Beyond Personal Monitoring

Aviom A-Net and A-Net Pro in networking

By Carl Bader and Chandler Collison

This Factory Direct was submitted by Aviom. Live Sound *makes every effort to eliminate any use of marketing inspired hyperbole.*

From the outset, Aviom has been defined as a distributed monitor mixing company. We began by analyzing the costs and frustrations of traditional stage monitoring, with its significant engineering talent and equipment costs and/or excessive on-stage volume, in addition to "never quite the right mix" and so on.

Our view was that there could be a better way, and apparently, many of you agree: a recent *Live Sound – ProSoundWeb* poll (*September 2004 issue*) showed almost 58 percent of respondents identifying monitor mixing as the hardest job in sound reinforcement.

The basic vision focused upon distributed control of personal monitor mixes on stage and in the studio. Instead of relying entirely upon a mix engineer who's often trying to satisfy many simultaneous (and competing) requests, the goad is putting more of the control of monitor mixes in the hands of the performers – the people who actually rely on the mixes. They could then make changes as they wanted/needed, leaving the engineer free to focus on fewer mixes, such as that of the lead singer.

This initial thinking led to the development of a new audio transport technology, called A-Net Pro digital audio transmission protocol (patent pending). In turn, this new protocol has led Aviom down a path far beyond in-ear personal monitor mixing.

A distributed monitor mixing system like the one we envisioned posed specific technical challenges. The

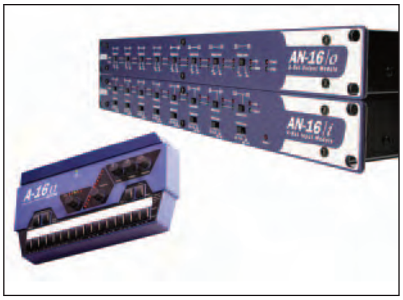

The AN-Series digital snake carries a combination of sends and returns from point-to-point, and to devices such as the Aviom personal monitoring system.

principal issue was finding or developing a way to distribute audio channels to the performers' mixers.

Analog distribution techniques were (are and probably always will be) expensive, complicated, and inconvenient. At the same time, existing digital technologies failed to meet our performance requirements. The primary goal was a system that was easy to use and understand – something performers and audio engineers alike could set up and use without having to master computer networking. High audio quality standards also needed to be satisfied.

During initial design, our objective for the distribution technology was that it support a system providing:

Real-time transmission with no perceptible latency;

Multiple channels of high quality uncompressed audio;

Plug-and-play set-up without computer-based configuration and associated costs;

Infinite network expandability through both serial and parallel connections.

Most current digital audio transport technologies are based directly on Ethernet. Certainly a powerful transmission technology, Ethernet has never been intended for use with high-density, constantly streaming, time-critical data such as audio.

An Ethernet network is a non-deterministic system, which we believe means transmission timing is unpredictable from packet to packet. The negative effect of this on audio systems can be mitigated by buffering a few samples at each device.

However, buffering adds delay, and with pro audio, timing is clearly critical. Most pro audio systems now include digital consoles, preamps, and signal processors, and each digital component introduces more delay from processing and transmission. Thus system-wide latency can very quickly grow to produce delays that simply can't be tolerated, requiring the latency of each component to be kept extremely low.

BREAKING IT DOWN

The Aviom design team stripped Ethernet down to its most basic level, the "PHY" layer. This is the part of the protocol that defines total data capacity, cable and connector type, and run length. When the rest of the protocol – those layers that deal directly with the management of the data packets – are pared away, what's left is simply an inexpensive yet robust, bi-directional data pipe capable of 100 megabits per second (mps) data throughput.

The key to achieving improved performance over Ethernet-based systems was to eliminate the need for these other layers. An A-Net network is fully deterministic: the path and timing of every transmission are determined from the outset, and so data pass through devices much more quickly, without collisions or bottlenecks and without the data management layers of Ethernet. This enables substantially faster data throughput.

One of the leading Ethernet-based audio protocols introduces 5.33 milliseconds of latency per device in its standard mode. The comparable delay with A-Net is less than 160 microseconds, or about 1/30th the delay (97 percent less).

