

Digital Audio Distribution Systems

Kevin Gross, Cirrus Logic, Inc.

Introduction

Routing and distribution of audio in a large facility can be a clumsy and expensive undertaking. This paper highlights the advantages of distributing audio digitally and examines the suitability of several computer networking technologies to the task of carrying real-time digital audio. The introduction of economical 100Mbit Ethernet standards makes the prospect for real-time audio distribution over computer networks more enticing as it provides the raw bandwidth to distribute upwards of 100 channels of full fidelity audio on a single cable.

Audio Distribution: Analog vs. Digital

Facility wide audio distribution has, up until recently, required numerous home runs of balanced analog signals to a central signal processing and routing system. The amount of cabling involved can be onerous. To minimize induced noise and cross-talk, extra care must be taken to segregate audio runs from other signals. The cable used for audio is generally of an expensive shielded twisted pair (STP) variety. With the introduction of the AES3 digital audio transmission standard it has become possible to distribute audio digitally. Distributing audio digitally has several advantages over an analog approach.

First, analog runs are not immune to noise. Noise, at any level, picked up along the run is added to the signal. Noise picked up in a digital run affects the signal only if its level exceeds a certain threshold. Furthermore, digital signals can be readily run over fiber which provides complete electrical isolation and noise immunity. The noise immunity advantage of digital transmission imparts a cost as well as a quality benefit. In most cases, digital runs can be carried on inexpensive unshielded twisted pair (UTP) cabling.

Second, digital data may be multiplexed more cost effectively than analog data. The utility of this is in allowing multiple discrete channels of audio to be carried over the same physical cable conductors. The savings in cabling costs, conduits, support and maintenance can be enormous.

Third, a digital transmission system operates as a closed loop. In a digital system, through use of error detection and correction codes, it is possible to guarantee or verify the integrity of received data before it is used. Closed loop analog transmission systems are limited in functionality and deployment. In an analog system, if noise is injected or a connection broken the resulting noise or silence is typically interpreted as valid signal at the receiving end.

Finally, there is an undeniable digital evolution in audio systems. As a larger portion of an audio system's end devices are digitally based, it becomes more attractive to perform the audio distribution digitally. Additional conversions between analog and digital

domains result in unnecessary expense, degradation in signal quality and emotional strife for the audio system designer.

The AES3 digital audio transmission standard modestly realizes all of these advantages. Two channels of full fidelity audio, as well as a modest amount of control traffic, can be carried distances of over 100 meters without any unreported degradation in signal quality.

Audio Networking

The AES3 transmission standard, being unidirectional, provides only a single access service. AES3 does not allow for a multiple access network topology wherein multiple sources can share the same transmission media.

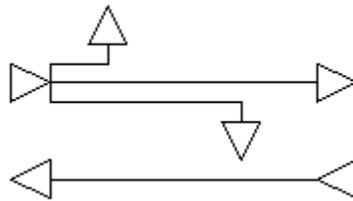


Figure 1: *Single access transmission systems like AES3 allow only a single source per segment.*

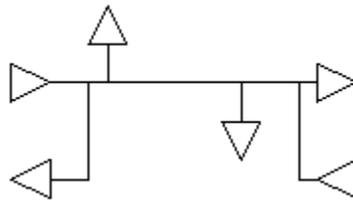


Figure 2: **Same distribution system as shown in Figure 1 implemented with a multiple access network. A multiple access network allows multiple sources per segment.**

In the context of audio routing and distribution, a multiple access network offers several important benefits over a single access system.

First, a network has intrinsic routing capabilities. A channel sourced onto the network may be broadcasted throughout the network. Individual receivers can dynamically select the channel(s) to which they wish to listen. Routing is thus accomplished in a distributed manner on the network by the individual receivers. For applications which desire central control of routing functions, the individual receivers may be configured remotely over the network from a central station giving the illusion of centralized routing. In eliminating the central router, distributed routing eliminates a potential single point of failure for the audio system.

Second, a multiple access audio network reduces total length and cost of cabling for an audio distribution system. Distributed routing eliminates home runs to a central routing

device. Furthermore, since a network is a bi-directional transmission system, in cases where sinks and sources are co-located, a single cable run can serve both devices. [Figure 2](#) demonstrates this savings.

Third, once a network infrastructure of sufficient bandwidth is installed in a facility the distribution system may be configured and reconfigured without needing to be rewired. At any network drop, audio may be sourced or sunk onto or off of the network for distribution to or from any other network drop(s) in the facility. To accommodate unexpected expansion or use, unterminated network drops may be sprinkled liberally throughout a facility at a minimum of cost. This is common practice when computer networks are installed. The network approach reduces the amount of tedious effort required to design an audio distribution system. The designer does not need to immediately specify the channel count sunk or sourced at each drop. The flexibility of the network approach also allows for agile response to changes, foreseen and unforeseen, in a facility's audio distribution requirements. This makes the network approach particularly attractive for dynamic installations such as convention centers, multiple use arenas and airports.

