www.audioxpress.com

audioxpress

ADVANCING THE EVOLUTION OF AUDIO TECHNOLOGY

Market Update—Digital Signal Processing

DSP Evolution
Audio Processing Powered by AI

By J. Martins

Automotive Audio
Interior Car Acoustics and DSP

By Roger Shively

R&D Stories

A DSP Audio Processor at the Heart

By Al Clark

Fresh From the Bench

SPL Phonitor One d
Headphone Amplifier System

By Stuart Yaniger

50 Years

Audio Electronics

Ultra-Linear Amplifiers: Why It Works And When It Will Not (Part 2)

By Morgan Jones

Amplifier Design:

Cascodes, Folded Cascodes, and Current Mirrors (Part 2)

By Morty Tarr

Audio Memory

The History of the American Magnetic Head Industry

By Scott Dorsey

Sound Control

From Open Office to A/V Recording Complex

By Richard Honeycutt

$9.00US $10.00CAN

NOVEMBER 2021

HAMMOND

MANUFACTURING / DESIGNING TRANSFORMERS & ENCLOSURES SINCE 1917

Hammond Manufacturing | www.hammondmfg.com

November 2021 ISSN 1548-0628

www.audioxpress.com

audioXpress (US ISSN 1548-0628) is published monthly, at $50 per year for the US, at $65 per year for Canada, and at $75 per year Foreign/ROW, by KCK Media Corp., P.O. Box 417, Chase City, VA 23924, US Periodical Postage paid at Chase City, VA, and additional offices.

Head Office:
KCK Media Corp.
P.O. Box 417
Chase City, VA 23924, US
Phone: 434-533-0246
Fax: 888-980-1303

Subscription Management:
audioXpress
P.O. Box 417
Chase City, VA 23924, US
Phone: 434-533-0246
E-mail: custserv@audioxpress.com
Internet: www.audioxpress.com

Postmaster: Send address changes to:
audioXpress
P.O. Box 417
Chase City, VA 23924, US

Advertising:
Strategic Media Marketing, Inc.
2 Main Street
Gloucester, MA 01930, US
Phone: 978-281-7708
Fax: 978-281-7706
E-mail: audioxpress@smmarketing.us

Advertising rates and terms available on request.

Editorial Inquiries:
Send editorial correspondence and manuscripts to:
audioXpress, Editorial Department
P.O. Box 417
Chase City, VA 23924, US
E-mail: editor@audioxpress.com

What's Old Is New Again

In this edition of *audioXpress*, Scott Dorsey offers an interesting historical journey into one of the most critical pieces of technology in the heyday of the audio industry: magnetic recording and playback heads. The people and the numerous companies involved played a vital role in shaping the music industry and popular music as we know it. From recording sessions at Sun Studios to 8-track Ampex recorders introduced in 1955 and explored by Les Paul, magnetic heads raised studio music production to unimaginable artistic heights. Computers and digital recording followed and enabled even more creative possibilities. All along, this interesting story serves to remind us how vital components and technologies pioneered in the audio industry dwarf in comparison to the massive business subsequently generated in other industries, and the formidable dynamics that creates.

In this same edition we jump to the cutting-edge of today's technology—the use of artificial intelligence in digital audio processing. As *audioXpress* focuses once again on the topic of digital signal processing, our Market Update feature discusses one such pivotal moment, with audio technologies now front and center in the priorities for consumer electronics and the powerful mobile communications industries. All the while raising unprecedented attention for automotive, security and even medical applications.

In this article, I tried to look ahead at the implications of machine learning, artificial intelligence, voice recognition, and the related consequences of transferring computing power and DSP directly to edge devices. This has massive implications in the audio industry, and particularly for the many companies that have pioneered those efforts and are showing practical results. From Texas Instruments to Analog Devices, DSP Concepts, Syntiant, and Synaptics, the history of DSP continues to be written in the US, where most of the algorithm and software companies are also strategically located.

From California to Massachusetts, an interesting history is being written of the efforts to raise DSP to a whole new level, on-device, at the edge, using AI/ML. The same way as digital recording and computer systems transformed the audio industry in the last 40 years, these (apparently farfetched) technology efforts are having a tremendous impact over audio applications—and they are just starting.

As detailed in this edition, evolution is happening faster thanks to mobile audio, voice recognition technologies, wearables/hearables, and the transition to spatial audio formats. Happening so fast that the companies we are just starting to recognize for their contributions are quickly absorbed in the frenzy of acquisitions that is already taking place. For developers, this poses a challenge since you might just be starting to explore a new hardware option, and learn how to code for an exciting platform, just to find out that your idea is readily available in existing DSP libraries and software modules for common consumer products.

That is why it is so important to cross-reference developments in multiple industries, and this edition includes two additional perspectives on audio processing. Roger Shively describes how dramatically things are changing when it comes to designing interior car acoustics and automotive audio applications, adding DSP to the equation. And Al Clark of Danville Signal Processing shares his journey, working with Analog Devices and DSP Concepts key DSP offerings to build a solid business in pro audio, which he recently expanded in a new versatile DSP solution for home audio and high-end enthusiasts.

Available technology applications demonstrating the power of audio processing by AI and ML are for now restricted to cloud-based implementations, mainly to improve and train the models. Speech-to-text for automatically generated subtitles and transcription in multiple languages is now using machine learning to improve accuracy, combined with image recognition and high-level context analysis. And AI is being applied to improve spatial simulation of sound sources, including for binaural processing. New SoCs with powerful neural engines launched this year are able to handle noise reduction, source separation, echo and reverberation removal, or applying regenerative algorithms to place sounds in different acoustic environments—all with audio processing directly on device.

I could also mention more interesting DSP stories from the professional audio market, where digital audio mixers and networked DSP processors are now able to offer dozens of channels of the best analog circuit emulation. Yes, even using AI-based tools for tape saturation to overcome limitations of standard DSP estimation of analog component interactions. You still need to master analog audio electronics, because now we can explore anything imaginable in that domain.

J. Martins
Editor-in-Chief

OUR NETWORK

audioxpress | circuit cellar | LIS LOUDSPEAKER INDUSTRY SOURCEBOOK | VOICE COIL

SUPPORTING COMPANIES

All Electronics Corp.	55
Audio Precision, Inc.	40
Clockworks Signal Processing, LLC	55
COMSOL, Inc.	9
Danville Signal Processing, Inc.	31
Hammond Manufacturing, Ltd.	3
JJR Acoustics	13
KAB Electro-Acoustics	33
KLIPPEL GmbH	43
Lundahl Transformers AB	15
Menlo Scientific, Ltd.	61
Oxford Digital, Ltd.	37
Parts Express International, Inc.	68
SB Audience	57
Syntiant	17
Triad Magnetics	67
Warkwyn	33
Western Electric	2
ZUMI Systems	19

NOT A SUPPORTING COMPANY YET?

Contact Peter Wostrel (audioxpress@smmarketing.us, Phone: 978-281-7708, Fax: 978-281-7706) to reserve your own ad space in the next issue of the magazine.

Contents

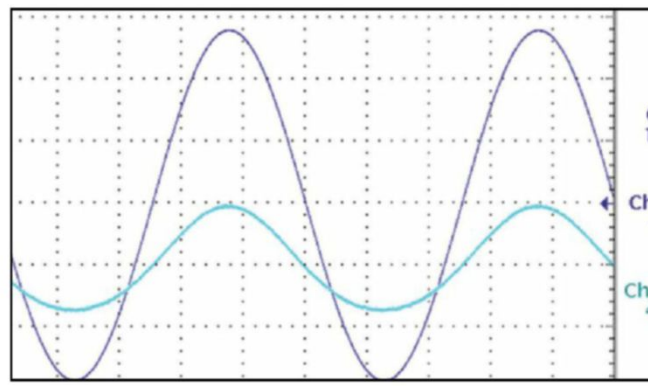

Features

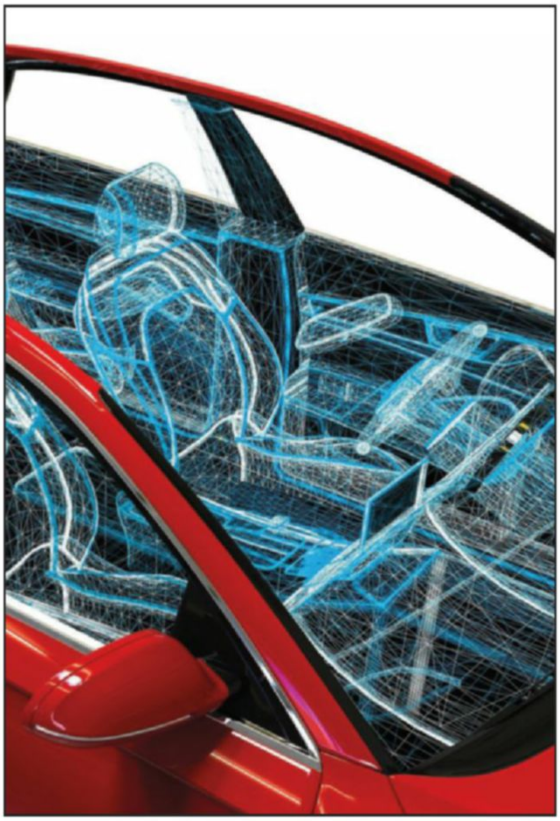

Market Update—Digital Signal Processing

Columns

Departments

Websites

www.audioxpress.com

www.voicecoilmagazine.com

www.cc-webshop.com

www.loudspeakerindustrysourcebook.com

www.circuitcellar.com

www.linuxgizmos.com

@audioxp_editor audioxpresscommunity

linkedin.com/company/audioxpress

SPL Phonitor One d Headphone Amplifier System

Entering the Matrix

By
Stuart Yaniger
Photos by Cynthia Wenslow

SPL of Germany has developed a new product line that combines the sound of its studio-oriented designs with a desktop-friendly option at an affordable price. First in its Series One is the Phonitor One d headphone amplifier ($799), which also features an acclaimed AKM premium USB DAC and balanced stereo line output. A great combination that Stuart Yaniger puts to the test and reviews.

"If you're talking about what you can feel, what you can smell, what you can taste and see, then 'real' is simply electrical signals interpreted by your brain."—Morpheus

SPL electronics (its idiosyncratic non-capitalization) is a German producer of gear targeted to recording studios and prosumer users. The company's product line is quite extensive, and includes mic preamps, recording interfaces, mixers, signal processors… about anything needed in studio electronics. SPL's products are designed and (unusually) built in Germany, and the studio gear has (for me) a very cool appearance. Its original claim to fame was an effects device called the Vitalizer, which uses a combination of level-dependent filtering, equalization, and harmonic generation to achieve more "punch" in a mix.

But SPL is also involved in the consumer market. According to Hermann Gier, Managing Partner, "We have a strict separation between the hi-fi and the Pro Audio markets. The products of the Professional Fidelity category are designed for the hi-fi audience whereas the products within the Studio/Mastering and Plug-Ins categories are for the professional users."

SPL's Phonitor line comprises headphone amplifiers for both the pro and consumer markets. It starts with the Phonitor se, its entry-level consumer headphone amplifier. The Phonitor x and the Phonitor xe are the top models in that category. The Phonitor 2 is targeted for professional studio and mastering. Compared to its amplifiers targeted to the consumer market, it has further functions that are needed in production such as Solo L/R, Phase Invert L/R, as well as Center Level adjustments of the Phonitor Matrix (more about that shortly).

The item under review here is SPL's new Phonitor One d, a combination DAC, preamp, and headphone amp for consumers, which has some interesting circuit twists. One of the technical differentiators running through SPL's product line is its 120V circuitry, which SPL calls VOLTAIR. In a nutshell, instead of the usual ±12V to 15V power

Phonitor One d
SPL electronics GmbH
P. O. Box 1227
41372 Niederkruechten
Germany
https://spl.audio/en
Retail Cost: $799

supplies common to most IC op-amp circuits, the SPL VOLTAIR products use ±60V rails. The higher power supply voltage potentially enables a larger output swing, and indeed, its studio gear runs signals up to 24dBu (about 12.3Vrms). The VOLTAIR amplifiers are used in the One d.

Product

The Phonitor One d (which I'll call the One d for convenience) is a solid and attractive unit (**Photo 1**). It has nothing of the plastic and flimsy metal of a lot of inexpensive headphone amps, but rather is solid-feeling, with an anodized black finish and a hefty all-metal case. The control knobs (volume and crossfeed) both feel velvety and luxurious in operation. The headphone output is a standard 1/4" TRS jack. Unusually (these days), the volume control here is implemented with an analog potentiometer controlling the preamplifier stage. Unlike digital control, this relies on tight matching between the potentiometer sections all through the rotational range to have stable channel-to-channel balance.

Photo 1: The SPL Phonitor One d has a simple and solid front panel layout with easy-to-see graphics.

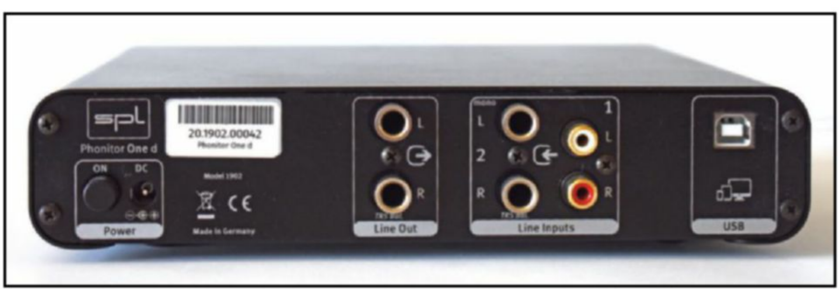

Photo 2: The rear panel of the One d has three sets of input connections and one set of fixed-level balanced outputs.