In any digital system, regardless of the transport protocol, A-to-D (analog to digital) and D-to-A conversion introduces more delay. With A-Net, the total analog-in-to-analog-out latency (including transmission and conversion on both ends) is rated as just 880 milliseconds – well below the speed an Ethernet-based protocol can attain for just a single module and without factoring in conversion delays.

An example of A-Net's efficiency is the speed with which one A-Net receiver passes a transmission on to another: less than 0.7 milliseconds of added latency. What this means, in practical terms, is that more than 150 A-Net devices could be daisy-chained and the longest delay – from analog in at the first device to analog out at the last box – would still be less than a single millisecond. (**Figure 1**)

An added benefit of the A-Net protocol is that cable runs can be lengthened. An Ethernet-based protocol can transmit 330 feet before the signal needs to be refreshed by another Ethernet device. Because A-Net does not require the back-and-forth data pings necessary in a non-deterministic system, cable runs can be extended to 500 feet between components. (Any A-Net device will refresh the signal for another 500 feet.)

Users who need added length can use third-party converters and transmit A-Net over fiber optic cable, for up to 1.6 miles on multi-mode fiber and 50 miles on single-mode fiber.

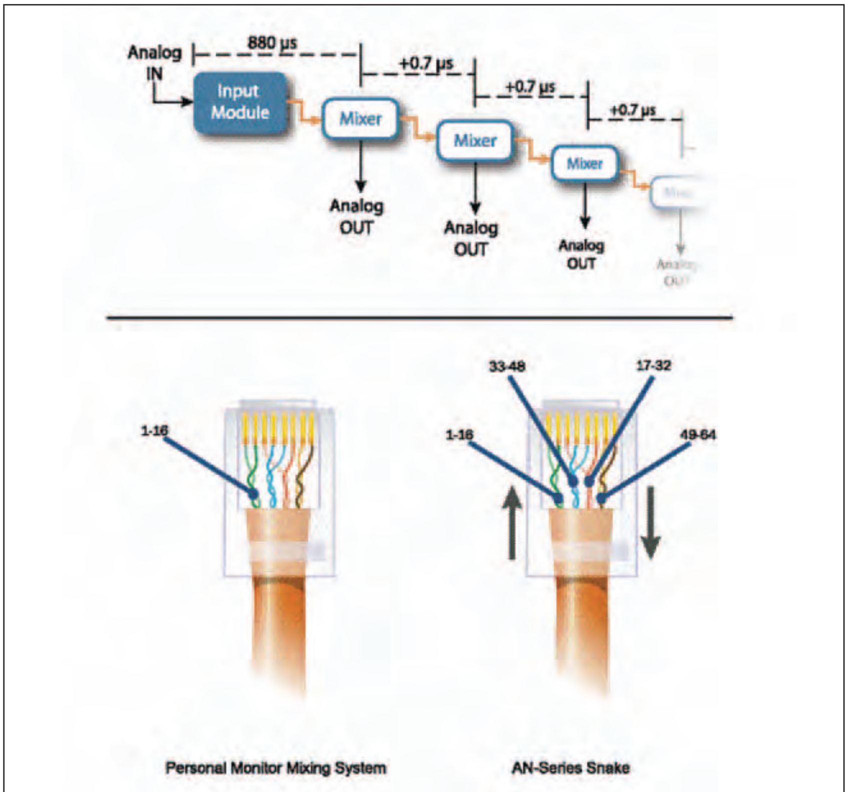

Figure 1: Latency in the personal monitor mixing system from analog in to analog out, show at top. Below, with a personal monitor mixing system, only one CAT5 pair is used, while with an AN-Series digital snake, all four can be used with the direction of each pair individually reversed.

BEYOND MONITORING

Design decisions for A-Net were made based on the requirements of the personal monitor mixing system. Thus it was originally designed to transmit in only one direction, from the A-Net input module out to the performers' mixers.

In addition, the protocol was built to carry 16 channels, not because of any bandwidth limitations, but because 16 channels provide users good flexibility in monitor mixing without overwhelming them with choices and controls. (In fact, A-Net utilizes only one of the four wire pairs in a CAT5 cable and does not fill even that single pair.)