Fourth, since a network provides a communications channel to and from each station, all network devices may be monitored and managed from a central point. In a single access system it is only possible for the transmitter to control the attached receiver(s) but, since a single access system is unidirectional, it is impossible to monitor the receiver(s). In fact, it is impossible to determine whether the connection to a given receiver has been broken. For this reason, single access systems are often augmented with a separate control network which is used to control and monitor end stations on the distribution system. On a multiple access network, this functionality is built in.

Finally, because all stations physically share the media on a multiple access network, some applications such as intercom and paging, can gain economy by sharing media bandwidth. In such applications, source devices need only transmit while a mic key is held or audio above a certain threshold is detected at the input. The bandwidth requirements for the network media is then limited to the number of concurrently transmitting sources on the network. Telephone companies have been realizing this economy since the first switchboards were installed.

Network Overview

An audio network should deliver multiple channels of audio data with minimal delay and without error over a variety of physical media. The network should also be able to carry control data. A purpose built digital audio distribution network engineered from the ground up has the potential of delivering a system which meets the specific requirements of audio distribution to a tee. Such purpose built systems are available. One such system could even be considered a multiple access network. Purpose built systems have the disadvantage of being sole sourced. The development cost for these systems is high and, for a product targeted to the professional audio industry, the volumes are relatively low. These systems are pricey and inclined to stay that way.

An alternate approach is to use existing computer network technology. An obvious advantage of this approach is economic. The market for computer networking arguably dwarfs the entire professional audio market. Due to the high volumes, off the shelf computer networking technology is much more affordable than any purpose built audio hardware. Additional advantages to using an established network technology include the availability of components from multiple sources and a general familiarity, acceptance and knowledge of the technology. The disadvantage is that computer networking technology is optimized for computer applications. An audio application as described has markedly different requirements when compared to computer applications. There are technical problems in applying computer networking technology to audio applications. Specifically, an audio distribution network, in addition to the audio data, needs to distribute a sample clock. Most computer networking technologies are essentially asynchronous in that each station uses its own local clock when transmitting. Receivers phase lock to the received data on a packet by packet basis. And, for successful real-time audio distribution, transmitting stations need timely and deterministic network access. Some network technologies yield solutions to these requirements more readily than others. Fortunately there are quite a number of computer networking technologies to choose from. Choosing the most appropriate technology will have a profound effect on the success of any digital audio distribution system based on computer networking technology.

Token Ring is a 4/16Mbit multiple access network. This network is electrically wired as a ring. Data on its way from one station to another is received and retransmitted by all intervening stations. Media access on Token Ring is controlled by circulating a "token" on the network. A station can not begin transmission until it receives the token. Receiving the token, the station issues its transmission and passes the token to the next station on the ring. The token passing network access scheme provides for deterministic network access and high potential bandwidth utilization. A pitfall of token passing is experienced if a station goes off-line before passing the token. In which case, the token is lost and must be regenerated in a process that interrupts data communications on the network.

FDDI is a fiber optic and/or UTP based network which is connected in a ring topology and like Token Ring, uses a token passing scheme to control network access. Through the use of intelligent (though expensive) hubs, the shortcomings of a token passing network have been overcome. FDDI is a high bandwidth network operating at 100Mbit/second. FDDI has basic provisions which allow it to carry isochronous data streams. Despite early claims that FDDI would reach the desktop, FDDI's high cost and the late arrival of a UTP implementation have limited its deployment primarily to network backbone applications.

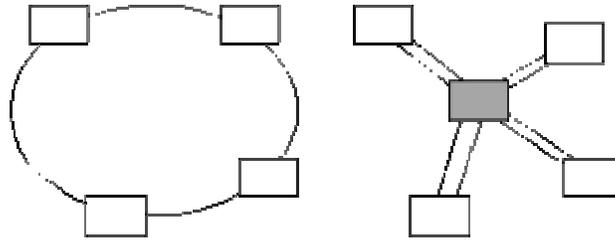


Figure 3: *Token Ring and FDDI use a ring topology to interconnect end stations (left). Both networks are typically physically wired in a star configuration using a central concentrator (shaded) for improved reliability (right).*

ATM has a heritage in telephony and as such has provisions for carrying audio. An ATM network relies on complex and expensive central switches to route data. ATM is an emerging technology which is currently being deployed as a backbone in many high capacity mission critical computer networks. This elite application does not realize the same commodity pricing as the end user LAN technologies. ATM is expected by many to eventually work its way to the desktop and even to the home in next generation cable television systems. Fast Ethernet and related technologies are currently giving ATM a good beating. It is currently unclear whether ATM will win the long term battle for world dominance its supporters, including many Fast Ethernet advocates, believe it should.

