Sync message contains the master's reference time information. The time at which this message leaves the masters system is t1. This time must be captured, for Ethernet ports, at GMII/MII interface
2. The slave receives the Sync message and also captures the exact time, t2, using its timing reference
3. The master sends a follow up message to the slave, which contains t1 information for later use
4. The slave sends a "Delay_Req" message to the master, noting the exact time, t3, at which this frame leaves the GMII/MII interface
5. The master receives the message, capturing the exact time, t4, at which it enters its system
6. The master sends the t4 information to the slave in the "Delay_Resp" message
7. The slave uses the four values of t1, t2, t3, and t4 to synchronize its local timing reference to the master's timing reference

Most of the PTP implementation is done in the software above the UDP layer. However, the hardware support is required to capture the exact time when specific PTP packets enter or leave the Ethernet port at the GMII/MII interface. This timing information must be captured and returned to the software for the proper implementation of PTP with high accuracy.

Originally released as IEEE 1588-2002, the updated IEEE 1588-2008 (or Version 2) specification increases the precision of the timing by adding several enhancements to the original specification. These enhancements included:

- Transparent clocks throughout a IEEE 1588 network which enables to expand the span of PTP networks over LANs
- Reducing the size of the synchronization messages
- Support for higher precision and faster updates
- An option for layer 2 transport
- Fault tolerance
- Quick reconfiguration responses to changes in the IEEE 1588 aware network

It is important to note that the IEEE 1588-2008 specification is not interoperable with the IEEE 1588-2002 specification. This will drive consumer products developed using the IEEE 1588-2002 specification and that require this higher precision timing to be updated to support the new IEEE 1588-2008 specification. This is because the IEEE 1588-2008 version supports Peer-to-Peer PTP (Pdelay) messages in addition to SYNC, Delay Request, Follow-up, and Delay Response messages. Figure 7 shows the method to calculate the propagation delay in clocks supporting peer-to-peer path correction.

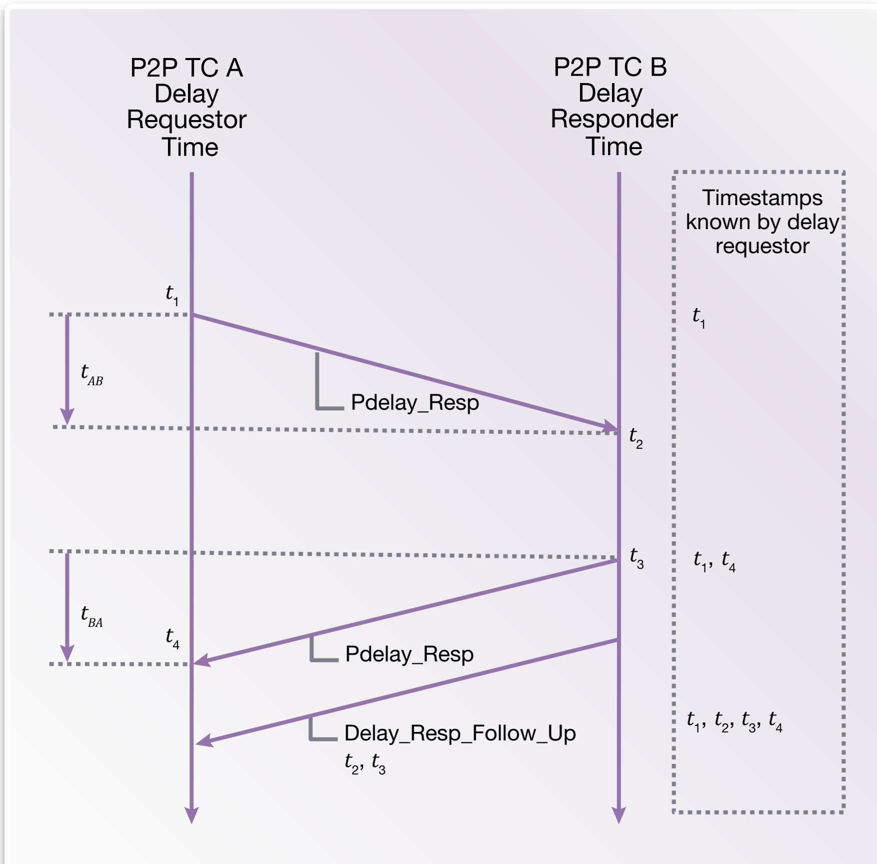

Figure 7: Propagation delay calculation in clocks supporting peer-to-peer path correction

As shown in Figure 7, the propagation delay is calculated in the following way:

1. Port-1 issues a Pdelay_Req message and generates a timestamp, t1, for the Pdelay_Req message
2. Port-2 receives the Pdelay_Req message and generates a timestamp, t2, for this message
3. Port-2 returns a Pdelay_Resp message and generates a timestamp, t3, for this message

   To minimize errors because of any frequency offset between the two ports, Port-2 returns the Pdelay_Resp message as quickly as possible after the receipt of the Pdelay_Req message. The Port-2 returns any one of the following:

   - The difference between the timestamps t2 and t3 in the Pdelay_Resp message
   - The difference between the timestamps t2 and t3 in the Pdelay_Resp_Follow_Up message
   - The timestamps t2 and t3 in the Pdelay_Resp and Pdelay_Resp_Follow_Up messages respectively

4. Port-1 generates a timestamp, t4, on receiving the Pdelay_Resp message
5. Port-1 uses all four timestamps to compute the mean link delay

Although originally targeted for industrial control systems, updates to the IEEE 1588-2008 specification include the timing synchronization required for the high-speed digital home as well as hardware implementation required to support the new Audio Video Bridging specifications in Networked devices.

IEEE 802.1AS – Timing and Synchronization

The full name of this working group, or Project Authorization Request (PAR) as referred to in the IEEE, is "Standard for Local and Metropolitan Area Networks - Timing and Synchronization for Time-Sensitive Applications in Bridged Local Area Networks". The goal of this working group is to specify the protocol and procedures that are used to synchronize components in a network running time sensitive applications. These include such common consumer applications where audio and video data is transferred across one or more components in a Local Area Network (LAN) with fixed and symmetrical transmission delays. An example of this for IEEE 802.3 Ethernet is a full duplex link.

The purpose of time synchronization is to provide the common 125μs cycle throughout the AVB cloud as well as a common time base so that the AVB source and destinations have a sense of how sampling and receiving times are related. This standard will ensure normal operation of the time sensitive network when adding, removing, and in the case of component failure, during network reconfiguration. This specification requires a hardware implementation of the IEEE 1588-2008 specification, and enables network stations attached to bridged LANs to meet the required jitter, wander and time synchronization needed in time sensitive network applications. It does not include synchronization to an externally provided timing signal such as UTC. Currently this specification is in version Draft V7.0 and appears to be close to final, with companies including these capabilities into their next-generation of Ethernet products targeted at home multi-media applications today. It will also be necessary that the hardware of Ethernet ports in the network has to be enhanced to support this specification along with the support for PTP stack in software.

IEEE 802.1Qav – Forwarding and Queuing Enhancements for Time Sensitive Streams

The full name of this working group, or PAR is "IEEE Standard for Local and Metropolitan Area Networks – Virtual Bridged Local Area Networks – Amendment: Forwarding and Queuing Enhancements for Time Sensitive Streams". What this boils down to is a new specification to guarantee that time-sensitive and loss-sensitive audio video data arrives at its destination correctly in a mixed wired / wireless LAN. This specification cites the priority ingress metering, priority regenerations and timing aware queue draining algorithms. It is being built on the timing derived from the IEEE 802.1AS. It will support both AV traffic as well as other bridged traffic on wired and wireless LAN. As bridges are increasingly added to home networks to add multimedia features, this specification will define the enhancements needed to bridge and guarantee correct transmission and reception to the audio video traffic with the required quality of service It will also be necessary for the hardware of the Ethernet port has to be enhanced to support this specification.

IEEE 802.1Qat – Streaming Reservation Protocol

The full name of this working group or PAR is "Standard for Local and Metropolitan Area Networks - Virtual Bridged Local Area Networks - Amendment: 9: Stream Reservation Protocol (SRP)". This standard will specify the protocols, procedures and managed objects usable by the existing higher layer mechanisms making up the AVB network. This allows network resources, like bandwidth and buffers, to be reserved for specific streams traveling across a bridged LAN. It will identify the traffic streams at a level needed for a bridge to determine what resources are required and provide a mechanism for dynamic maintenance of the resources.

This specification will allow AVB connections to be registered, enabling source-to-receiver management of the data and ensuring that the QoS needed for AVB traffic is met. A standard signaling protocol will facilitate the registration, de-registration and retention of the AVB networked components. This is a key component of the AVB specifications and is needed to ensure automatic configuration of AVB components in a LAN.

This specification will be software implemented and will not require hardware updates to the Ethernet AVB components.