The success A-Net met in this initial application sparked interest from users (and other companies) in other, expanded applications of the protocol. An exploration commenced into use of A-Net for other types of distributed audio networks.

The first expansion introduced bi-directional streaming of four times as many audio channels. In the monitoring system, 16 channels are transmitted on one wire pair. CAT5 cable has four such pairs, and there's no reason A-Net can't use all four pairs simultaneously. (also **Figure 1**) This is the basis of the bi-directional 64-channel AN-Series digital snake product line.

The restriction of the AN-Series: because each wire pair carries 16 channels, sends and returns must be allocated in blocks of 16 channels. Thus, the AN-Series Snake can be operated in 64 by 0, 48 by 16, and 32 by 32 (or smaller) configurations.

In many applications, 16-channel granularity is all that's needed. But it's clearly not the most efficient way to make use of the available bandwidth.

This led us to A-Net Pro, a more full-featured version of the A-Net protocol. Clearly, it will cost more, but with substantial added power – and – without sacrificing the central performance specs of original A-Net. (For the record, we plan to continue to fully support and introduce new products based on A-Net.)

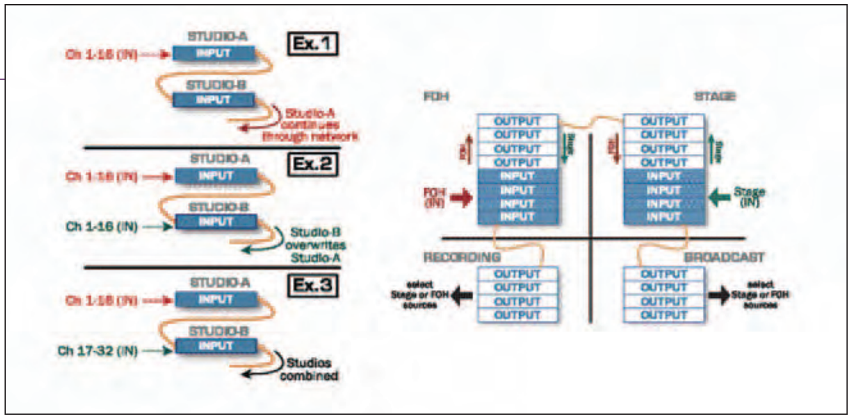

Figure 3: Left, three different ways of controlling an A-Net stream in Manual mode. Right, a system using Manual distribution control of 128 simultaneous channels.

VARIABLE SAMPLE RATES

A-Net Pro carries up to 64 audio channels on a single wire pair, still at 24 bits and with ultra-low latency. In addition, it supports variable sample rates within three basic ranges: 39 kHz to 52 kHz (for 44.1 kHz and 48 kHz applications), along with 2x (96 kHz) and 4x (192 kHz) ranges.

In the higher ranges, channel counts are reduced by a corresponding factor because the increased sample rate requires additional bandwidth for each channel. This support of variable sample rates provides the ability to track to a device whose sample rate is not perfectly locked, such as a digital tape deck or a SMPTE video desk.

A-Net Pro also offers flexible clocking options, with any device on the network able to be set as the audio clock master, or the system can be synced to an external master clock from any point in the network. Thus, an A-Net Pro system can smoothly operate with other digital device, for applications such as synchronizing to digital consoles or in facilities with house clock.

Perhaps the principal flexibility comes from the level of control users have over the distribution of audio. Each A-Net Pro64 ASIC – the A-Net Pro chip – has a full crosspoint switch at both input and output.

While original A-Net has 16-channel granularity, the Pro64 ASIC delivers channel-by-channel control: any audio channel can be loaded into any one of the 64 slots in an A-Net stream. Similarly, any audio output can be drawn from any one of the 64 A-Net slots.

The Pro64 ASIC supports two basic modes of distributing the audio. The first is as a 64-channel "omnidirectional" audio network, with no upstream or downstream.

In a traditional system, output modules must come after input modules in the signal flow (in other words, downstream). This can be problematic, especially in permanently installed networks where directional limitations can substantially restrict where certain signals can be transported.

