



The

Broadcasters' Desktop Resource

www.theBDR.net

... edited by Barry Mishkind – the Eclectic Engineer

DIGITAL AUDIO

Choosing a Networked Digital Audio Technology



Andrew Calvanese & Scott Johnson
Wheatstone Corporation

[May 2010] Anyone who has shopped for a broadcast audio console in the last few years is no doubt aware that networking is the watchword. Virtually every console manufacturer in the industry now offers some form of networked digital audio console and router solution. Making some sense of that is the goal of this article.

Networking gives us tremendous advantages, simplifying and streamlining the way we do our daily work. It allows us to have “any source, anywhere” flexibility, integrates the control surface into the routing system for maximum resource availability, and provides the framework for automation and other enhancements of station workflow.

FROM TDM TO AOIP

During the early years, beginning in 1998 or so, Time Division Multiplexing (TDM) was the only technology available for networked digital audio.

Recently, advances in information technology and industry developments have made possible a second type of audio networking. Audio over Internet Protocol (AoIP), often called IP Audio or simply IP, is now offered by several manufacturers and is a rapidly emerging technology.

This leaves many engineers and managers with many questions. What are the advantages and disadvantages of each technology? Which will best serve the needs of my station, staff, and workflow? For us to arrive at meaningful answers to those questions, we will need to arm ourselves with some basic knowledge of both technologies.

COMMONALITIES WITH TDM AND IP AUDIO

Before we get into what is different, let us look at what is not different. The TDM and IP technologies have many common features and characteristics. Once understood, these are factors we can safely disregard in deciding between the two.

Both types of systems:

- Digitize or sample rate convert data as it enters the system, handling it completely in the digital domain thereafter.
- Have the capacity of switching a source anywhere in the system to a destination anywhere else.
- Control surfaces are digital controls only and carry no actual audio. Digital signal processors (DSP) located in a rack or network I/O unit perform the actual busing, mixing, level control, metering, and processing in accordance with data sent to them by the control surface's faders, knobs, encoders, and pushbuttons.
- Use a star topology.
- Allow us to distribute the routing system's inputs and outputs to various locations within the station, making access extremely convenient and reducing the need for bulky, expensive point-to-point multi-pair cabling.
- The number of I/O points provides the ability to make the system redundant, or at the very least, fault-tolerant. (As we will see, though, the degree of redundancy available does vary widely between the two technologies.)
- In most cases, a single cable connects the control surface to the rest of the system.

A MAJOR COMPATIBILITY PROBLEM

At the time of this writing, both technologies are closed, proprietary systems. There are no standards for TDM audio, and each manufacturer has its own unique approach to the technology. There is no direct interoperability between systems from different companies.

Although there are some standards beginning to emerge for IP-based audio networking, at present none of the manufacturers have fully adopted any of these standards. The data transmitted by IP network devices adheres strictly to Internet Protocol standards so that it will pass properly through the necessary switches to its destination, but the *contents of the packets themselves* (the payload) is proprietary.

While it is generally true that the smaller, simpler, set-and-forget IP systems that can utilize an existing network infrastructure will be slightly less expensive, it may still come as a surprise to know that the cost difference between systems of similar capability in each technology is usually not significant.

On the other hand, as we scale up to larger systems that require multiple audio formats and full redundancy, the balance tips slightly in the other direction, and soon after that we reach a level where it is impractical to implement a system with IP technology, and TDM becomes the only choice.

So, where do TDM and AoIP differ?

TDM – TIME DIVISION MULTIPLEXING

Time Division Multiplexing is a way of putting multiple signals onto a single connection or bus by dividing the bus into time slots, and allocating one time slot to each signal.

By synchronizing with one particular time slot on the bus, we can put a signal into that slot, or read data from it. From that data, we can reconstruct the original audio. This technology is not new; in fact, telephone and computer networking systems have been using it for decades. Only its application to professional, digital audio systems is relatively recent.

TDM systems use building blocks called card racks which are, in and of themselves, routers. Racks come in various sizes, but most have around 20 card slots. Because of their high component density, some racks require cooling fans, and all have fairly hefty power supplies.