Photo 3: The One d's internal layout and component quality are excellent.

The One d has three sets of inputs—balanced and unbalanced analog and USB—selectable from the front panel with a three-way toggle switch. SPL provides an ASIO driver for the USB connection, which during my use was absolutely stable and non-buggy. The rear panel (**Photo 2**) has all three input source connectors—RCA phono plugs for unbalanced and TRS plugs for balanced analog, as well as the USB connector. There's also a fixed-gain 1/4" TRS balanced analog line output. A pushbutton on-off switch and power input from the external supply complete the picture.

Opening the case, the quality of the innards is consistent with external impressions (**Photo 3**). The circuit boards are well laid out and populated with high-quality components. Op-amps used for the analog circuitry include SPL'S 120V-rated OPX for the headphone amp and the low noise and low distortion OPA2134 for the preamp. Headphone output devices appear to be paralleled MJE340/MJE350 complements. The DAC is the well-regarded 32-bit 768kHz-capable AKM AK4490EQ, from its so-called "Velvet Sound" series.

The power supplies feature active regulation, with 7812/7912 three-pin regulators for the lower voltage portions and TL783 used for the ±60V rails.

Problems to be Solved

I will start by confessing that I'm not an avid headphone user. Some really excellent ones have passed through here (most notably several versions of the Stax electrostatic headphones). The 1MORE Triple Driver Over Ear phones (see Resources), which aren't state-of-the-art, but are still quite good, are my daily driver. And I try to take the opportunity to listen to various current high-end units at trade shows. So it hasn't been lack of exposure to quality headphones that has relegated them to a distant second place in my heart.

How do I not love them? Let me count the ways. First, physical discomfort, but (for example) the 1MOREs seem to fit my ears nicely and are good for 2 to 3 hours before I have to rip them off my head and yell, "ENOUGH!"

Second, the imaging and soundstage follow the user's head motion instead of being fixed with respect to his or her surroundings.

Third, headphone images seem to come from inside the user's head, stretch unnaturally widely, and evince a weird stuffy feeling when material is pan-potted far left or right; this is all too common in the music of my youth (the 1960s and the 1970s) and so-called LCR panning on contemporary mixes.

The first objection is highly user dependent, and of course, will depend on the match between the headphone's physical design and the user's ears and head. Given the stubborn refusal of humans to standardize their head and pinna shapes, it will be impossible to design a universal solution, but the diversity of headphone mechanical design ensures that there's likely to be something out there that you can comfortably wear.

The second issue can have technological solutions. Keeping the image fixed with respect to surroundings requires head tracking, but it's been accomplished, albeit not inexpensively, by several manufacturers, most notably Smyth Research (see Resources). With modern accelerometers mass produced for cellphones, do not be surprised to see that kind of function built into headphones and amps in the next year or two.

The third issue, pan-potting sounding odd and uncomfortable, has in the past been addressed by reasonably simple approaches. There's a fundamental difference between loudspeakers and headphones in that the left and right loudspeakers in a stereo setup are not isolated; that is, the acoustic output of the right speaker reaches both your right ear and your left ear. There is a slight timing difference due to the width of your head and the frequency response and level of the sound from the right loudspeaker impinging on your left ear has some alteration because of the acoustic effect of your noggin in the way of the sound waves (**Figure 1**). Of course, by symmetry, the same is true for the acoustic output of the left speaker. Because of the timing, frequency response, and level differences between channels, your brain synthesizes a stereo image.

Besides the crosstalk between ears from the direct sound, unless you and your speakers are in an anechoic chamber, the room reflections also affect the mixing of channels between the ears, with even more timing, level, and frequency response alterations. This provides your brain with even more clues that the sound is emanating from outside your head and allows it to synthesize a virtual image.

By contrast, headphones completely isolate the left and the right channels, so there is a massive difference in the way the stereo image is presented to your ears, and your brain will synthesize a very different stereo presentation.

The Phonitor Matrix

In order to overcome this effect, over the years there have been many attempts to electronically create a crosstalk signal, which simulates the acoustic crosstalk inherent in speakers and rooms, while accounting for head effects. The earliest analysis and engineering solution that I'm aware of was a passive line level circuit proposed by Ben Bauer (see Resources), which added to each channel a portion of the other channel with a delay of slightly under a millisecond and some level adjustment and equalization. Over the years, this basic concept was expanded and refined by Siegfried Linkwitz (see Resources) and many others, to try to more accurately model the effect of acoustic crosstalk around your head.

A general implementation of this circuit is shown in **Figure 2**. The passive circuitry shown in the dashed box is meant to be generic—this is not necessarily the exact passive frequency and time-shaping circuit used in the One d, but it is a fairly common one giving similar frequency responses as the Matrix circuit used in this headphone amp. The rest of the circuit is taken from SPL's schematics, and the op-amps used as inverting amplifiers with

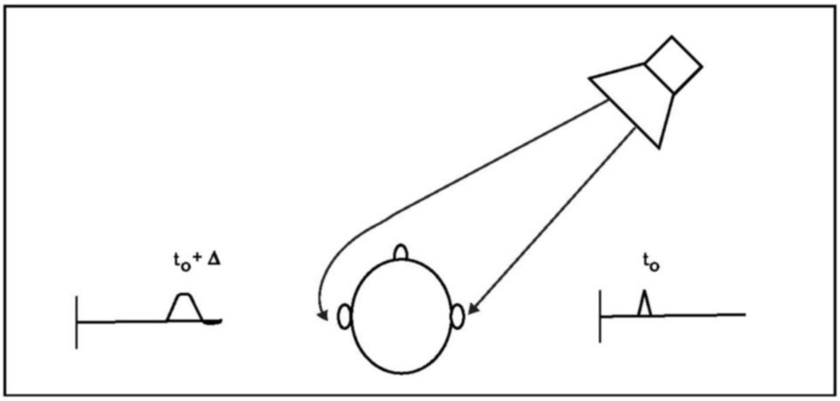

Figure 1: With stereo loudspeakers in a room, each channel feeds the opposite ear with a time delay and altered frequency response.

About the Author

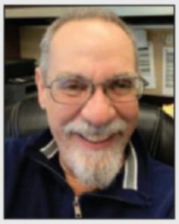

Stuart Yaniger has been designing and building audio equipment for nearly half a century, and currently runs a technology consulting agency in western New York. His professional research interests have spanned theoretical physics, electronics, chemistry, spectroscopy, aerospace, biology, and sensory science. One day, he will figure out what he would like to be when he grows up.

the crosstalk circuit sandwiched between them are, in SPL's implementation, the high voltage VOLTAIR amplifiers.

SPL claims that its Phonitor Matrix will localize sounds in the manner of loudspeakers, and even suggests that a headphone-wearing mixer can use the Matrix to allow mixing for loudspeaker playback. According to SPL, the virtual angle of the sound is equivalent to a ±30° loudspeaker spread. But this being more of a consumer product, the Matrix's principal use would be for increased listening enjoyment from commercial recordings.

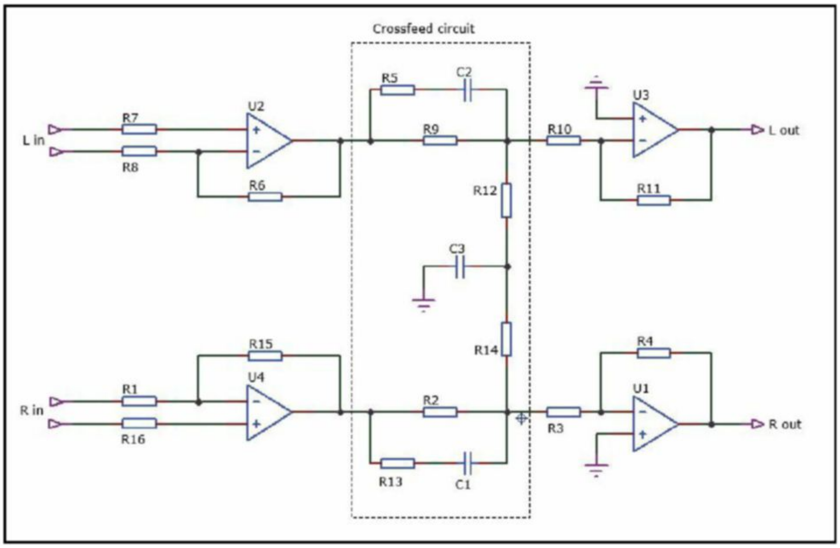

Figure 2: In the Phonitor One d, the crossfeed circuit is sandwiched between two high voltage inverting amps. The particular RC passive circuit shown here may or may not represent the actual matrix circuit used by SPL. (Illustration courtesy of Jan Didden)

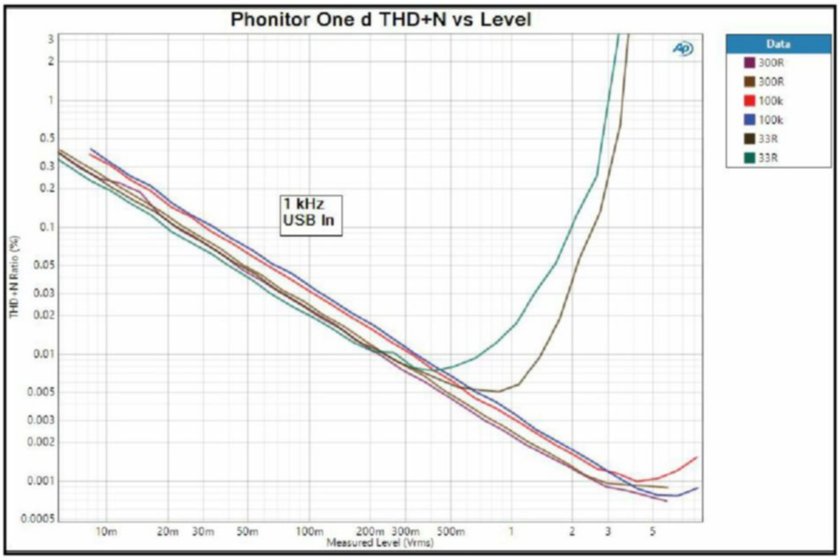

Figure 3: The Phonitor One d has sufficient drive capability for headphones having normal sensitivity.

The Phonitor One d in Use

My headphones are relatively low impedance (32Ω nominal) so are something of a worst-case load, albeit with middle-of-the-road sensitivity (104dB/mW). To take them to eardrum-bursting 124dB will require 2V, so anything past that is gravy for me. With the One d, I had no problem playing music far louder than I found comfortable. Lower sensitivity headphones at similar impedances will require proportionately more drive, so when I substituted a friend's Grado SR325s, I had to raise the volume a bit, but everything was still clean as a whistle at high volume. I was unable to borrow a real challenge like the notorious Audeze LCD-X, but I would want to try that combination before buying. For more mainstream headphones, there's no lack of driving ability.

When it comes to electronics, I have found that the audible differences are usually so minor that they can't easily be detected with the interference of changeover time between one component and another. So these days, if I think I hear a difference, I'll try to confirm it using a high-resolution capture of the output to create a WAV file and compare it to the original signal with essentially instant switching using the ABX facility in DeltaWave or foobar (see Resources).

Comparing the One d to my reference DAC/headphone amp (an RME ADI-2fs b), I thought it might sound a trifle soft. Comparing WAV files via ABX, I scored 10/12 correct identifications—statistically significant, but pretty subtle. As will be shown shortly, this very slight tonal aberration will be expected to be headphone-dependent, so I would not characterize it as inherent or a means to tame any excessive treble from other headphone models.

I had a lot of fun with the Matrix control. Interestingly, with recordings that were single-point or otherwise simply miked, the effects were very subtle. For example, "Hey, Mr. Mumbles" from Clark Terry's *Live At The Village Gate* (Chesky) barely altered subjectively until I reached nearly the maximum crossfeed setting. By contrast, songs like "The Jimi Hendrix Experience" (Exp) from *Axis Bold As Love* (Reprise) or the Beatles' "You Never Give Me Your Money" from *Abbey Road* (Universal—Apple Corps) went from almost unlistenable to absolutely delightful.

The images didn't sound to me like they came from loudspeakers in front, but they didn't give me the stuffed ears sensation. And instead of the hard left-right sounding like a confusing interrogation, the image was brought in subjectively closer to the sides of my head. This feature is, in my opinion, a very worthwhile one. I should mention that although

this is frequency response shaping, I did not get any sense of the tonality changing with the crossfeed level. This is a problem with some crossfeed setups, and its absence is tribute to SPL's choice of response shaping.

It should be noted that you can get plug-ins with various crossfeed features, so although the implementation here is analog electronic, you can achieve the same ends in the digital domain.

Measurements

As usual, my measurement setup comprises an Audio Precision APx-525 analyzer with an APx-1701 acoustic interface. This is backed up with a Hewlett-Packard 3466A digital multimeter, a Vicnic 1kHz precision sine wave source, and a homemade twin-T notch filter for ultra-low distortion measurement.

I started by running the One d through some basic measurements. Using the single-ended analog input, the gain to the balanced line outputs was 5.89dB and 5.91dB for left and right channels, respectively. Gain was similar using the balanced analog input. The DAC output to the fixed level balanced output was 4VRMS at 0dBFS.

