

Only a few years ago, computer networks were complex beasts tended by special acolytes and running on different standards. Today they have become commonplace in many homes and offices, simply plugged together using Ethernet technology. The same revolutionary change is coming for Audio/Video (AV) networking, as AVB (Audio Video Bridging) products that run over the same network, begin to enter the market.



Xmos for AVB: Ethernet based networking for Audio/Video

Putting together networks of AV equipment for professional and consumer use, or for use in vehicles, is about to become simpler while also delivering better quality. No longer will specialist connectors and cables be needed to create a rats' nest of connectivity. Instead Audio Video Bridging (AVB), a set of international standards, will make setting up and managing networks almost as simple as just plugging together the different elements. Sound and video sources will be mixed and distributed to screens and speakers, with high quality, low latency and tight synchronization. Furthermore, the connectors and cables are cheap, widely available off-the-shelf and easy to route.

Problems with Audio/Video

Audio/Video has become more complex, particularly in the areas of professional AV in studios, concert venues, consumer AV and AV in automobiles. Two significant problems particularly stand out; the rats' nests of wiring caused by point-to-point communication and the great care required to ensure that all speakers are in phase in multi-speaker, large venues.

There have been several attempts to re-create the simplicity of today's computer networks for AV, but these have been only partly successful. Few people understand the complexities of successfully building an AV network and there are no widely accepted standards. The physical elements of a proprietary network, such as the controller chips for each element in the network, can be expensive as they are made only in small quantities and often don't meet the standards required for high quality audio. The AVB standards take widely used computer networking technologies, particularly Ethernet, and add additional capabilities to make them better suited to AV networking.

How Ethernet Works

Within Ethernet, data is transmitted between devices (such as a computer and a printer) in packets. Each packet carries one or more addresses for its destination. Like a postal packet traversing the postal system, the network has no knowledge of what is in the packet, but uses the address to pass the packet to the next point in the network.

In an Ethernet based network, each endpoint (computer, storage element, printer etc.) is identified by a unique address and has a single connection to the network, through an Ethernet switch. The switch may have multiple devices connected to it, as well as connections to other switches.

Whilst a computer previously used a dedicated link to print a document, now, with Ethernet it breaks the description of the page into packets, throws the packets into the Ethernet and gets on with its activities. Each packet is labeled with the network address of the printer and the electronic equivalent of "Packet 3 of 17". In a simple network the packets go from the computer through a switch to the printer. In a more complex network there may be more steps between computer and the printer. The higher the number of switches and the longer the distances between them, the longer the packets take to traverse the system. In a complex network different packets of the same document may be routed through different routes, arriving out of sequence at the destination. The printer reassembles the page description and prints it; if a packet is missing, the printer asks for a resend. As packets travel through the network they are interleaved with other network traffic, including computer to computer communication, emails, and so on.

Ethernet gives no guarantees; packets may not arrive, may arrive out of order or take a long time to cross the network. For printing pages of a document this doesn't matter greatly—resending a packet is not a big deal and timing measured in seconds is adequate for most computing needs. For sending AV, however, timing is crucial.

Ethernet technology

Ethernet is the dominant form of wired network, and has more or less completely replaced other versions of networking in computing and communication. It is a simple and low-cost way to build an easy-to-use basic network, yet can be scaled to serve large networks of computers, printers and other peripherals.

At first glance Ethernet might seem unsuitable for AV networks as it includes no concept of time (see *How Ethernet Works* callout). Ethernet can, however, assign priorities to packets and set up dedicated links—rather like those in an old fashioned telephone network, where a dedicated circuit was created at the beginning of a call and torn down at the end.

The new approach, Audio Visual Bridging (AVB) exploits these techniques and adds timing to make Ethernet a suitable base for creating AV networks.

AVB

AVB is a collection of standards developed by IEEE that build on Ethernet and other networking technologies to make them suitable for the demands of all applications, even high-end AV (see *How AVB works* callout). AVB is being promoted commercially by the AVnu alliance (see *AVB Bodies* callout).