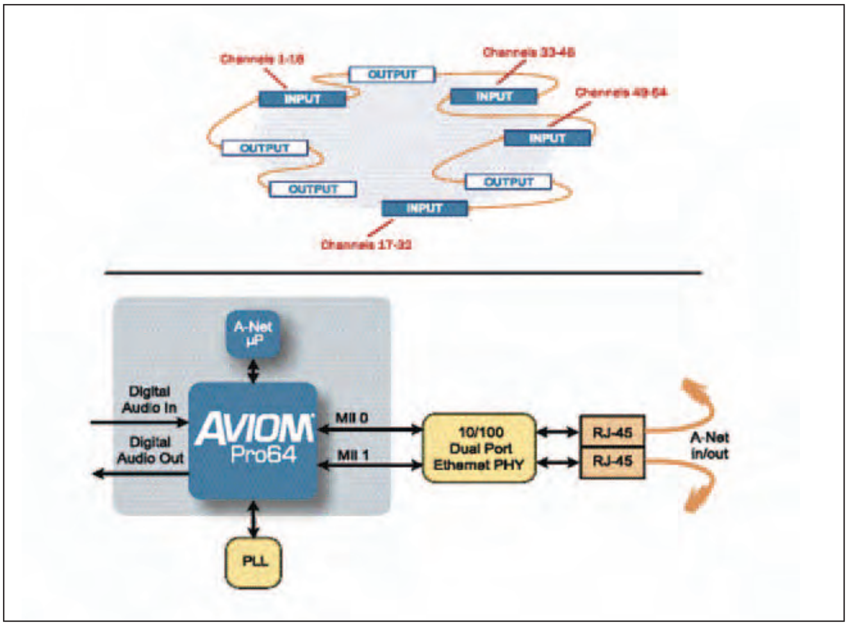

Figure 2: Above, a basic network with all 64 audio channels available at each of four output nodes. Below, the architecture of the A-Net Pro64 ASIC from digital audio in and out to A-Net in and out.

EVERYTHING EVERYWHERE

A-Net Pro eliminates this restriction, making 64 channels simultaneously available at every point in a network – something we call "everything everywhere." (**Figure 2**) This is A-Net Pro's Auto mode: the user simply connects the devices and sets any one unit as the system "Control Master." The system will then self-configure into a 64-channel everything-everywhere distributed audio network.

Alternatively, A-Net Pro allows users to manually control the composition and distribution of the two A-Net streams. Every Pro64 ASIC is fully bi-directional. (also **Figure 2**) The chip receives two simultaneous streams and retransmits them after completing its assigned functions (such as loading more audio into the stream).

In Manual mode, up to 128 discrete channels are available on a cable: 64 channels on one wire pair and 64 different channels on a second pair, heading in the opposite direction. But because signal distribution is controlled manually, there is, in fact, no limit to the total number of channels that can be in the network.

For instance, a group of channels can be transported from one studio to another, where some are replaced by new signals. This new A-Net stream, comprised of selected signals from the two studios, then continues through the network. (**Figure 3**)

Selective overwriting of channels already in the A-Net stream removes network-wide channel limits. The only restriction is that no more than 128 audio channels (64 in each direction) can be in the A-Net stream at any point in the system. (also **Figure 3**)

Further, in both Auto and Manual modes, any number of audio devices can be connected to the network. The crosspoint switch in the ASIC allows the user to decide which signals get distributed. This provides important flexibility, because any audio not loaded into an A-Net stream is still live audio and can be monitored and processed locally.

Thus, a band can be working in Studio A, for instance, connected to the network but not loaded into the A-Net stream, while another band is working in Studio B, using the network space. For the next set of sessions, Studio A can utilize the network without physical re-patching.

A key element is the ability to interface with a PC. Since its inception, Aviom has designed technology that does not require PC-based configuration or operation; rather, the focus is upon plug-and-play, stand-alone operation. A-Net Pro will continue this – the protocol does not require the user or installer to connect a computer at any point.

However, the Pro64 ASIC's Application Processor Interface will allow PC-based system management. Aviom and its corporate partners will develop software applications for use with the A-Net Pro protocol.

MORE THAN AUDIO

A-Net Pro distributes more than a lot of audio channels. Ethernet-based audio distribution protocols are able to transport more than audio as well, but their non-audio data compete with audio for network resources. Performance can be compromised.