The really odd measurement was the output impedance. It was relatively high at 20Ω, but flat across the audio band. I asked SPL about it to confirm I didn't do something stupid or that there was a product defect, and they felt that the number should be much lower. SPL was kind enough to send an AP plot showing source impedance under 1Ω,

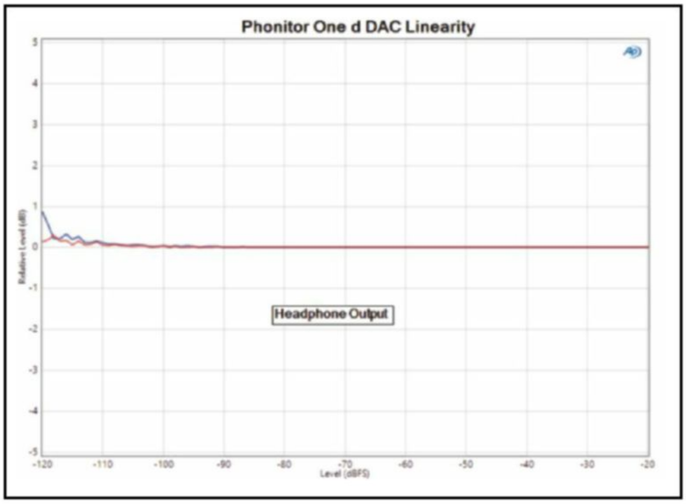

Figure 4: Used as a DAC, the One d's linearity is excellent even at very low levels.

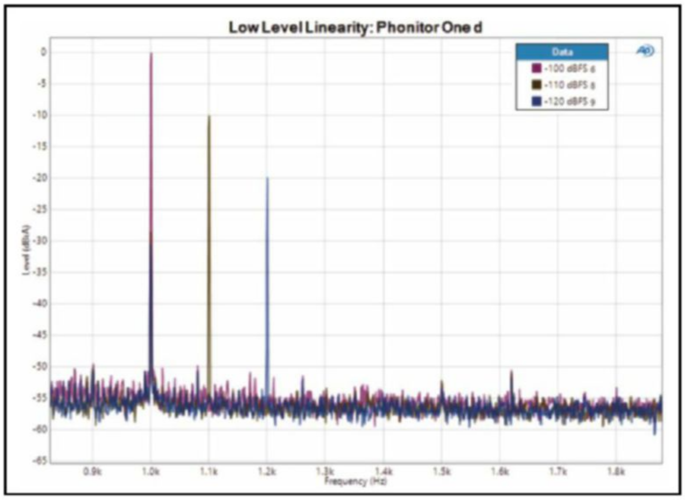

Figure 5: Ultra-low-level signals are reproduced at exactly the correct level, an excellent DAC implementation.

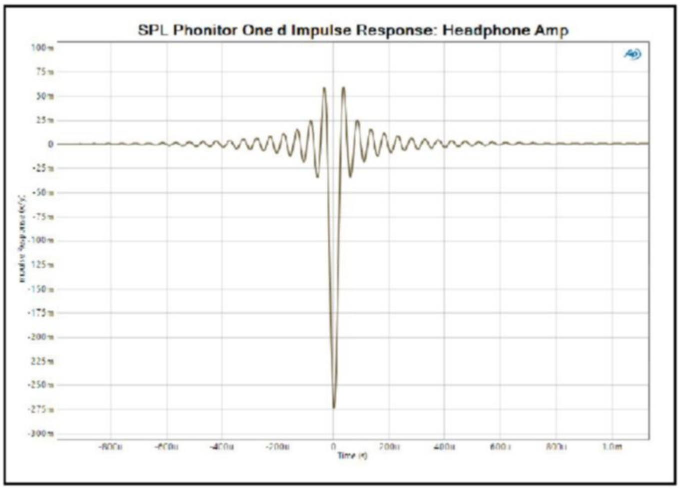

Figure 6: The One d's impulse response appears to be linear phase.

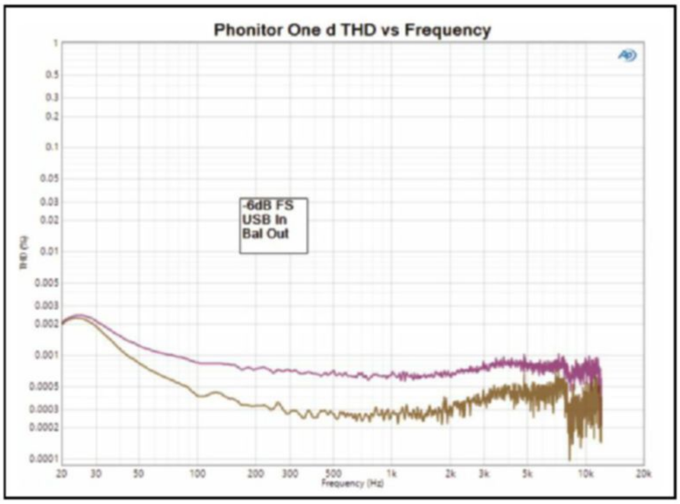

Figure 7: The distortion of the One d's DAC remains low across the audio band.

but no details about the test setup. I measured it three different ways (using the AP Impedance Utility, using a voltmeter with the headphone amp at a fixed volume setting and changing the load resistance, and by attaching a reactive headphone load and seeing the frequency response variations) and got consistent results from each method. I am confident that my measurement is accurate. In general, you want the source impedance to be very low compared to the load impedance, so this was the first suggestion to me that the One d is more suitable for higher impedance headphones.

The first job of any headphone amplifier is to drive headphones! **Figure 3** shows the drive capability of the One d into different impedances. For higher impedance loads, the One d is capable of sourcing about 8V before the distortion rises above 1% or so. Lower impedance loads are a bit trickier since the source impedance will limit the drive voltage via a divider effect. In this case (33R), the knee in the distortion versus drive voltage occurs between 500mV and 1V.

Although the One d's intended function is driving headphones, it also can function as a DAC. One important basic measurement for DACs is linearity. **Figure 4** shows that the One d's DAC is quite linear to better than -115dBFS, an excellent performance. Another way to visualize this is playing fixed low volume tones and seeing what happens as the level is decreased.

Figure 5 illustrates the output at -100dB, -110dB, and -120dB relative to full scale output. I changed the frequency a bit for each measurement in order to display the levels more easily. As can be seen in this graph, the tracking at these very low levels is essentially perfect.

The frequency response at both the balanced outputs and the headphone output was flat to within better than ±0.1dB. The impulse response is shown in **Figure 6**, which is typical of linear phase brick-wall filters. That type of filter is the best choice to minimize images and aliasing. A sweep of distortion versus frequency at -6dBFS is shown in **Figure 7**. Although it rises a bit at low frequencies, where hearing and system response are most forgiving, it remains low and the worse of the two channels is still better than 0.001%, indicating good engineering in the analog sections.

Just out of curiosity (not at all relevant to real-world use, but indicative of the engineering), I looked at the spectrum of an 80kHz signal at 2.6V from the balanced line output using the available 768kHz sampling frequency (**Figure 8**). Although not as beautifully clean as the performance in the audible band, the distortion still remained better

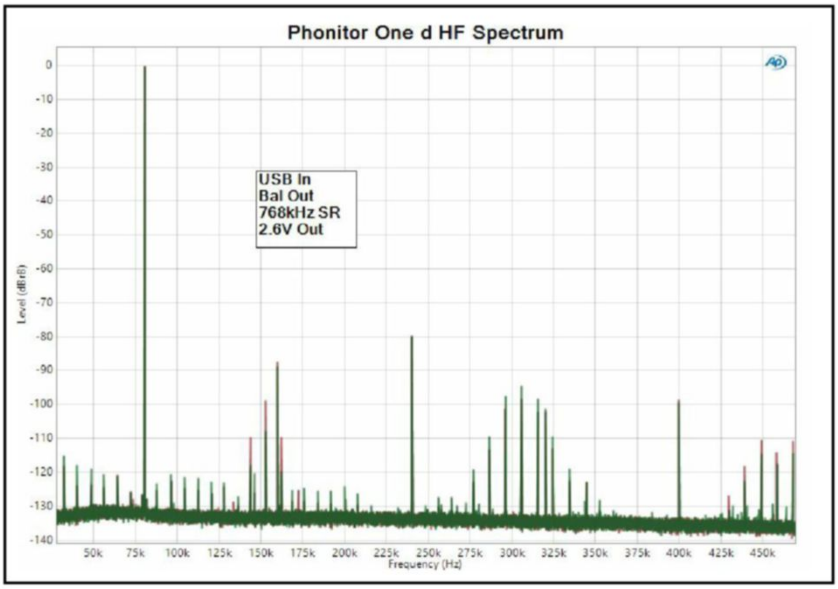

Figure 8: An 80kHz ultrasonic sine wave reproduced with a 768kHz sampling rate has distortion and noise components below -80dB (0.01%).

than -80dB (0.01%). And it shows good performance from the circuitry at the highest sample rates relevant only to bats.

The Matrix circuit is clearly of interest here. I made several measurements of its functioning by injecting signals into one channel, then looking at the output of both. **Figure 9** shows the time-domain output from a sine wave injected into the left channel (blue) and the right channel (red) undriven, with crossfeed at the fully on position.

Besides the reduced amplitude of the crossfed signal, the delay is also evident, a bit under 0.1ms at this frequency. Sweeping the frequency, the interchannel phase is shown in **Figure 10** at 50% and full crossfeed. The turnover frequency increases with increasing crossfeed, but in both cases, the asymptotic phase shift is 90°, giving some clues as to the reactive crossfeed circuit.

Similarly, the frequency response with the crossfeed at 50% (dashed lines) and 100% (solid lines) is shown in **Figure 11**. Note that at full crossfeed, there's less than 3dB difference between main signal and crossfeed below 1kHz or so.

Also note that the spectral balance is maintained by increasing the main channel's treble to compensate for the treble roll-off in the crossfeed. As I observed earlier, to my ears, this is done very successfully.

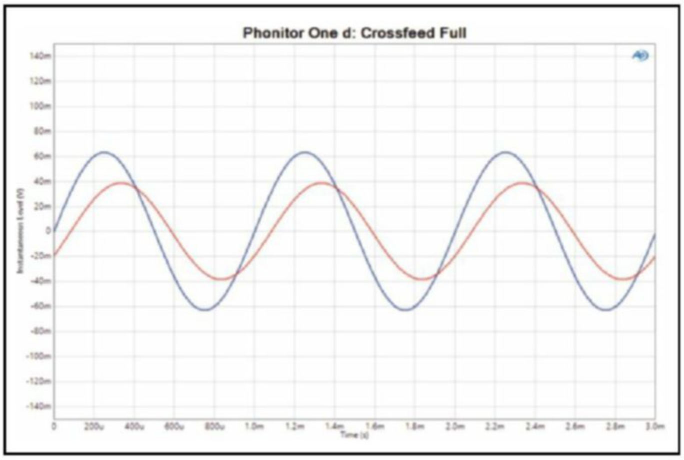

Figure 9: The One d gives this time domain response when fed a signal only into the left channel. The Matrix crossfeed circuit produces a reduced amplitude and delayed signal into the opposite channel.

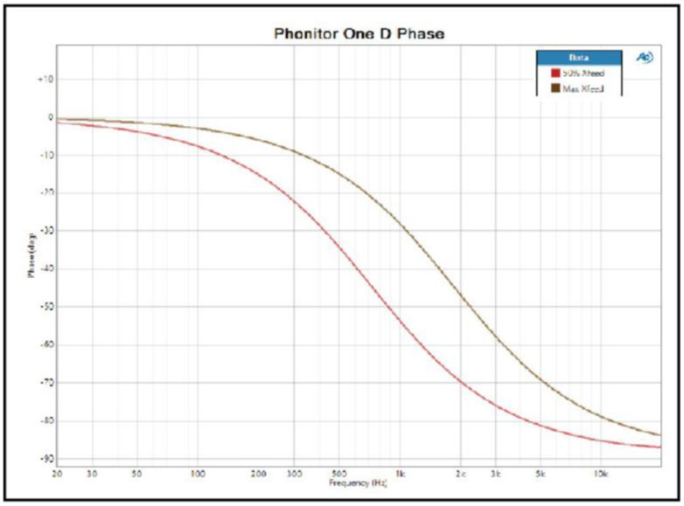

Figure 10: The turnover point for the interchannel phase increases with increased crossfeed.

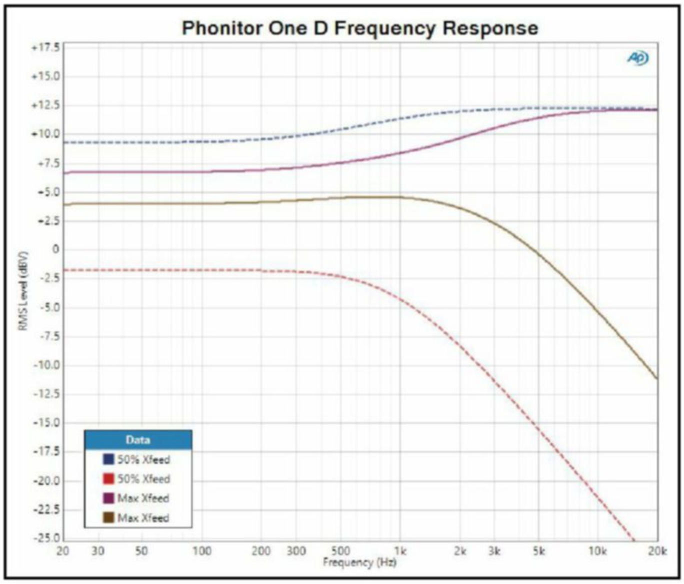

Figure 11: The frequency response of the main signal is changed to compensate for the spectral balance of the crossfeed signal (50% crossfeed is dashed lines, full crossfeed is solid lines).