IEEE 802.1 BA – Audio Video Bridging Systems

This specification, in an early draft version, gives an overview of the various components in an AVB system and how they interact with each other and combine to form an AVB system. It defines the profiles that select features, options, configurations, defaults, protocols and procedures of bridges, stations and LANs which are necessary to build networks that transport time sensitive audio and/or video data streams. It will help manufacturers of LAN equipment to select the profiles required to develop AVB-compatible LAN components. It will also help users to build a network, using those components, that does not require configuration to provide working Audio and/or Video services.

This standard describes some of the consequences and limitations for AV operation that result from the incorporation of non-AV capable devices in an AVB system and provide guidance in terms of meeting the end-to-end latency requirements for successful AVB operation. In short, it defines the conformance requirements for AVB systems, in terms of the standards to which conformance is required for the various system components and the optional features of those standards that are required to be implemented. These conformance requirements address the guaranteed delivery, end-to-end latency, and time synchronization requirements for successful AVB operation.

Other IEEE AVB Related Activities

Given the enormous market potential of 802 LAN audio video applications, it is no surprise that in addition to the specifications already discussed there is additional on-going work that will have further positive impacts on audio video LAN applications going forward. Two key groups include the IEEE 1722.1 and IEEE 1733 working groups, each of which is summarized below.

IEEE 1722 Layer 2 Transport Protocol Working Group for Time Sensitive Streams

This group is defining the AVB transport protocol (AVBTP), further building on AVB by adapting the IEEE 1394 specifications comprehensive suite of media formats, encapsulations and synchronization mechanisms for use in Ethernet AVB networks. Products using IEEE 802.1 AVB and AVBTP standards will enable the construction of highly interoperable 802 networks capable of streaming audio and video maintaining perfect QOS.

IEEE 1733 Layer 3 Transport Layer for AVB

This group is specifying the protocol, data encapsulations and connection management, as well as the presentation time procedure to ensure interoperability between the audio and video end station using standard networking protocols defined in IEEE 802 networks. This standard leverages the Real-time Transport Protocol (RTP) and IEEE 802.1 AVB protocols.

Synopsys Ethernet QOS Solution

Synopsys provides a complete Ethernet solution for the next generation of AVB compatible network devices. Building on the DesignWare Universal Ethernet core, Synopsys provides a very configurable Ethernet MAC with AVB and IEEE-1588-2008 easily selectable from a configurable menu of options. To implement an AVB design you would simply add the needed AVB configurations to your MAC design using the coreConsultant tool.

Adding the AVB blocks into your Ethernet design is a simple configuration of the Ethernet QoS core. When you select the AV Features Support option from the list of optional modules, the required IEEE 1588-2008 modules to implement the specification will automatically be enabled. The selection of audio video support is shown in Figure 8.

Figure 8: DesignWare ethernet QoS audio video configuration menu

When selecting AV you can select none, one or two additional channels into your Ethernet configuration to be dedicated to AV traffic separately for both transmit and receive. This way you can optimize the implementation of your design as per specific application. An example would be if you are designing a system-on-chip (SoC) to support only AV receivers, you would configure the core to support no additional channels for transmit and one or two channels for receive, again, depending on your product's AV requirements. If no additional channel is selected for transmit path, the traffic management function (Fig 10) for scheduling the transmission of AV traffic as specified by 802.1-Qav is automatically removed from your AVB configuration.

Next you can select some optional or additional configurations for the IEEE 1588-2008 features. When you go to this configuration option you will see that IEEE 1588 Time Stamping and IEEE 1588 Advanced Time Stamping are already selected and cannot be unselected. This is because these features are required to support the AV features previously selected. This is a feature of the DesignWare Ethernet QoS core that allows designers to focus on their unique design requirements and not have to worry about selecting every model required when another module has been enabled. Figure 9 shows the IEEE 1588 Time Stamping configuration block. You can select to support external time stamping or enabling the IEEE 1588 higher word register in the Ethernet QOS core.