On the other hand, A-Net Pro reserves space both for audio and for other types of data. As a result, these data are never in conflict to possibly compromise system performance. An A-Net Pro network is a hybrid system in the sense that it can carry multiple kinds of data, but the various kinds of data never compete for network space. (**Figure 4**)

These other streams fall into three main categories. One is command and control, used by the system for network monitoring and systemic functions. A-Net Pro also has a bi-directional talkback channel for intercom purposes.

However, the most powerful non-audio data stream is A-Net Pro Virtual Data Cables (VDC), which can be assigned to multiple simultaneous different data types with the information available throughout the network. Using this VDC, a user can, for instance, connect a MIDI controller to an A-Net Pro device and transmit MIDI data to an unlimited number of MIDI slaves connected to other A-Net Pro devices.

Because the VDCs are integrated into the protocol, an infinite number of 500-foot runs are possible for a range of data types without running additional cable. (also **Figure 4**)

NOW & THE FUTURE

In the third quarter of 2005, Aviom will introduce the first products in its own line of Aviom-branded networking products based on the A-Net Pro protocol. However, the useful applications of this protocol extend beyond a single company's product designs.

With this in mind, we're now debuting the A-Net Pro64 ASIC System, offering A-Net Pro protocol offering full crosspoint switching at both input and output with single-channel granularity, as well as optional PC interfacing, VDCs and more.

As a result, other audio manufacturers will be able to easily integrate the protocol into their own products – mixing consoles, loudspeakers, digital signal processors, instruments, etc. – so the size and scope of the A-Net audio network can be expanded.

Aviom modules are designed to handle any necessary A-to-D and D-to-A conversion and transport signals throughout a network – from stage to front-of-house, studio to broadcast – any location to any other location throughout a system/building.

With the Pro64 ASIC System, the network can begin right at the mixing console or digital keyboard and extend through the signal processors and recording consoles all the way to the system's power amplifiers and/or loudspeakers.

All of this on a single (and inexpensive) cable. ■

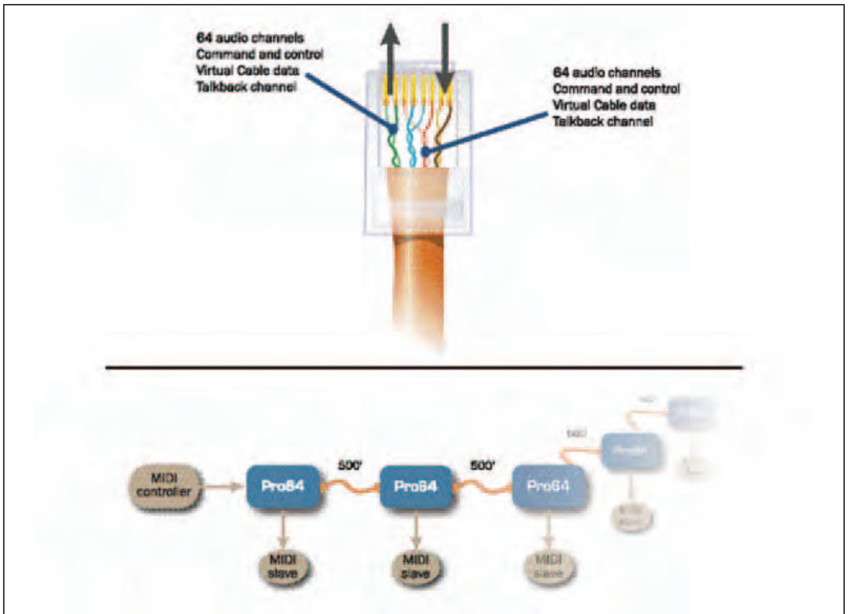

Figure 4: A look at A-Net's use of wire pairs in a CAT5 cable (above). Using VDC, signals from a MIDI controller can be distributed 500 feet to an unlimited number of MIDI slaves using A-Net Pro modules (below).

Carl Bader is president and co-founder of Aviom, while Chandler Collison heads up new business development for the company. Reach him at ccollison@aviom.com.