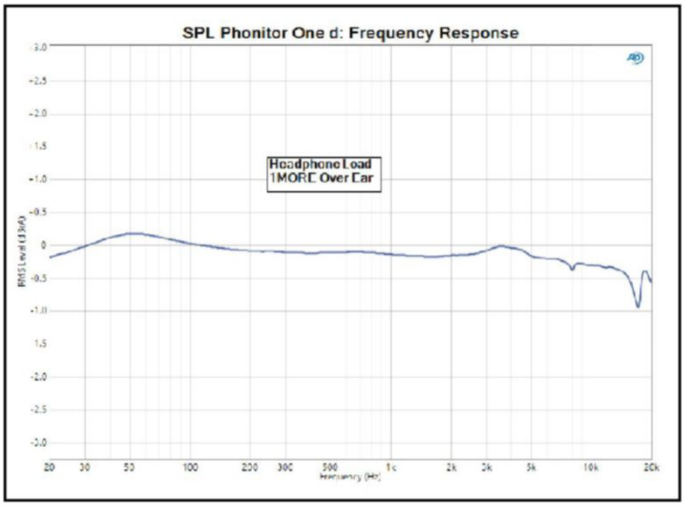

Figure 12: The source impedance of the headphone amp interacts with the reactive load impedance of headphones to slightly alter the frequency response.

Finally, I checked the frequency response into a real-world load, specifically my headphones. This is shown in **Figure 12**. The change in frequency response from flat is due to the voltage divider effect between the One d's source impedance and the non-constant impedance of the headphones as a function of frequency. This correlates with the very slightly audible softening in the treble that I observed. Of course, with other headphones, this response will look quite different, and for 300Ω nominal headphones, the frequency response aberrations are likely to be well below a human's detection threshold.

Wrap-Up

The Phonitor One d is a solidly put together unit offering an excellent DAC. The headphone amp portion works essentially perfectly for loads of 300Ω or higher, but its relatively high source impedance limits the drive to low impedance headphones, and will cause the unit to have a load-dependent sonic signature with low impedance headphones. Even so, distortion performance is well below any audible detection thresholds, and the available drive is sufficient to take most headphones to ear-splitting levels.

The Matrix circuitry is very effective at removing the "stuffed ears" effect in recordings where the production is enthusiastic about pan-potting, but falls short of the ideal of removing in-head localization. The ASIO drivers are stable and bug-free.

At $799, it is not inexpensive, but the overall performance is excellent and the construction and track record of SPL promises high reliability. If I were an avid headphone user, the One d would be on my short list of DAC/amplifier options.

Resources

B. Bauer, "Stereophonic Earphones and Binaural Loudspeakers," *Journal of the Audio Engineering Society* (JAES) Volume 9, Number 2, pp. 148-151; April 1961.

"DeltaWave Audio Null Comparator," https://deltaw.org

J. Didden and S. Yaniger, "The Smyth Realiser A8," *Linear Audio*, Volume 7, pp. 199-206, March 2014.

"foobar2000 ABX Comparator," www.foobar2000.org/components/view/foo_abx/releases

S. Linkwitz, "Improved Headphone Listening," *Audio*, pp. 42-43, December 1971, www.linkwitzlab.com/headphone-xfeed.htm

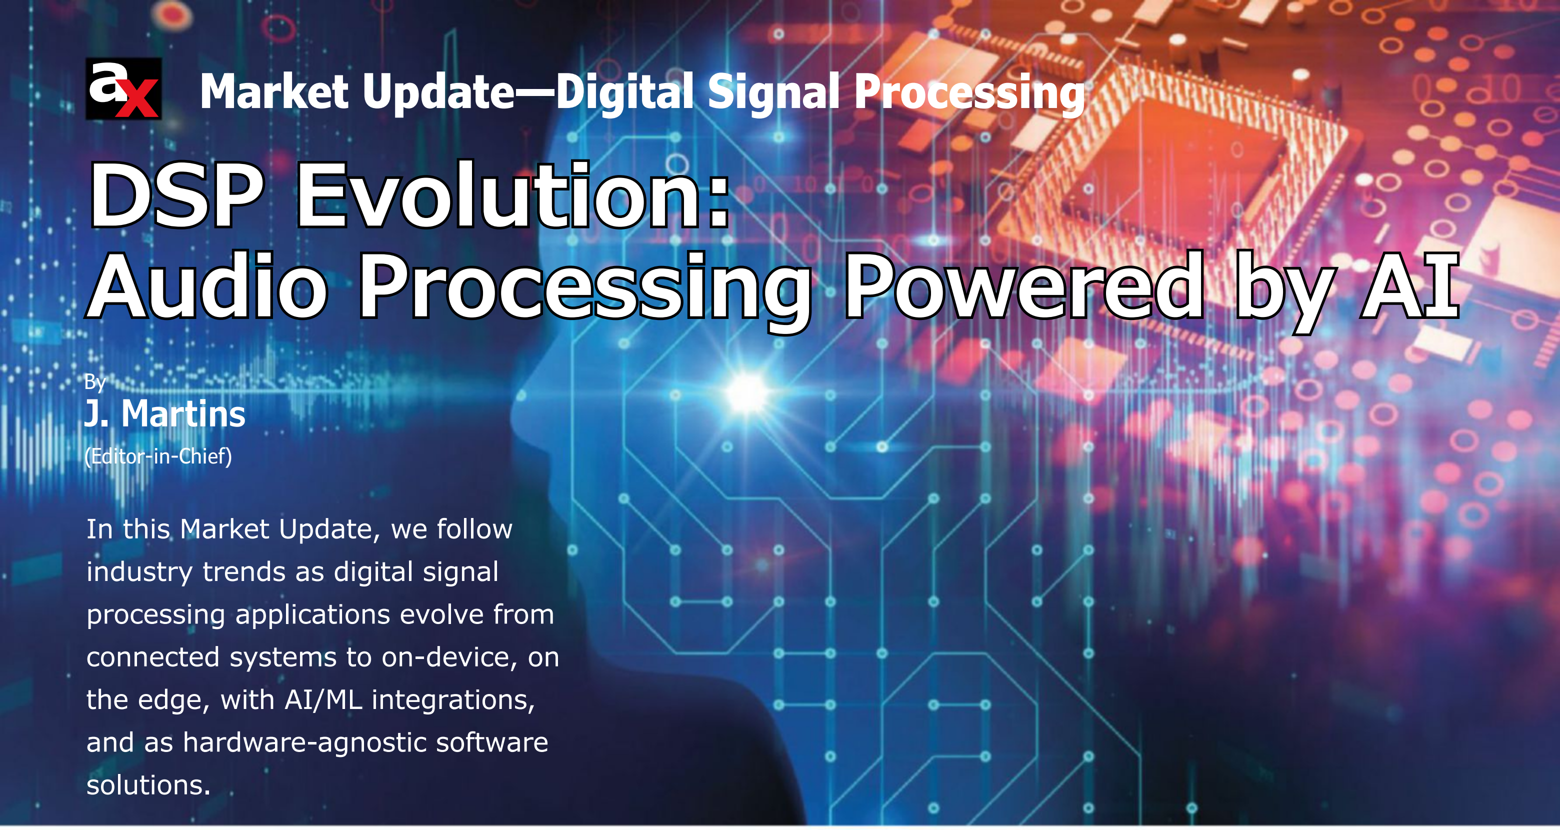

Market Update—Digital Signal Processing

DSP Evolution: Audio Processing Powered by AI

By
J. Martins
(Editor-in-Chief)

In this Market Update, we follow industry trends as digital signal processing applications evolve from connected systems to on-device, on the edge, with AI/ML integrations, and as hardware-agnostic software solutions.

I believe the most interesting story when it comes to audio digital signal processing (DSP) in 2021 is the introduction of Edge-AI platforms and how those are changing market dynamics, inspiring mergers and acquisitions that promise to redefine the audio industry landscape.

Underlying that trend, we have two other major technology shifts. The transition toward mobile edge-computing in all product categories, particularly in wearables and hearables, and the cutting-edge research and development being conducted in the field of voice recognition, voice interfaces, and digital assistants, using artificial intelligence (AI) and machine learning (ML).

DSP applied specifically to audio signal processing plays a fundamental part in all of those technology shifts, and is increasingly fully system-integrated. It's all about how DSP is no longer dependent on a series of pre-programmed fixed functions, and how the processing operations are now controlled dynamically by external input (sensors). Controlling a DSP chip with a microcontroller is a requirement in connected devices and gradually software processes were transferred to the MCU itself.

Audio signal processing with DSP today is widespread in mainstream consumer products, and with mobile and wearable devices, DSP features that were previously user-controlled, personalized options, are now increasingly sensor-based. That is until the volume of decisions are just too many for a user to control, or fixed functions are not effective under continuously changing variable conditions. Then, data-dependent operations will become more effective, if controlled by machines that are able to learn, and eventually even use artificial intelligence (AI) inference.

The recent introduction of AI-based edge processors is enabling a new field for DSP to be implemented in a much more effective and useful manner. Allowing the processing to be managed and orchestrated within an integrated system, closer to the data source, and quickly responding to data input.

To better understand the concept, let's look at the example of the Knowles AISonic "mobile-centric audio edge processor," as the company defines it. This is a completely new (and for some, surprising) field for a company that is mainly recognized for its microphones, transducers/microspeakers, and sensors. In recent years—and publicly since 2017—Knowles has evolved toward the integration of audio processing, and this evolution is now visible with the introduction of the Knowles AISonic range of processing solutions—the IA8201 Dual Core, and IA8508 Quad Core—specifically designed for advanced and power-efficient audio and machine-learning applications. AISonic solutions place signal processing at the center of integration, allowing robust voice activation and multi-microphone audio processing with the compute power needed to perform sophisticated audio signal processing, such as required by context awareness, gesture control, and other sensor-responsive operations.

This doesn't make Knowles a "DSP" company per se, but it places DSP integration at its core. As an example, the AI8201 includes a 128-bit core (DMX) with a Knowles proprietary instruction set and a Tensilica HiFi3 core (HMD), both with Knowles audio enhancements. The DMX is a four-way floating-point SIMD processor targeted toward efficient, high-performance computing, that is ideal for advanced audio processing in beamforming, barge-in, Acoustic Echo Cancellation (AEC), and more, while the HMD is targeted toward efficient, low-power, wake-on-voice applications with a two-way floating-point SIMD processor. Both cores contain dedicated accelerators for FFT, peak finding, and Deep Neural Networks.

The quad core IA8508 platform adds three heterogeneous Tensilica-based, audio-centric DSP cores and an ARM Cortex M4 for general purpose controller capability. This core "brain" is integrated with a rich set of audio, and general-purpose, high-speed interfaces for interfacing with digital microphones and sensors, and a host for further processing.

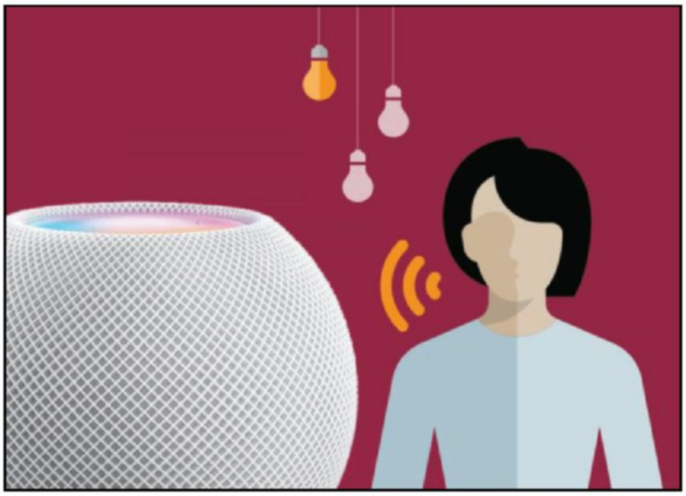

With voice interfaces and smart speakers, audio DSP applications now include real-time, adaptive, and even predictive processing, spurring the use of machine learning and artificial intelligence as the supporting engine for DSP.

The AISonic platform in itself provides a robust ecosystem for custom and third-party solutions, and Knowles promotes an open DSP partner program, which enables leading contributors to offer solutions for a variety of use-cases. Partnerships with software and algorithm companies (e.g., Waves, Sensory, Alango, and Chatable, among others) have been already announced.

Since I mentioned Tensilica, it's also relevant to see this evolution from the product development perspective. Cadence Design Systems recently unveiled its latest Tensilica FloatingPoint DSP family, optimized for power, performance and area (PPA).

With mobile, automotive, and consumer markets in sight, Cadence optimized the new DSP IP cores from small, ultra-low power to very high performance, as required today for a broad array of applications, from energy-efficient solutions for battery powered devices, to AI/ML, sensor fusion, object tracking and augmented reality.

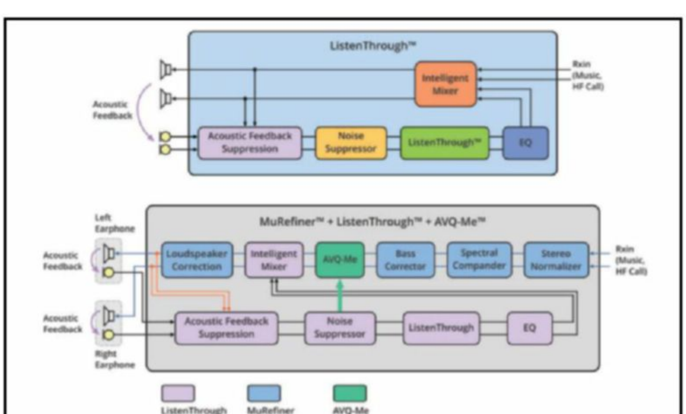

Alango Technologies licenses its advanced digital signal processing technologies for voice and audio applications, including the ListenThrough environmental audio awareness technology, which can be integrated with the company's own Audio Enhancement Package audio enhancement solutions.