With AVB, AV products attached to the Ethernet are called endpoints, which communicate with each other through specialist endpoint controllers. The endpoint controllers sit between the AV dedicated electronics and the network interfaces, handling the timing and traffic priority issues needed for AVB.

Choosing endpoint controllers

Many companies are contributing to the technology for implementing endpoint controllers. Developers must make some difficult judgements as they assess the controller options available, including ease of use, cost and how well they cope with the inevitable changes in future standards for AVB.

Using a classic dedicated chip (an ASIC) requires a specification that is tightly locked down. Once the chip has been designed and the first samples manufactured (itself an enormously expensive job with today's technology) any changes will be time consuming and expensive. Since ASICs must be

How AVB Works

Existing Ethernet cabling and Ethernet endpoints (computers etc) can continue to exist on an AVB Ethernet, but the switches and bridges responsible for routing the traffic have to be AVB compliant. Any new AV equipment added to the network must also have AVB functionality. AVB co-exists with other AV Ethernet protocols, including CobraNet, EtherSound, Audinate Dante, AES50 and Q-LAN.

AVB network endpoints are described as talkers, listeners or both. A talker might be a CD player while a listener might be a loudspeaker. A mixing desk both listens and talks. Within the endpoint is a dedicated AVB controller chip. In a talker, the controller takes the digitized audio signal through a standard interface, such as I2S, and converts it to an AVB Ethernet data stream, with timing data and the destination address. It then passes the stream to a Physical Interface Device (PHY), which moves the Ethernet out of the endpoint through a standard RJ45 connector and on to the network. At the listener the process is reversed; the signal travels from the network through the PHY to the dedicated controller chip, which unpacks the Ethernet signal and carries out the relevant timing operations before passing it to the audio circuitry.

AVB adds three new properties to Ethernet and other networks:

- The concept of time
- A reservation protocol
- Queuing and forwarding rules

AVB also adds a master clock to the network, to which all AVB elements within the network (endpoints, switches/bridges) are synchronized. Endpoints can then carry out operations simultaneously, for example, speakers can be kept in phase, or audio and video streams can be synchronized (lip-synch).

While Ethernet is packet based, AV data must be streamed—AVB provides a reservation protocol for creating streams of packets. A talker tells the network that it wants to send a stream of data to one or more listeners, and reserves a path and resources through the network for the time needed for the stream.

To make sure that this stream works, there are queuing and forwarding rules that allow priority to the stream while making facilities available to other traffic on the network.

With these techniques, the delay through the network (network latency) can be accurately predicted. Even with 100 Mbps Ethernet, delay is 2 milliseconds for traffic going through seven switches, and lower still for Gigabit Ethernet.

Synchronization of streams and of devices can be as accurate as 1 microsecond.

sold in very large quantities to recover the development costs, there is no scope for individual companies to design-in additional features that can differentiate their product in the market.

Another approach is to use FPGAs (Field Programmable Gate Arrays), general purpose chips that can be programmed to carry out specific functions and reprogrammed to cope with changes. The design process for an FPGA, however, is complex (it can be as complex as creating an ASIC), individual devices are expensive and they can consume large quantities of power.

The XMOS solution

XMOS has developed a new approach to designing electronic systems. The xCORE multicore microcontroller is an event-driven multicore architecture used as the platform for developing systems entirely in software.

xCORE multicore microcontrollers combine the best features of three normally distinct classes of semiconductor devices. The most familiar devices are microcontrollers, the cousin of the processors used in a PC but tailored for use in control applications. Microcontrollers are programmed using conventional software development tools but suffer from inflexible architectures and interfaces. A second class of device is the Digital Signal Processor (DSP), a specialized class of processors designed for dedicated processing of data streams, such as audio processing. Like processors they are relatively easy to program but, again, are inflexible in their predefined hardware capabilities. Both processors and DSPs can be costly, but not as expensive as the third class of devices, Field Programmable Gate Arrays (FPGAs), which are the reverse of processors in that the underlying architecture is very flexible. An FPGA can fulfil a multiplicity of applications, but the design process is complex (often very complex with larger and more powerful FPGAs) and unpredictable.