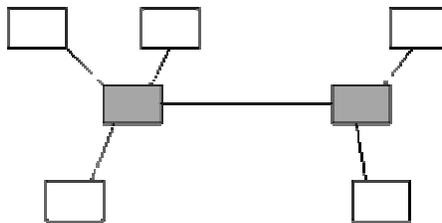


Figure 4: *ATM uses central switches (shaded) to route data between end stations.*

PC peripheral networks such as SCSI, Fiber Channel, Fire Wire and USB can have high throughput and some even have built in provisions for carrying audio and other isochronous data. These systems all currently have severe length limitations which make them inapplicable to facility wide audio distribution.

Standard Ethernet or 10BASE-T is the dominant local area network technology with 50 million nodes installed worldwide. Ethernet is generally wired in a robust start topology. Ethernet is, however, electrically a bus. As such, a transmission from any station, traveling near the speed of light, reaches all other stations almost simultaneously. Because of its bus architecture, Ethernet poses few obstacles to sample clock distribution. The problems for carrying audio over Ethernet lie mainly in its media access control scheme. Media access on Ethernet works like a telephone conference call. Anyone who

wishes to speak (transmit) first waits until no one else is speaking. The person will then begin speaking. If someone else on the call had simultaneously made the decision to chime in, there will be a collision. Both parties will detect this, interrupt themselves and wait a random period before trying again. On a busy network, collisions can occur quite regularly and the time required to resolve them can eat into the useable bandwidth. Since the process of collision resolution is a statistical one, access to the media can be said to be non-deterministic. There is a real probability that a station wishing to transmit on a busy network will wait an inordinate amount of time to do so. Furthermore, there are technical problems with Ethernet's implementation of this media access control scheme which can further exaggerate Ethernet's non-deterministic behaviors.

Fast Ethernet or 100BASE-T is standard Ethernet with all timing shortened by a factor of 10. To accommodate the faster signals, 100BASE-T requires higher grade CAT5 cable whereas standard Ethernet may be carried on CAT3 cable. Fast Ethernet has all the strengths and weaknesses of standard Ethernet and provides 10 times the bandwidth of standard Ethernet. Fast Ethernet products are perceived to be the most compatible with Standard Ethernet giving it a crucial edge in the huge Ethernet upgrade market. One downside to Fast Ethernet is that due to tight timing specifications, the standard restricts the length of CAT5 cable runs to 100 meters. By bridging to fiber, distances can be extended beyond 2000 meters.

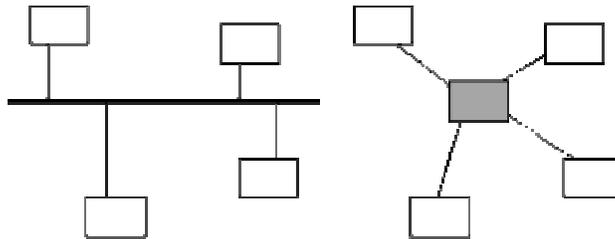


Figure 5: *Ethernet uses a bus topology (left). For reliability, Ethernet is most commonly wired in a star configuration to a central hub (shaded) in which the electrical bus is contained (right)*

Switched Ethernet is a deployment of Standard or Fast Ethernet utilizing central switch devices in a manner similar to ATM. In a partially switched installation, several stations may share a single connection to a switch (see [Figure 6](#) [right]). In a fully switched Ethernet installation, each station is wired directly to a port on a switch (see [Figure 6](#) [left]). In this configuration, if supported by both the switch and the end station, the connection can be run in Full Duplex Ethernet mode. In this mode, a station can simultaneously transmit and receive potentially doubling the bandwidth of the link. There is no possibility for collision on a Full Duplex Ethernet segment. The performance of a Switched Ethernet network rests heavily on the performance of the switch(es). Many switches are not designed with a priority given to long term real time performance as is required for audio distribution. Switches often deliver poor performance when it comes to delivering broadcast messages. In the most straightforward implementation, an audio distribution system would make use of broadcast messaging to distribute a sample clock

and to allow a single audio channel to be delivered to multiple locations most efficiently. A broadcast message is a packet which is not addressed to any specific station. Any station on the network should be able to receive all broadcasts if it is so configured. At best, broadcast messaging defeats the performance advantages gained by going to a switched architecture. In computer networking, broadcast messaging is used primarily for administrative functions, not data delivery. As such, many switches regard broadcast messages as lower priority transmissions and do not deliver them in a timely manner and possibly drop them entirely.