The new Cadence Tensilica family consists of four cores: the Tensilica FloatingPoint KP1 DSP, the Tensilica FloatingPoint KP6 DSP, the Tensilica FloatingPoint KQ7 DSP, and the Tensilica FloatingPoint KQ8 DSP. The technical details for these products are out of the scope for this market update, but suffice it to say that these are perfect examples of solutions that fill today's requirements of energy-efficient, cost-effective and high-performance DSPs, designed specifically for floating-point-centric computation.

DSP for ANC

Digital signal processing applied to audio has seen innovative uses introduced, many of which are now generalized. This is directly related to the new product categories and user applications introduced in just the past five years. Active noise cancellation (ANC) is an excellent example. It's been around for quite a while, but what we have perfected in audio products since the introduction of true wireless earbuds, for example, is a totally different level of systems architecture, serving what is essentially advanced audio processing matched with the latest generations of sensors with direct microcontroller integration into the overall electroacoustic design.

And, ANC is evolving toward new use cases in virtual acoustics and automotive audio, which means it will expand into less power-constrained environments, where more robust computing power can be explored. Before we will be able to see the benefits in mainstream home audio, we are likely to see this implemented in the car, given the large investments that are being channeled in that direction.

Most of the progress achieved in very sophisticated DSP is seen as gradually incrementing, until some product success stories become obvious, such as was the case with ANC applications.

Since the Bose QuietComfort headphones were introduced to the market and adopted early by frequent flyers, triggering a whole new product class in headphones, ANC became a sought-after mainstream feature. Introduced only at the end of 2019, Apple's AirPods Pro brought ANC to a whole new level and inspired a complete transformation in the true wireless category, even inspiring a model architecture that is now powering an entire wave of hearables and wearables. Under $100 USD TWS earbuds offering effective ANC are now available in the market. The big difference is the AirPods Pro benefits from a dedicated SoC, with dedicated DSP and CPU cores that are able to orchestrate the whole system in innovative ways, and are even able to adjust to different environments as well as evolve with new algorithms and software updates.

From basic audio signal processing designed to cancel continuous airplane engine noise, implemented and used in products for pilots from Bose and Sennheiser in the 1980s, to consumer products used by millions while traveling, commuting, exercising, or to improve concentration while studying or to avoid noise exposure at work, active noise cancellation has significantly evolved. Today, it is a part of a sophisticated system that enables audio enhancement and augmentation, and more importantly, adaptive audio processing, powering improved communications, voice recognition, and digital assistant interactions.

ANC, together with spatial audio processing, are exciting fields in audio DSP applications. The transition from adaptive (reactive) processing to intelligent (predictive) processing—or as the industry now likes to call "smart" processing—is the frontier that is now inspiring the use of ML and AI as the supporting engine for DSP. The transition from pure computational, algorithmic, and logical instructions to deep neural networks, ML and AI also conveys an important systems architecture transition.

It's important not to be confused or scared by those terms. All of them are in fact just extensions of computations required in adaptive systems that essentially emulate human reactions, but when implemented in systems that also act faster than humans can, are able to process data from sensors faster, and are even able to anticipate conditions.

As discussed by Michael I. Jordan, a leading researcher in AI and ML at the University of California, Berkeley, machine learning is not about the imitation of human thinking but serves to augment human intelligence, via painstaking analysis of large data sets in much the way that a search engine augments human knowledge by organizing the Web (see Resources). That is what we are doing with current "intelligent" or smart processing systems—augmenting human intelligence, which is a step toward AI.

The realization of how fast things are evolving and that we are finally starting to see real AI technology implementations comes from the fields of voice recognition and human-machine voice interfaces. Those systems, as Jordan defined them, involve high-level reasoning and the ability to form the semantic representations and inferences of which humans are capable. This is molding strategic decisions by some of the largest industry players.

Recently, I noted that STMicroelectronics acquired Cartesiam, a software company based in Toulon (France), founded in 2016, and specializing in AI development tools, enabling ML and inferencing on ARM-based microcontrollers.

The Cartesian team includes data scientists and embedded signal processing experts, with significant experience in delivering standard and custom solutions. Its flagship and patented solution, NanoEdge AI Studio, allows embedded systems designers without prior knowledge in AI to rapidly develop specialized libraries integrating machine-learning algorithms directly into a broad range of consumer and industrial applications.

With this acquisition, STMicroelectronics reinforces its AI strategy and technology portfolio. The NanoEdge AI Studio solution is fully complementary to STMicroelectronics' STM32Cube.AI toolset, which offers the ability to map and run pre-trained artificial neural networks on STM32 microcontrollers, creating a platform that is ripe for innovation in audio applications.

But let's leave pure AI and look at two companies working in this field and how that research relates to DSP and more specifically to audio processing.

Voice AI

Kardome is a company from Israel working in the pure domain of speech and voice processing. While speech recognition enables a computer to recognize, understand, and translate human speech into text, using natural language processing and ML, voice recognition also applies to a machine's ability to identify a specific user's voice. Kardome combines audio processing, ML , and AI to improve voice recognition and voice user interfaces, exploring the full potential of these technologies.

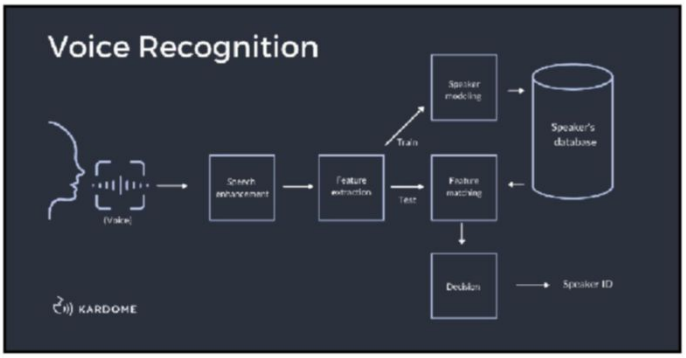

Kardome speech and voice processing solutions depend upon the company's implementation of advanced audio processing, including AI-driven signal separation and noise reduction.

The company combines an experienced team in acoustics, signal processing, and ML that is now leveraging advanced AI algorithms to develop a solution that makes individual voices understandable for both speech recognition engines and human listeners. And Kardome is specifically exploring applications in environments with difficult acoustics, such as the inside of the car.

While focusing on voice user interface design technology, Kardome addresses the complexity of acoustic environments and improving voice signals in general. The company develops AI-driven signal separation and noise reduction technology to facilitate voice interactions. And its patented technology simultaneously enhances speech signals from multiple speakers in real-time, isolating the target speech.

As the company explains, the key differentiator between Kardome and other audio front-end solutions lies in the fact that its technology clusters/aggregates speech signals based on location. Kardome's source separation software treats each person in any environment as if that person is the only person talking. This focus contrasts with direction-based technology (e.g., beamforming).

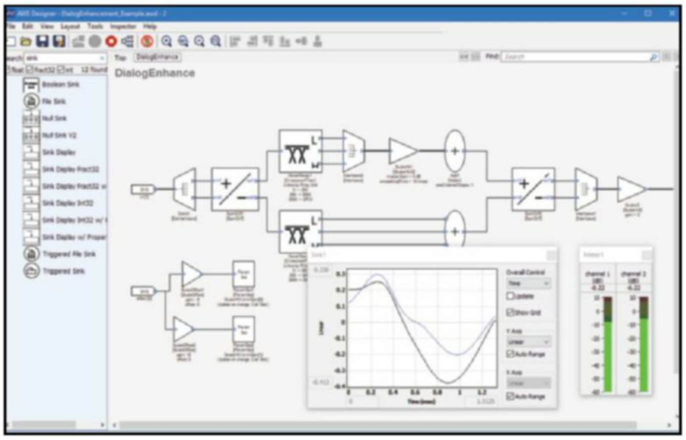

In Audio Weaver, DSP Concepts offers four basic output processing blocks, including improving the sound field of stereo playback, extending the perceived bass response, boosting dialog intelligibility of speech in mixed content (pictured), and controlling the dynamic range of content in real-time.

Yamaha's latest flagship and truly unique YH-L700A headphones introduced an advanced ANC implementation with a series of smart optimizations, resulting from the combined microphones, sensors, and DSP circuits built-in. The headphones continuously monitor headphone air leakage and fit, compensating for more or less passive isolation and adjusting the frequency response to the changes—every 20 seconds. The DSP processing in the YH-L700A also powers a 3D Sound Field feature, which emulates a spatial audio mode from standard stereo audio.

In a closed environment (e.g., automotive applications), Kardome's technology enables multiple people to give clear, understandable commands. As the company details, its core technology includes Speech Separation and Noise Reduction (SSNR) modules that facilitate reliable Automatic Speech Recognition (ASR) performance in noisy and multi-speaker scenarios. In summary, providing the "smart" software components to audio front-ends, with algorithms for localization, source separation, and noise reduction.

Taking this technology a step further, we find companies such as Canadian company Fluent.ai, which is also focused in speech understanding and voice user interface solutions based on ML and AI. The difference is that Fluent.ai's speech-to-intent technology employs unique neural network algorithms to directly map the incoming speech of a user to that person's intended action without the need to perform speech-to-text transcription.

Dutch lifestyle brand Fresh 'n Rebel introduced the Clam Elite ANC headphones, a large circumaural design, featuring Audiodo' Personal Sound hearing profiles. The Audiodo software libraries and algorithms are built into the Bluetooth chipset of the headphones so that the audio customization remains permanently available for the user, and the signature sound of the headphones will be retained.

Fluent.ai uses trained models that directly associate semantic representations of a speaker's intended actions with the spoken utterances. In a way, those models are based on the concept of vocabulary and language acquisition in humans. Unlike conventional automatic speech recognition (ASR), Fluent.ai's technology does not require phonetic transcription and the models are able to recognize a new language from a small amount of data, and enable end-users to interact with the devices in a language of their choice. Also, users are able to choose words of their preference.

As the company explains, spoken language understanding (SLU) without highly constrained conditions is the next frontier for AI technology. In current voice commands, users must say the exact words or phrases for which the system has been trained in order to guarantee a high recognition accuracy. With natural language interactions, these systems are able to deal with several variations of each intent or command. Fluent.ai is creating end-to-end SLU applications, directly extracting the intent from speech without converting it to text first.

More recently, Fluent.ai launched its Automatic Intent Recognition (Air) engine, a direct speech-to-intent spoken language understanding system, which offers fully offline, accurate speech understanding. Running completely on-device, this solution ensures privacy and is ideal for natural voice interaction in any language or combination of languages with appliances, industrial, or navigation systems.

Fluent.ai's approach enables users to design speech understanding models that are much smaller in size, yet provide high accuracy even in noisy environments. Fluent.ai systems are capable of recognizing up to thousands of intents in a small model size of hundreds of kilobytes. Furthermore, Fluent.ai's ability to build multiple languages into a single model means that users can seamlessly switch between languages when interacting with their devices, without the need to configure language settings in between. This has enabled the company to expand its product and business development plans after introducing the first successful implementations in connected home devices and factory robot automation.

Fluent.ai Air offers OEMs and device manufacturers the technology for developing their own faster and robust voice user interfaces, with the ability to learn from users' preferences and feedback. Of course, these systems are dependent on the quality of the voice signals provided by audio front-ends, and in those developments Fluent.ai partners with companies such as DSP Concepts, XMOS, or CEVA, and supports a wide range of hardware platforms.

This gives a broad overview of the power contained in the systems used for such applications. Both Kardome's and Fluent.ai's solutions run on ARM Cortex platforms and many digital signal processors (DSPs). Fluent-ai is porting their algorithms to AI chips with specific neural network accelerators because the frontiers between these architectures are clearly blurring.

Fluent.ai's suite of speech-to-intent technologies has been ported and optimized for CEVA's low power audio and sensor hub DSPs, providing a high performance, robust solution for wearables and consumer devices in general. CEVA's DSPs, including the CEVA-X2, CEVA-BX1, CEVA-BX2 and SensPro family, enable the full suite of speech-to-intent technologies to seamlessly run in always-on mode. These same DSPs can also run other software and algorithms that further enhance the performance and feature set, including CEVA's ClearVox front-end noise reduction, MotionEngine for sensor fusion, and the SenslinQ framework for contextual awareness.

CEVA's audio and sensor hub DSPs are optimized for sound processing applications ranging from always-on voice control up to multiple sensors fusion. They have been specifically designed to tackle multi-microphone speech processing use-cases, high-quality audio playback and post-processing, and on-device sound neural network implementations. In addition, CEVA offers a large third-party ecosystem of audio/voice software, hardware and development tools.

Multiprocessing

As we wrote in previous Market Updates on this topic, the technology transition that is taking place to support the exponentially more complex calculations required by AI and ML, is also now enabling systems to support the most common audio DSP operations, such as equalization and bass enhancement, filters, adjustable crossovers, signal delay, signal summing/distribution, and more. All these foundation blocks become available in real time and at the highest level. User adjustments and even some complex algorithm optimizations can be remotely performed, with the support of a tablet or smartphone and a connection to the cloud, but once defined at the local level, they can be flawlessly performed without latency and freely adjusted in real time by the user.