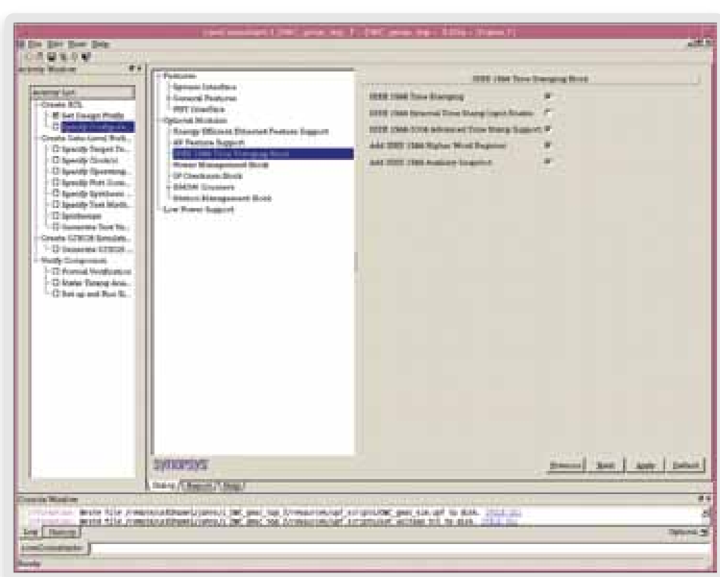

Figure 9: DesignWare ethernet QOS IEEE 1588-2008 configuration menu

The Ethernet QoS is now configured to include the required logic to support AV and IEEE 1588-2008 Time Stamping. Figure 10 shows a block level implementation of the DesignWare Ethernet QoS core for AV traffic with 2 additional AV channels in both transmitter and receiver.

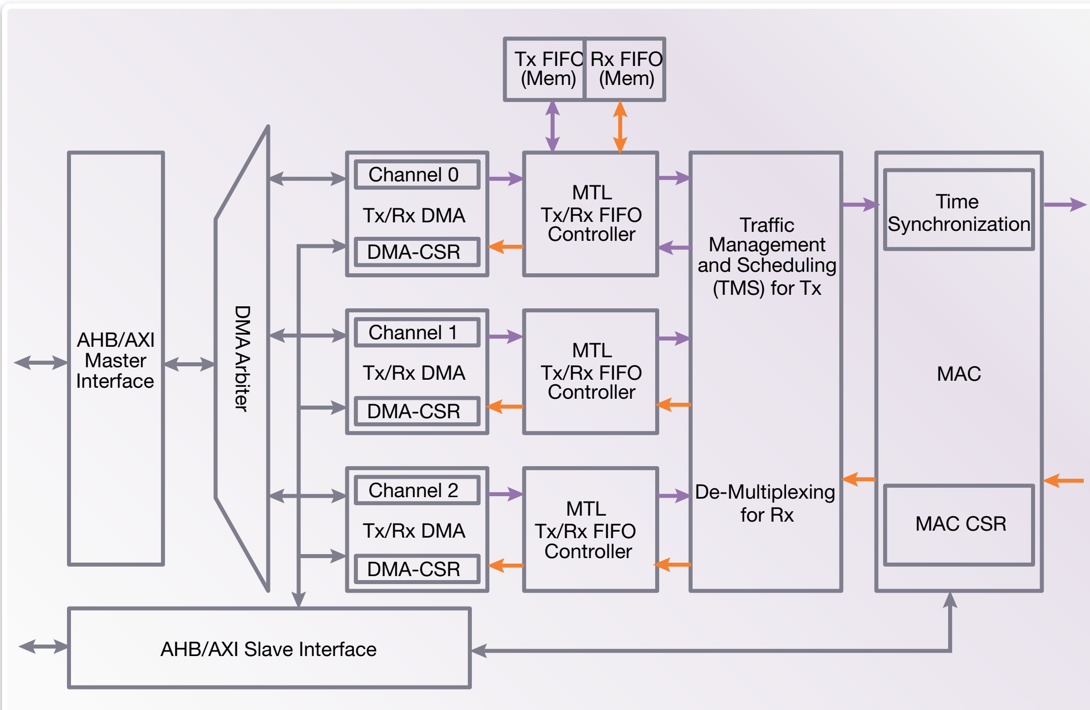

Figure 10: Top level block diagram of DesignWare ethernet QOS for AVB

After completing the rest of your designs configuration options, such as power savings, PHY interface, checksum offload options, etc., you simply click on the apply button shown at the bottom right of Figures 8 and figure 9 to generate the RTL targeted for your Ethernet QoS implementation. The RTL is now optimized for your selected configuration and only the logic needed to implement your configuration is included ensuring an efficient RTL implementation of your Ethernet AVB design. You can now integrate your design into your SoC.

Summary

Never bet against Ethernet! As a long lasting and proven technology, these updates to the widely used Ethernet protocol will drive a new generation of designs that will simplify cabling and network support. These designs will not just be in consumer products either, but will also migrate into carrier networks and other markets like automotive. Synopsys will continue to provide updates to its Ethernet product line, meeting customer requirements for this and future evolutions of Ethernet.

Synopsys
Predictable Success
Synopsys, Inc. • 700 East Middlefield Road • Mountain View, CA 94043 • www.synopsys.com