Inside each rack is a digital TDM bus, a simplified version of which is shown in Figure 1. Input, output, and audio processing cards plug into this backplane and communicate with it.




Figure 1: The TDM Bus

A PRACTICAL ILLUSTRATION

I once ate at a sushi restaurant in Los Angeles that reminded me very much of the way TDM works. A continuous conveyor ran the full length of the sushi bar and through the kitchen. The chefs placed items onto a spot on the belt, and customers would simply reach out and take any desired item from the belt.

The TDM bus is like that conveyor. It has a number of time slots available, one for each discrete audio signal in the system. An input card can place an incoming signal into a time slot on that bus, and an output card can retrieve signal from any time slot on the bus and send it out of the rack.

TDM bus and rack sizes vary, but the principle is the same. If the bus has 512 time slots, then that rack will be able to switch 512 discrete, individual audio signals.

Input and output cards are available in a variety of types covering every imaginable audio signal format, analog or digital. By plugging the desired number and type of cards into the rack, the user populates the rack with inputs and outputs. These cards are generally hot-swappable, easing troubleshooting and making installation of spares very convenient.

The TDM racks generally have very high I/O capacity, and are thus very efficient in terms of rack space and power.

THE TRICK IS IN THE TIMING

There is one aspect of TDM that is unlike the sushi conveyor, though. This digital conveyor is extremely time-accurate. In fact, it is so accurate that we know exactly what time each spot will be in front of each customer!

Because of this very accurate synchronization, a card can easily find and read its time slot without having to identify it in any other way - *just by its timing*. Each card also can depend on receiving a steady, regular stream of data, so there is no need for buffering or storage of the incoming audio. Latency (the

delay between audio entering and leaving the system) is thus very, very low, which in some situations is a big advantage.

SCALABLE

TDM is a highly scalable technology. By connecting racks together, we can build an arbitrarily large routing system consisting of thousands of inputs and outputs.

Interconnection of racks can be done via fiber or metallic (CAT5) links. The many options that are available make it is easy to build extremely robust, fully redundant systems with automatic, seamless failover using the TDM technology. And because of its time-slot-based operation, bandwidth can be rigidly planned and “reserved” during the design phase – link bandwidth is thus guaranteed - in a type of system that engineers call “deterministic.” There never are concerns about traffic spikes affecting network throughput or latency.

As mentioned before, TDM is a closed, proprietary system. No third party equipment can act as a weak link and disrupt the reliability of the network. In fact, IT network requirements are minimal, consisting primarily of connections to allow engineers to reconfigure the system as necessary.

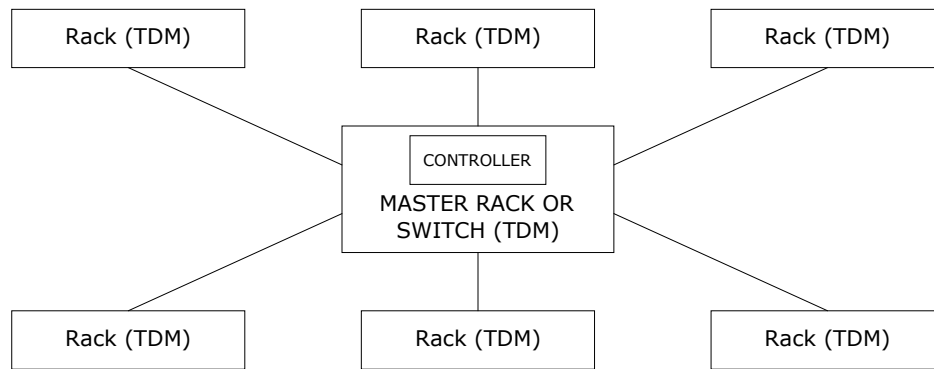


Figure 2: Star Topology of the TDM Network

In the TDM system, there is generally a single controller, in which all of the system’s intelligence resides; this controller may be backed up by a redundant one to guard against a single-point failure. All of the racks within the system communicate with this central controller so that routings can be set and interrogated and commands can be carried out. Generally the controller resides within a central rack or switch, which is at the center of a “star” topology as shown in figure 2.