What was once the realm of dedicated DSP chips, ASICs, or FPGAs, is now increasingly available in software that is able to run in any type of platform for embedded audio development. That is the case with DSP Concepts' Audio Weaver and one of the reasons why its use has exponentially increased across multiple industries and applications. Audio Weaver is purpose-built for audio processing and offers a comprehensive audio development framework with hundreds of ready-to-use modules. More importantly, it enables design and development without hardware, immediate evaluation of audio results, and deployment on the hardware of choice for the application. Deployment models extend from ARM Cortex-M4 to Intel processors, Cadence Tensilica, CEVA, and ARC-embedded cores, to Texas Instruments DSPs, and Analog Devices SHARC processors.

In fact, DSP Concepts even supports multicore processing for multiple hardware platforms, from Texas Instruments to Analog Devices or ARM-based solutions and more recently launched AWE Core OS, a new variant of the audio-processing engine powering its Audio Weaver platform, optimized for multicore development. Among other possibilities, AWE Core OS enables moving audio processing between cores and even between engines in the SoC. This is key for automotive audio development, but as mentioned, could also be strategic for key emerging applications in adaptive and virtual acoustics, personal audio zones, and spatial audio processing in home audio systems.

We also discussed that evolution and transition in our 2020 Market Update on DSP Products and Applications—recommended reading as a preamble for this update.

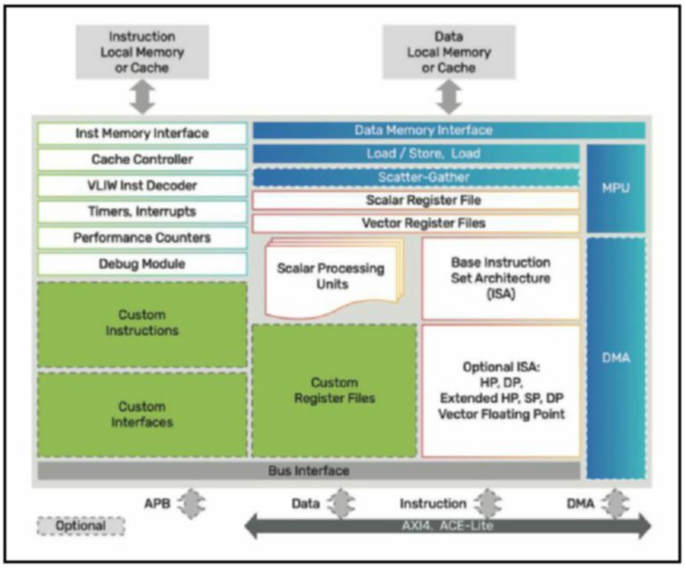

A key player in the DSP space, Cadence Design Systems recently unveiled its new Cadence Tensilica FloatingPoint DSP IP cores, which extend from small, ultra-low power to very high performance in a broad array of applications. This is the block diagram of Tensilica FloatingPoint KQ7 and KQ8 DSPs.

Adaptive Software DSP

As applications require being able to support full audio processing on-device, increasing systems integration allows developers to create solutions that are able to adapt to variable conditions. This is key in ANC, and also in voice recognition and communications implemented in personal wearable designs, as an example.

An interesting processing tool for those designs is the ListenThrough "smart algorithm" for headphones and the new generation of true wireless earbuds and hearables, offered by Alango Technologies. This is a software solution that was specifically designed for those portable devices, adding to both passive and active isolation from the surrounding environment, to allow important sounds to be passed through.

Alango ListenThrough complements music playback with environmental alertness, by mixing ambient sounds into the Bluetooth music stream. The algorithm suppresses background

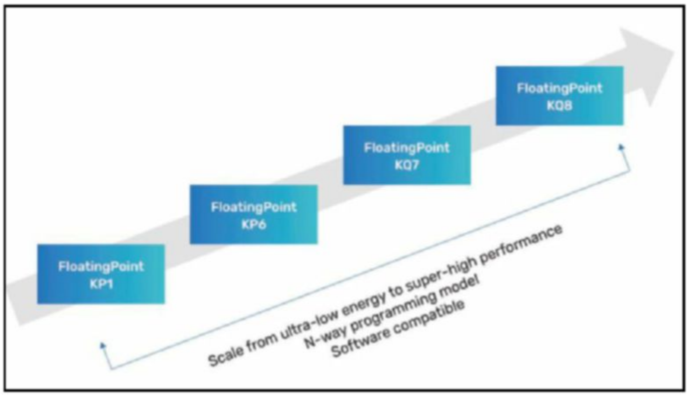

The new Cadence Tensilica FloatingPoint DSP family is comprised of four cores: the Tensilica FloatingPoint KP1 DSP, the Tensilica FloatingPoint KP6 DSP, the Tensilica FloatingPoint KQ7 DSP and the Tensilica FloatingPoint KQ8 DSP.

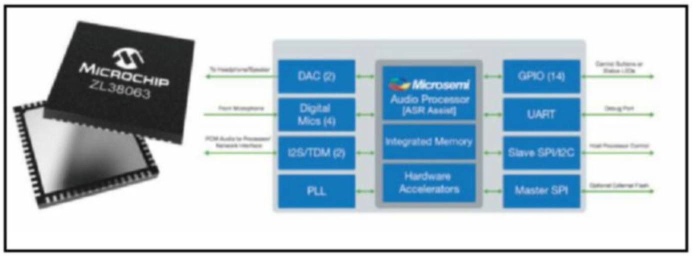

With the acquisition of Microsemi, Microchip now offers the ZL38063, the most powerful chip in the Timberwolf audio processor family, which combines noise reduction, stereo echo cancellation, and de-reverberation with support for up to four digital mic inputs, and two internal 16Ω speaker drivers.

noise while passing through sound events (e.g., an alarm, horn, train, or talker) that exceed the preset background noise level. This is achieved with low latency and without content analysis of the environmental sounds—an area also being addressed by other companies.

In parallel, Alango implements a solution called Intelligent Mixer, which receives a stereo audio signal and the signal from the ListenThrough block, and mixes them according to the respective amplitudes. The Intelligent Mixer ensures that sound-events are mixed at suitable levels for the user. Acoustic feedback technology is implemented since the sound-events are amplified by as much as 30dB during loud music playback. The solution also combines Acoustic Feedback Suppression (AFC) to eliminate feedback loops between the transducers and the microphones, noise suppression that automatically detects and attenuates stationary and transient noises, and an equalizer for adjustment of microphone frequency response with high-frequency resolution.

This solution can be integrated with other Alango Technologies, including MuRefiner, a set of integrated audio enhancement DSP technologies, and Automatic Volume and Equalization control for

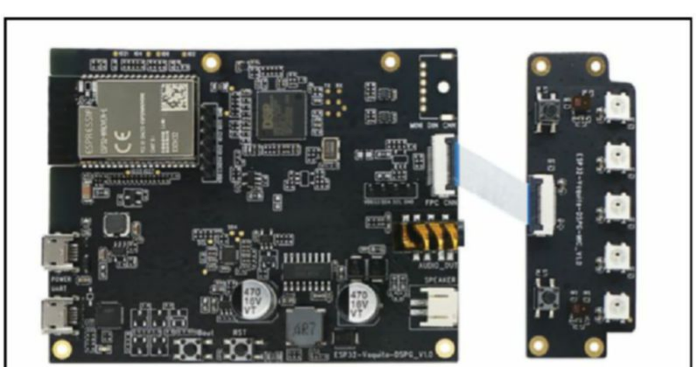

Espressif Systems offers a range of ESP audio development boards, powered by ESP32 and ESP32-S2 wireless SoCs, for applications (e.g., smart speakers and smart-home solutions). Pictured is the ESP32-Vaquita-DSPG solution with advanced Wi-Fi + Bluetooth dual-mode, DSP Group's DBMD5P audio SoC, and a two-mic array with 360° pickup, ready for Alexa Voice Service (AVS) integration.

Music and Entertainment (AVQ-Me) technology. AVQ-Me amplifies and equalizes music content according to ambient noise characteristics providing equal perceptual loudness, intelligibility, and sound coloration in conditions where noise level and spectrum dynamically change (e.g., in mobile conditions).

This is clearly a set of apparently conventional DSP functions that are being combined and applied to systems that are now much more sophisticated in design.

Among other platforms, Alango offers these solutions, together with its extended Voice Communication Package (eVCP), a universal software package of digital signal processing technologies for voice applications, running on DSP Group's DBMD7 SmartVoice processor.

These multicore DBMD7 processors are a part of the SmartVoice family of audio, voice AI, and edge AI processing solutions from DSP Group and are a highly scalable, low power, and cost-effective platform to bring always-on voice control and voice call performance to power-sensitive applications. They are used by developers with Google TensorFlow, Audio Weaver, and now also with the effective processing solutions from Alango.

Specifically for voice processing, Alango's eVCP includes echo cancellation, optional dual- or multi-microphone beamforming, noise reduction, equalization, and gain/signal level adjustment. The signals are combined in eVCP's intelligent mixer block based on the respective background noise, signal quality, and voice activity in multiple frequency bands of each signal—resulting in a clean, noise-free, echo-free signal with the desired voices enhanced.

The DBMD7 processors support up to eight microphones and together with Alango eVCP offer a computationally efficient solution with large on-chip memory to cover a wide range of applications and microphone configurations. And the available SDK for this range of processors enables developers to cost-effectively deploy high-performance ML inference at the edge to avoid network latencies, minimize power consumption, ensure end-user privacy, and free up scarce network bandwidth.

DSP Group and Synaptics

The DBMC2-TWS advanced adaptive processing audio codec for true wireless stereo (TWS) headsets from DSP Group is another example of the company's latest generation of audio-centric solutions. To dynamically configure settings and filter parameters on-the-fly, the DBMC2-TWS integrates a DesignWare ARC EM processor from Synopsys, together with abundant internal memory and numerous digital interfaces, including SPI, I^{2}C, and UART.

The DBMC2-TWS powerful hybrid ANC processing supports advanced audio features in a single, compact, low-power device, while the scalable family of DSP-enhanced ARC EM processors can be tailored for the optimal balance of performance and power consumption, enabling the deployment of advanced sound processing features (e.g., adaptive ANC, user ear placement characterization, and smart ambient awareness control).

The DBMC2-TWS integrates SoundChip's (now part of DSP Group) patented Soundflex technology for 32-bit/192kHz high-resolution audio playback with a suite of advanced features, including Framed Adaptive sound processing and Smartaware augmented audio.

This suite also enables the deployment of advanced sound processing features, including Framed Adaptive ANC, Leakage Compensation, and Wearer Detection. Unlike high-latency digital-only ANC methods, Soundflex technology combines zero-latency, digitally controlled analog feedback ANC processing with low-latency, high-order, digital feed-forward ANC filters to achieve the highest levels of ANC performance.

In June 2020, DSP Group gained access to this valuable suite of technologies with the acquisition of SoundChip, allowing the company to offer a complete set of ANC technologies, engineering services, design tools, and production-line test systems.

By combining SoundChip's proven capabilities in hybrid ANC with DSP Group's SmartVoice advanced low-power voice processing platform, algorithms, and mixed-signal expertise, the company streamlines the delivery of cutting-edge wireless and TWS earbuds, from concept through to manufacturing. Among its clients, DSP Group's technologies have recently been adopted by Google, Panasonic, and Technics in their latest TWS headset models. The SoundChip acquisition also turned DSP Group into a very strong provider of technology and expertise for hearables and smart hearing enhancement solutions, a market that is expected to rapidly grow.

That combination certainly influenced the recent announcement that DSP Group was in turn acquired by Synaptics. Synaptics is the company from San Jose, CA, that in 2017 acquired Conexant—one of the industry's leading audio and voice technology specialists—and more recently acquired DisplayLink and the wireless connectivity business of Broadcom. While originally founded in Herzliya, Israel, DSP Group is also headquartered in San Jose.

The company's leadership positions across multiple markets in the connected audio devices, unified communications and collaboration, and wireless connected devices made DSP Group quite relevant to Synaptics, and promises to create a leading player in the audio processing space.

But in line with the trends already discussed, Synaptics also recently announced its own Low Power Edge AI initiative with the introduction of the Katana Edge AI platform, directly targeting battery powered, connected devices for consumer and industrial markets. The Katana platform combines Synaptics' proven low-power SoC architecture with energy-efficient AI software, enabled by a partnership with Eta Compute. The Katana solution is optimized for a wide range of ultra-low-power use cases and supports object recognition and counting, visual, voice or sound detection, asset or inventory tracking, and environmental sensing—far surpassing the needs of audio processing.

Katana features a multi-core processor architecture optimized for ultra-low-power and low-latency voice and audio with domain specific processing cores, including leveraging neural engines for edge AI combined with wireless connectivity, an area where Synaptics now also leads. Combined with Eta Compute's highly optimized software, this is a key platform for edge intelligence, with audio- or voice-based detection and classification.

"A catalyst for innovation… (enabling) the proliferation of AI edge inferencing in a wide range of existing applications, as well as new applications that were never thought to be possible," as Ted Tewksbury, the CEO of Eta Compute, describes it.

CEVA offers an extensive family of audio SoC and DSP solutions for audio development including the CEVA-BX1 audio/voice DSP for ultra-low-power designs (e.g., TWS earbuds), light AI processing and always-on sensor fusion; and the CEVA-BX2, a high-performance audio/voice DSP, designed for more intensive applications (e.g., smart speakers).