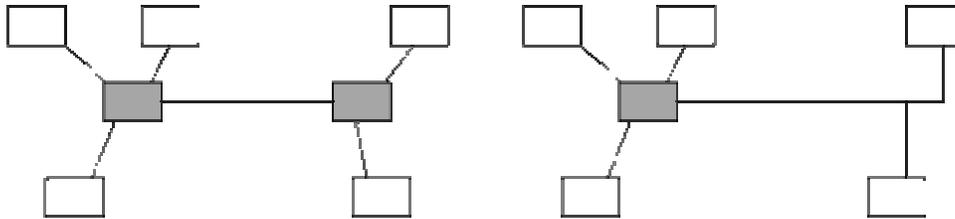


Figure 6: *Fully switched (left) vs. partially switched network (right).*

A high bandwidth network developed by Hewlett Packard known as 100BASE-VG or AnyLAN which uses a shared media, like Ethernet, but employs intelligent hubs to control access to the media. This networking technology has provisions for carrying multimedia traffic. Unfortunately, in a bitter contest which is reminiscent of VHS vs. Beta, AnyLAN has recently lost a number of acceptance battles to Fast Ethernet (including the right to call the product "Ethernet"). This trouncing is primarily due the Ethernet's tremendous inertia in the computer networking industry.

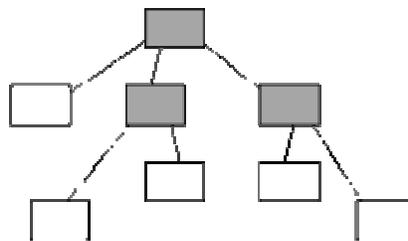


Figure 7: *AnyLAN uses a tree structure with intelligent hubs (shaded) to control access to the media. Once access is granted, transmissions may be distributed to all stations as if the network were wired in a bus topology.*

Isochronous Ethernet is a 16Mbit network standard which like Standard Ethernet can carry 10Mbit of computer traffic plus an additional 6Mbit of isochronous traffic primarily for the purpose of integrating a LAN to Wide Area Network (WAN) services like ISDN and T1 lines. Though it employs RJ-45 connectors and CAT3 wiring Isochronous Ethernet is about as similar to Ethernet as is Token Ring. Isochronous Ethernet has a mode of operation allowing all 16Mbits of bandwidth to be dedicated to Isochronous data

transmissions. Even so, the 16Mbit capability of this network is somewhat disappointing for audio distribution applications.

Selecting a Network Technology

There are limited applications in digital audio distribution for systems with only 4, 10 or even 16Mbits of bandwidth. This makes Token Ring, Standard and Isochronous Ethernet somewhat unattractive. The difficulty of distributing a sample clock on a ring topology strikes a second blow to Token Ring and also takes FDDI out of the running. Switched Ethernet can also create clock distribution difficulties. Due to its current immaturity, cost and complexity, it may be wise or necessary to avoid ATM in the short term. As mentioned before, the peripheral networks are currently distance limited making them inappropriate for audio distribution.

Media professionals cannot dismiss AnyLAN easily as datacomm professionals have. AnyLAN offers an honest deterministic and prioritized network access with 100Mbit bandwidth. Use of this network is strongly indicated for applications bringing high quality multimedia to desktop workstations. Downsides to AnyLAN are not technical in nature. Manufacturers of critical network controller chip-sets are embracing Fast Ethernet over AnyLAN at a ratio of 10 to 1. Network design consultants report a similar acceptance gap between the technologies.

Peak Audio has developed a system for distributing audio over Standard and Fast Ethernet. This technology, called CobraNet™, achieves deterministic performance and accurate clock distribution over an Ethernet dedicated to audio distribution. We believe the best network for digital audio distribution is a CobraNet™. The datacomm industry is rushing to embrace Fast Ethernet at the expense of other networking technologies such as AnyLAN and ATM which may even be technically superior. Fast Ethernet has the bandwidth required for audio distribution in medium to large installations. In certain applications, where long distances, electrically noisy environments, and few channels are involved, Standard Ethernet will be viable. Standard Ethernet may also be viable for some paging and intercom applications where low duty cycle and lower audio quality requirements may allow reasonably large applications to fit on 10Mbits of bandwidth.

Conclusions

Due to their many advantages, digital multiple access networked audio distribution systems will become increasingly prevalent. Affordable computer network technology with sufficient bandwidth for digital audio distribution applications has recently become available. Peak Audio has developed the technology which allows a dedicated Standard or Fast Ethernet network to be used as the basic infrastructure for an audio distribution system. The Peak Audio system, CobraNet™, has solved the problems of clock distribution and deterministic network access over Ethernet. A solution, such as CobraNet™ which utilizes components and technology from the computer industry has real-world advantages over proprietary solutions.