Most TDM systems come pre-configured by the manufacturer, according to a design agreed upon between manufacturer and customer. Because of the concentrated network links between racks and the density of the racks themselves, fewer cables are required than in an IP system. All this tends to make installation a great deal more trouble-free.

AOIP – AUDIO OVER INTERNET PROTOCOL

IP is a technology which leverages the ready commercial availability of Ethernet routing and switching equipment normally used in Information Technology (IT). Through the use of proprietary hardware, audio signals are digitized or sample rate converted, buffered, and assembled into packages called packets, suitable for transmission over a local area network via the Internet Protocol.

“Internet Protocol” is a deceptive term as used here. AoIP routing systems cannot generally be made to work over the Internet, because they demand a level of guaranteed bandwidth which the Internet cannot reliably provide. The versatile addressing scheme of IP is the reason that we, as well as most local area networks (LANs), use it.

Unlike TDM, an IP network does not have pre-allocated time slots, or really any organized way of distributing bandwidth. Instead of the sushi conveyor, what we have on an IP network is a bit like a major highway, with the packets as cars on that highway, as shown in Figure 3.




Figure 3: IP Network Traffic

IP networks are constructed from units which we will call access points. As with TDM, each of these access points can place audio data onto the network or receive audio data from it. An advantage of this technology is that computers and other network-connected devices can also participate as access points on the network to insert or extract audio, requiring nothing more than the installation of a driver to do so.

Each packet has a header which gives detailed information about it – where it is going, where it came from, what time it was sent, when it will “expire” if not delivered, and so forth. This header adds considerable overhead to the data stream, and makes transmission less bandwidth-efficient overall, but it is necessary so that standard switches and routers can properly switch the data to its destination. The packet’s audio content, or “payload,” is of a proprietary structure.

Each packet is handled separately – there is no fixed timing relationship between them. They are simply vehicles on the same roadway and, as on many of our superhighways, delays and even collisions can occur if traffic becomes heavy enough. There is no guaranteed bandwidth here, so heavy traffic generally results in either much higher latency (longer delays) or the complete loss of some packets. Several standards for the transmission of audio over IP have been proposed or announced (EBU N/ACIP, ITU H.323, and IEEE 802.1 AVB are examples), but as mentioned, no manufacturer has fully implemented any of them.

Most of the time, the packets are directed automatically to their destination points via common Ethernet switches, whereupon they are broken down and reassembled into the original audio streams. Because the structure of each packet is standards-based, to networking components it just looks like any other packet of data.

DEALING WITH THE PACKETS

Like TDM, IP is based on the star topology. However, rather than being centered on a proprietary TDM audio switch, it is instead connected via a large, managed Ethernet switch, which is off-the-shelf IT hardware.

Managed switches, unlike the more typical unmanaged switches seen in small offices, can be specially configured to prioritize certain types of data. This is important because IP networking is statistical rather than deterministic. Bandwidth is not reserved for a particular link, but is instead allocated dynamically according to priority.

There is no guaranteed timing associated with the flow of packets. In fact, the packets are *not even guaranteed to arrive in the same order* in which they were transmitted.

This means that in order to allow the receiver to “ride through” any gaps in the incoming data and to sequence it correctly, the audio must be buffered. A certain amount of incoming data is stored in a buffer so that there is always a pool of data waiting to be processed. This increases latency; in other words, in an IP system, there is a longer delay between data entering the system and the same data leaving the system.

MULTIPLE ACCESS POINTS

Access points in the IP systems tend to be smaller and more numerous than racks in TDM networks. There are exceptions – one vendor even offers a one rack unit (1 RU) micro TDM rack – but in many TDM systems the racks are between four and six RU, and some have external power supplies.

Networks built on IP technology are extremely simple to upgrade and expand. Adding another access point is as simple as naming and configuring it, then plugging it into the switch.