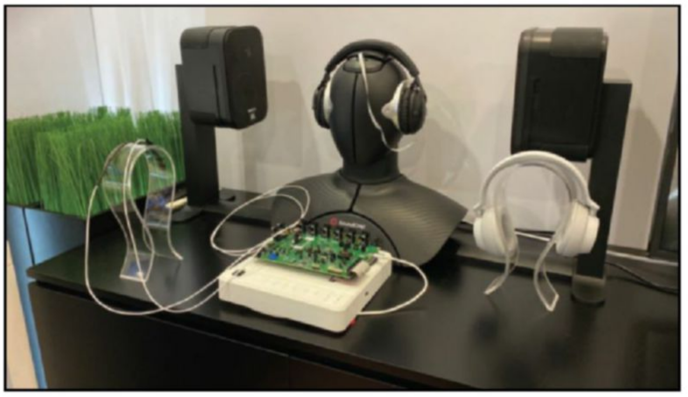

With the acquisition of SoundChip in 2020, DSP Group considerably increased its scope of audio processing technologies and the levels of services offered to the industry. Synaptics recently acquired DSP Group.

Resources

M. I. Jordan, "Artificial Intelligence—The Revolution Hasn't Happened Yet," *Harvard Data Science Review*, July 1, 2019, https://hdsr.mitpress.mit.edu/pub/wot7mkc1/release/9.

K. Pretz, "Stop Calling Everything AI, Machine-Learning Pioneer Says," IEEE Spectrum, March 31, 2021, https://spectrum.ieee.org/amp/stop-calling-everything-ai-machinelearning-pioneer-says-2652904044

Sources

Analog Devices, Inc. | www.analog.com

Cartesiam.AI | www.cartesiam.ai

Ceva, Inc. | www.ceva-dsp.com/app/audio-voice-and-speech

DSP Concepts | www.dspconcepts.com

DSP Group | www.dspg.com

Espressif Systems (Shanghai) Co., Ltd. | www.espressif.com

Fluent.ai | www.fluent.ai

Kardome Technologies, Ltd. | www.kardome.com

Knowles Electronics, LLC | www.knowles.com

STMicroelectronics | www.st.com

Synaptics, Inc. | www.synaptics.com

Synopsys, Inc. | www.synopsys.com

Cadence Design Systems

MPEG-H 3D Audio Runs on Tensilica HiFi DSPs

Tensilica HiFi DSPs are the most widely licensed audio, voice, and AI speech processors with more than 150 partners in the comprehensive hi-fi audio ecosystem. More than 125 top-tier semiconductor companies and system OEMs have selected Tensilica HiFi DSPs for a vast array of products.

Cadence is now offering decoder implementations supporting the new MPEG-H 3D Audio Baseline Profile to designs using its Tensilica HiFi DSP, which opens up the potential to bring interactive and immersive sound to millions of users worldwide.

Tailored to the needs of next-generation broadcast, streaming, and high-quality immersive music delivery, the MPEG-H 3D Audio Baseline Profile was created to meet industry requirements for truly immersive experiences and unmatched advanced Next-Generation Audio features including user interactivity and accessibility. As a subset of the existing MPEG-H 3D Audio Low Complexity Profile, the Baseline Profile enables maximum interoperability with existing devices that have implemented that particular profile, while at the same time significantly reducing the implementation and testing effort.

Cadence was the first company to offer a licensable MPEG-H audio decoder implementation for DSPs in 2016. Adding support for the MPEG-H 3D Audio Baseline Profile with its streamlined set of tools to the Tensilica HiFi DSPs enables manufacturers to meet market demand for Next-Generation Audio features.

Cadence also joined the MPEG-H Audio System Trademark Program administered by Fraunhofer as a licensed decoder provider to offer tested components to end-product manufacturers. The program signals to consumers that MPEG-H products have been verified to work with each other and support all necessary MPEG-H Audio features.

www.cadence.com
www.iis.fraunhofer.de/en/ff/amm.html

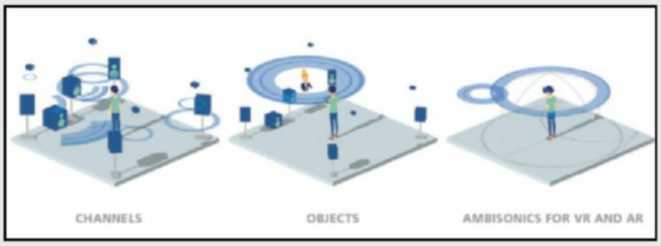

MPEG-H Supported audio formats.

Syntiant

Second-Generation NDP120 Deep-Learning Processor for Audio and Sensor Applications

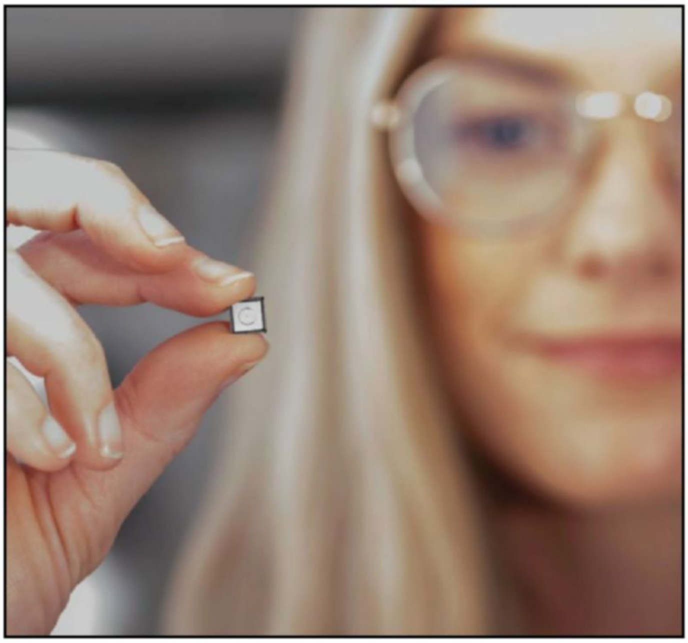

Syntiant Corp., a company from Irvine, CA, specializing in AI for edge devices, announced the availability of its latest Syntiant NDP120 Neural Decision Processor (NDP), a new generation of special-purpose chips for audio and sensor processing for always-on applications in battery-powered devices. The low-power Syntiant Core 2 technology in the NDP120 processor delivers 25× the tensor throughput of the core technology found in the earlier Syntiant NDP100 and Syntiant NDP101 devices that are currently shipping in high volumes.

The Syntiant NDP120 applies neural processing to simultaneously run multiple applications with minimal battery-power consumption, including echo-cancellation, beamforming, noise suppression, speech enhancement, speaker identification, keyword spotting, multiple wake words, event detection, and local commands recognition.

The first of a family of semiconductors using the new Syntiant Core 2 tensor processor platform, the NDP120, brings always-on neural processing to all types of consumer products, including earbuds, smart speakers, and smart home devices. Ideal for use cases where audio filtering and echo cancelation are required for far-field speech processing, the NDP120 packages the Syntiant Core 2 ultra-low-power deep neural network inference engine with a highly configurable audio front-end interface. The NDP120 also supports multi-modal sensor fusion, including sensors for infrared detection, multi-axis acceleration, tilt, magnetic field, and pressure.

For audio developers, the NDP120 includes support for up to seven audio streams, I^2S/TDM output audio interface for streaming audio output, including post-processed audio, embedded user programmable Tensilica HiFi-3 DSP, and a deeply embedded ARM Cortex M0 microcontroller for device management with 48KB SRAM, dual-timers, and UART functionality. This allows the Syntiant Core 2 neural processor to perform in stand-alone operation, with onboard firmware decryption and authentication, quad SPI target and controller interface, I^2C serial interface for sensor applications and system interface, and up to 26 GPIO pins.

www.syntiant.com

Analog Devices

Audinate Dante Embedded Platform for SHARC Audio Digital Signal Processors

Audinate, the developer of the industry-leading Dante Audio-over-IP networking technology, has collaborted with Analog Devices to introduce Dante Embedded Platform for SHARC audio digital signal processors. This means that manufacturers using SHARC DSPs are now able to benefit from Dante audio-over-IP as on-chip software and a streamlined product development cycle.

Dante replaces traditional analog cables by transmitting perfectly synchronized audio (and video) signals across large distances, to multiple locations at once, using nothing more than an Ethernet cable. Audinate has been offering development solutions in cooperation with Analog Devices since the introduction of a software-based Dante reference design kit for its popular ADSP-SC589 DSP + ARM processor, and the Australian company becoming a member of the Analog Devices Alliances program in 2019.

The ADSP-SC589 board, combining SHARC DSPs with an ARM CPU and the Dante Embedded Platform (DEP), provides a cost-effective and flexible means to deliver Dante technology in modern software-driven devices. The software development kit (SDK) for ARM-based audio products, which allows manufacturers to deploy Dante audio networking as on-chip software in their products, has been available since January 2021.

The collaboration with Analog Devices is made possible by Audinate's introduction of DEP, which offers Dante audio-over-IP as on-chip software. By adding DEP to a powerful, compact, audio feature-rich, and multi-core application processor—such as the SHARC ADSP-SC5xx—manufacturers can deliver Dante connectivity in products with high performance, low latency, and deterministic audio processing.

With minimal additional hardware and low marginal costs, Dante Embedded Platform provides access to Dante audio networking and interoperability with the thousands of Dante-enabled products—from hundreds of OEMs—already installed in countless installations.

www.analog.com
www.audinate.com

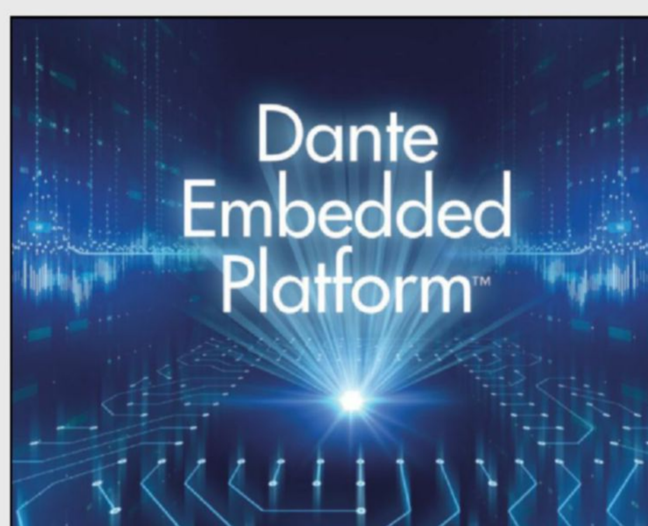

The Dante Embedded Platform for Analog Devices' ADSP-SC5xx DSPs combines an ARM Cortex core, up to two high-performance SHARC audio DSPs with high on-chip RAM, and audio connectivity peripherals. The SHARC DSPs provide a rich ecosystem of advanced audio processing algorithms enabled by Analog Devices' Sigma Studio graphical programming environment that shortens time to market.

Automotive Audio

Interior Car Acoustics and DSP

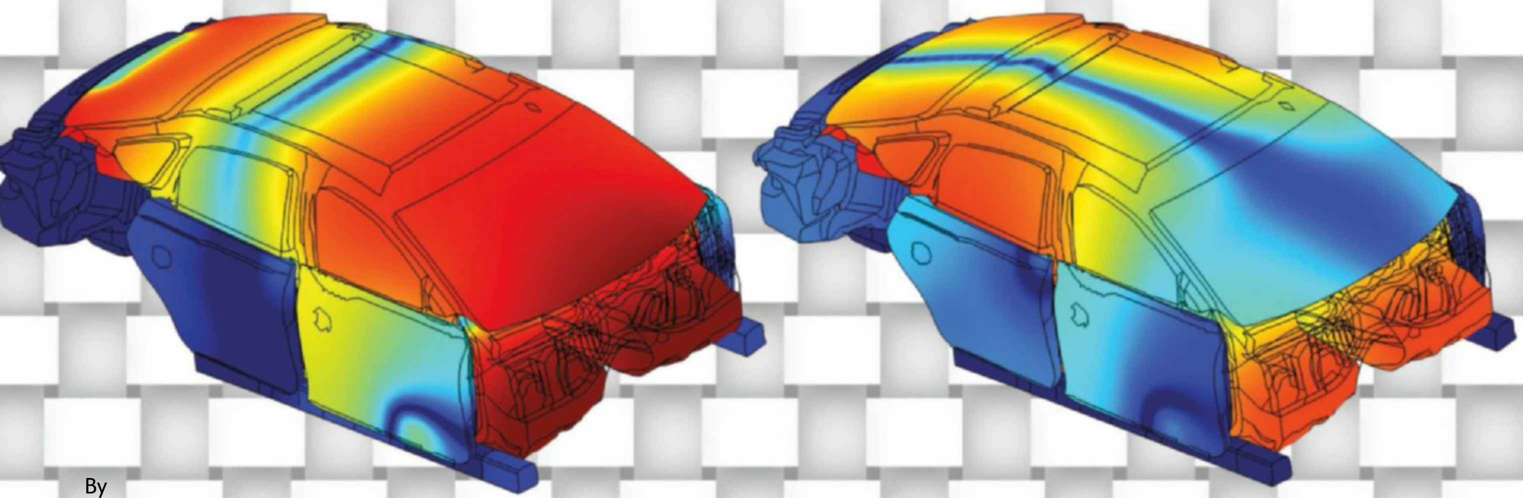

By
Roger Shively

In the unique acoustic space of the car interior, there has always been a need to control the sound field either to simply balance the tonal quality of the audio playback or to reproduce the auditory illusion of a larger event space where the apparent width or placement of a source can be enhanced, and the listener can feel they are naturally enveloped in the soundscape of the original recording or broadcast.