The xCORE combines the underlying flexibility of FPGA with the simple design process of the processors, but at a much lower price than any of these.

Since an xCORE is entirely programmed in a C-based high-level programming language, any device can be re-configured during the development process, or even when deployed

in the field. This is perfect for building devices when a specification is evolving, or where there is a requirement to meet a standard in the final stages of development, just as the situation is with AVB.

This flexibility and low price means that xCORE devices are already being used as the basis of a number of AVB designs from some of the leading AV companies.

AVB Bodies

Two bodies are deeply involved in Audio Video Bridging (AVB). The technical standards, which define how AVB actually works, are being created within the IEEE by the IEEE 802.1 Audio/Video Bridging Task Group. The work required to turn the standards into commercial products is being driven by the AVnu alliance.

The AV Bridging Task Group

Like other IEEE standards activities, the job of the task group is to create a set of standards which are then implemented in the commercial sector. Members of the Task Group are drawn from a range of organizations, but within the IEEE environment they are expected to act as objective individuals. While the Task Group is working on AV Bridging as a technology independent of the underlying network, the most attention has been on Ethernet, as defined by IEEE 802.3 documents. There is also work on AVB for 802.11 (Wireless LANs or Wi-Fi).

The IEEE standards most directly used within AVB are IEEE 802.1BA Audio Video Bridging Systems, IEEE 802.1AS Precision Timing Protocol, IEEE 802.1Qat Stream Reservation Protocol, IEEE 802.1Qav Queuing and Forwarding Protocol, as well as IEEE P1722 and IEEE P1733.

AVnu Alliance

The AVnu Alliance is a group of leaders in silicon chips and high-quality professional/consumer audio, who promote Ethernet AVB. Currently it is concentrating on the automotive, professional and consumer electronics markets.

Members of AVnu are developing the tests and procedures that demonstrate that different devices work together. They organize plug-fests, where members bring AVB equipped equipment and demonstrate that it works with other manufacturers' products.

The AVB Endpoint Reference Design

Developers building an AV product to attach to an AVB network generally want to concentrate on the product. They require a controller that is low-cost, flexible and easy to implement, rather than a device that requires them to engage with the details of AVB standards.

To meet these requirements, XMOS and the design and engineering company Attero Tech have created a reference design for an AVB audio endpoint.

The reference design uses a low-cost XMOS XS1-L16 device as the basis for running a software-only implementation of an AVB audio endpoint, capable of both talker and listener modes, running up to 8 duplex channels of audio. It integrates the full range of AVB protocols, digital audio interfaces and control software, in a single board with a range of

analog and digital I/O. Developers can buy a kit with two boards, two Ethernet cables, power supplies and a xTAG-2 adapter for JTAG debugging.

The software included in the reference design means that developers don't need to have detailed knowledge of Ethernet or the AVB standards to create a product that meets these standards and operates efficiently—with the reference design, a developer can rapidly and easily incorporate an audio endpoint into a product, while retaining the flexibility to change the firmware to incorporate any changes in the hardware or software specification. Refining the design through firmware changes provides the speed and predictability of an implementation that has been proven both through plug-fests and field experience.

More about XMOS and AVB

XMOS is a young fabless semiconductor company based in Bristol, UK. The wide range of xCORE multicore microcontrollers with Hardware-Response provide a unique combination of parallel multitasking, deterministic performance, support from a large library of software-defined peripherals, allowing you to meet your exact design requirements. This makes xCORE multicore microcontrollers ideal for demanding embedded applications in audio, automotive, consumer and industrial products, where other microcontrollers struggle. XMOS provides xTIMEcomposer Studio, an easy to use development system, that makes it simple to design complex embedded systems.

- XMOS Website <http://www.xmos.com>
- Attero Tech Website <http://www.atterotech.com>
- AVnu Alliance: <http://www.avnu.org/>