In most cases, an access point is a one RU unit which handles the same number of inputs and outputs one or two cards in a TDM card rack. By virtue of their size, these units can be placed in more locations, often in spots where a full TDM card rack would not be practical. The access points require no fans for cooling, but because each has its own power supply and its own onboard intelligence, efficiency is quite a bit lower and more power is consumed.

IP access points are not as versatile as card racks. In fact, most can accept only analog, AES digital, or microphone-level signals. Other formats, such as MAD1, are not supported on IP networks, nor is embedded 601-format digital audio.

Intelligence in at least one IP-based system is distributed. Each access point has its own onboard intelligence, maintaining a full copy of the entire system’s configuration. This provides multiple levels of redundancy and makes the system somewhat fault-tolerant.

It is important to note, however, that because of the nature of IP networking, it is difficult – if not impossible – to design truly redundant systems with seamless automatic failover. While IP networks can be self-healing, the capability is vendor-specific, and healing takes place on an IT time scale. It can take from seconds to minutes for links to re-form around a failure point.

IP networking is cutting-edge technology; that fact is undisputed. It is at the center of current industry development efforts, and therefore it is likely that advancements in features and capabilities will appear in this technology first. The technology has the potential to be more “open” in the future, primarily

because the data packets are of a standards-based structure. This same structure also raises the potential for long-distance transmission of audio data using inexpensive IT-industry solutions, as well as the potential for allowing third-party interfaces to access the network.

WEIGHING THE ADVANTAGES

So, now that we have discussed the workings of both systems and some of the assets and liabilities on each side, we are left with a large number of factors to consider when we decide which system will best serve our needs.

What it comes down to, in the end, is that each system has situations where it is clearly the right choice. To clear the waters a bit, we can summarize most of what we have just learned in tabular form.

Category	TDM	IP
Latency	Minimal	Significant
Stability	Excellent	Good
Cabling costs	Low	Higher
Rack Space Economy	High	Low
Redundancy Options	Many	Few to None
Bandwidth Allocation	Static, Guaranteed	Dynamic, Not Guaranteed
Intelligence	Central	Distributed or Central
Standards-based	No	Partially
Expansion	Easy	Easier
Troubleshooting	Easy	Challenging
Power requirements	Low	Higher
Scalability	Huge	Large
I/O formats	Virtually All	Few/Basic
Built-in Fiber	Yes	No
Support	Manufacturer	Manufacturer plus IT supplier
Fans	Maybe	No
Status of Technology	Stable, Mature	Emerging, Leading-Edge

MAKING THE RIGHT CHOICE

While each technology has its strengths and weaknesses, some patterns do emerge from the above information.

For large, mission-critical systems where redundancy, reliability, and scalability are important factors, TDM technology remains the best solution. Likewise, if you require the ability to directly use audio formats other than AES or analog, you will probably want to go with a TDM system.

Customizable logic configurations and large audio I/O complements in minimum space also call for TDM designs. TDM requires the least cabling and is therefore the most efficient in terms of cable tray or conduit space.

If you are determined to have the most cutting-edge, modern technology possible in your plant, you will probably want an IP system. In smaller systems, IP technology can be highly cost-efficient and easy to update and expand when the time comes.

Surprisingly, you already probably know which technology is best for you. If you are more comfortable with control surface functions, audio levels, and sound quality than you are with driver updates, network topology, and switch port configurations, then TDM is for you. IP technology is not for novices.

On the other hand, if you would rather that your audio devices all looked and acted just like your other computers, then IP is for you. If you are trained and experienced in IT and are already comfortable managing a complex network, then adding IP-based audio to the network is a natural progression.

In either case, your best bet is to approach the selection with an open mind, and work closely with the manufacturer as you move toward making a choice. There is a fair chance that whatever system you choose will be one you will live with for many years to come. For that reason, above all, be sure that the technology you pick is one that with which you are -- and are likely to remain -- most comfortable.

- - -

Andy Calvanese is the Engineering VP and Scott Johnson is a Systems Engineer for Wheatstone. Email contacts are acalvanese@wheatstone.com and scottjohnson@wheatstone.com

[Return to The BDR Menu](#)