We spend more quality listening time in our homes and cars than we do in an auditorium or concert hall, yet a larger amount of research exists that helps us understand the acoustics of concert halls and factory spaces. A large effort is now being put into the science of sound in homes and cars, because that is where we spend our time and feel, or search for, an emotional attachment to our experiences there. How that is being accomplished is rooted in the challenges that the acoustic space of a car interior creates. It is the environment of the car interior that imposes the boundary conditions on the remaining elements of system design.

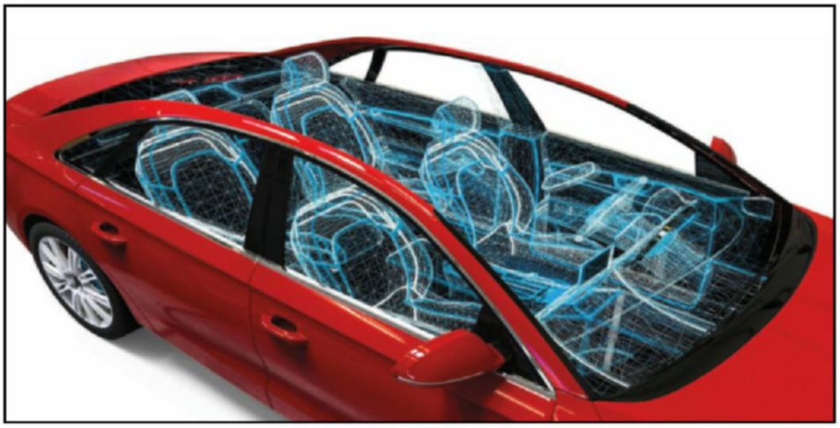

Photo 1: The unique acoustic space of the car interior. In terms of spectral characteristics alone, many factors can contribute to a vehicle's real and perceived frequency response.

The automobile is not an ideal listening environment. It is harsh and challenging. The one redeeming fact is that the seat position for each listener is known. Otherwise, the automobile is a small noisy environment with several negative influences on the spectral, spatial, and temporal attributes of a reproduced sound field. Though the science is still evolving for defining acoustic behavior inside a car, no one could think that the small listening room of a car could be closely related to a larger event space (e.g., a concert hall). Understanding what qualities we appreciate in those larger spaces, can help us optimize our auditory experiences inside of a car. We can look at the interior of a car as if small room acoustics apply (**Photo 1**).

Spectral Characteristics

In terms of spectral characteristics alone, many factors can contribute to a vehicle's real and perceived frequency response: The interior volume of a vehicle, the size and shape of interior boundary surfaces, absorption characteristics of all the boundary surfaces, and the location of speakers in the vehicle relative to the listener and relative to nearby boundaries. We need to solve the same problem we would for any other listening room. We need to understand and predict the room response. In any room, there is a frequency region where the room dominates and one where the loudspeaker dominates, and there is a transition region in between. Below that transition region, in room acoustics terms, is the modal region, understood by wave acoustics—think room modes, eigen values, and so forth. This is the region where the room dominates. And, above that transition region, is the statistical region, which can be understood with geometrical acoustics—think ray acoustics, and other analogies with optics. This is the region where the loudspeaker dominates. As rooms get larger, the statistical region gets larger, and the loudspeaker assumes greater control.

In concert halls, the entire audible frequency range is in the statistical region, and the diffuse-field theory of acoustics reigns almost perfectly. As the room gets to be the small size of a car interior, the room that is the enclosure of the car assumes greater control.

Above 300Hz, and up to approximately 1kHz, the large coupling of modes is reduced and modal resonances with Q values of less than 3 are numerous. In this region there is a great deal of spectral coloration due to the broader spacing of the modes, and variations in amplitude as high as 8dB SPL can be observed. The volumes of the automobile's trunk and doors where speakers are sometimes mounted can contribute resonances as well between 150Hz to 500Hz, with amplitudes anywhere from 4dB to 8dB SPL (**Table 1**).

Below 1kHz, good sound is the result of good transducers, appropriately located to even out room modes of the car and seat to seat variations, with proper equalization based on good measurements and only predictable through simulation for real cars. This is the process of soundfield management to produce a smooth, even frequency response, before equalization and time alignment can be applied to improve the frequency response even more.

Above 1kHz, the loudspeaker can dominate and the mid and high frequencies can be somewhat predictable from anechoic data that characterizes the frequency response and directivity of the loudspeaker. In this region, good sound is the result of good transducers, correctly located, aimed, and mounted, with proper equalization based on good measurements.

Vehicle interiors consist of glass, plastic, carpet, cloth, or leather on surfaces that are at various compound angles with respect to each other, the listener, and the loudspeaker. Reflections off of these many and varied boundaries result in complex interference and diffraction effects, which cause frequency response aberrations and sound coloration at frequencies above 300Hz to 500Hz. Interference occurs between direct and reflected waves. Diffraction will occur through speaker mounting holes and grille coverings, and from mounting locations that are drastically off-axis from the listener. Because of these boundary effects in the car, simulation is required in addition to predictive anechoic data.

Loudspeaker Characteristics

The spatial and temporal characteristics of loudspeakers in a car have an equally important role on reproduction and perception of sound. Because of the small dimensions of the vehicle interior, the reverberation time is virtually zero. In this size environment, a reverberant field cannot form. The sound field in a vehicle consists entirely of direct energy and early reflections, which are quickly absorbed or dissipated. For the typical vehicle, the decay time for a 60dB signal reduction is 30ms to 50ms. This provides a very dead space for sound reproduction. The large number of early reflections produce a type of diffuse sound field where 90% of the energy occurs within the first 10ms of the direct wave arrival. Effectively the listener is experiencing a diffuse and frontally incident sound field, inside of which the first 10ms of sound incidence are somewhat busy (**Figure 1**).

For example, in a left-hand-drive driver's position, a right side speaker can arrive at the left ear 1ms to 2ms later than a left side speaker. The arrival from that right side speaker can have a greater amplitude than the left side speaker. The left side arrival will decay naturally, but the right side arrival will have a reflection from the driver's side window that arrives 4ms to 5ms later, and then it decays after that. The amplitude of that reflection is also nearly as large as the initial right side arrival. The listener integrates this subtle information and creates an impression of the width and breadth of the sound stage. This is not an experience common with a typical listening room. There are no nearby boundaries to smear the time signature of the sound field and confuse the lateral staging or clarity. Quantifying and controlling the lateral information and the direct and reflected information can lead to greater control of the spatial envelopment of the listener in the car.

Digital Signal Processing

Digital signal processing (DSP) is key to controlling this complicated raw sound field and creating not only a pleasant but engaging audio landscape. Companies such as DSP Concepts and Dirac employ data measurement and sound field management tools to do just that to overcome the spectral challenges of a car. The industry leaders of Harman and Bose employ such advanced techniques themselves. And

Average Mid-Size Car	3.5m³
Acoustic Resonances in a Car	
Large Uniform Coupled Modes	80Hz to 300Hz (12dB)
Broadly Spaced Resonances	300kHz to 1kHz (Q<3)
Trunk and Door (Speaker Enclosure) Resonances	150Hz to 500Hz

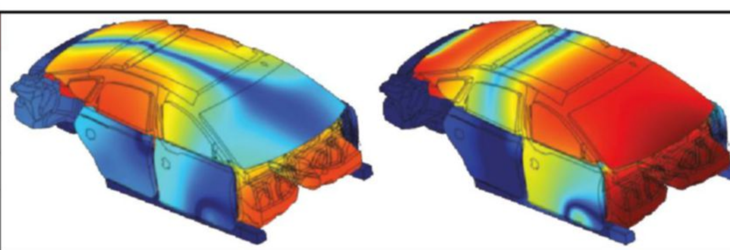

Table 1: Typical transverse and longitudinal modes in automobiles

companies such as Dolby with Dolby Atmos; Xperi with DTS:X.; and Sony, with Sony 360 Reality Audio (which uses an MPEG-H scheme) are promoting content that requires employing techniques to translate spatial envelopment introducing specific challenges for correct translation. The ultimate control of spatial aspects in a car leads us to what we have come to know as "Immersive Audio." All the industry providers are working in this realm now.

Immersive audio is a fairly generic description of audio formats that add height channels to channel-based surround configurations, and evolved to object-oriented solutions, a process which first appeared around 2010, and which was given some criteria standards for content production and distribution by the Society of Motion Picture and TV Engineers (SMPTE) just a few years ago.

Immersive audio is driving new automotive audio speaker architectures—more speakers in overheads and headrests—and finding new ways to use existing architectures. As new content is being created with immersive audio formats, translating sound objects within a certain number of channels along with metadata that describes where the sound should be located in the audio landscape requires careful design in automotive environments.

Automotive audio speaker configurations are well established to deal with existing channel-based audio formats, where the sound sources are recorded and delivered in separate pathways, the most common being stereo (left and right) and the home theater 5.1 layout (left, right, center, surround left, surround right and subwoofer). The object-oriented application of immersive audio opens up greater possibilities not only to translate content delivered in such formats, but also to accommodate new approaches, such as personal sound zones, unconventional speaker layouts (no door speakers), or uses in Advanced Driver-Assistance Systems (ADAS), enabling directional sounds from anywhere.

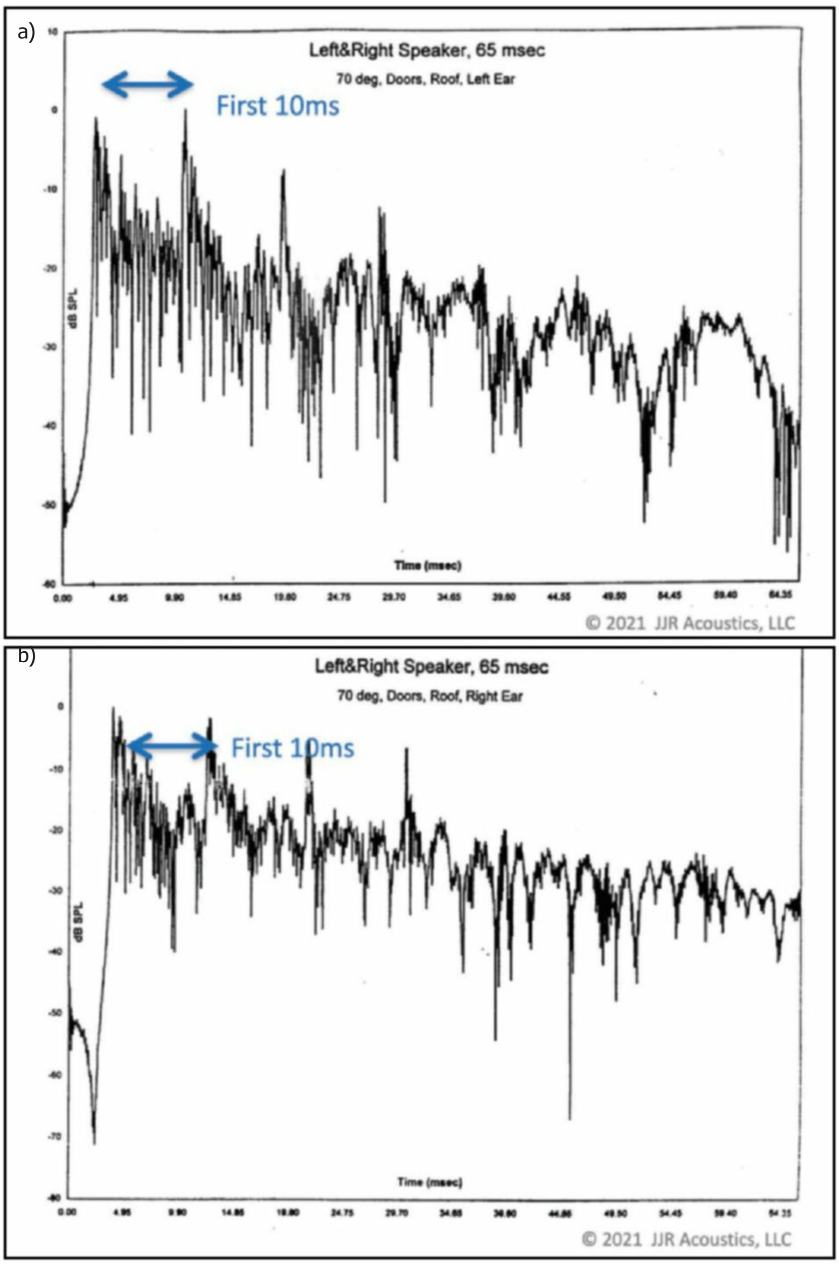

Figure 1: In the time domain, the initial sound arrives (the first peak) followed within 10ms by the first reflection in the car. At the driver's left ear (a) and at the driver's right ear (b). Ninety percent of what we need to know about the spatial qualities of the sound we are hearing happens within those first 10ms.

About the Author

Roger Shively is a co-founder and principal of JJR Acoustics (Seattle, WA). He has more than 34 years of experience in engineering research and development, with significant experience in product realization and in launching new products at OEM manufacturers around the world. Before co-founding JJR Acoustics in 2011, Roger worked as Chief Engineer of Acoustic Systems and as functional manager for North American and Asian engineering product development teams in the Automotive Division of Harman International Industries Inc; a journey that began in 1986.

Roger received his degree in Acoustical Engineering from Purdue University in 1983, and finished post-graduate work in the field of finite element analysis. He is a member of the Audio Engineering Society (AES), Acoustical Society of America (ASA), and Society of Automotive Engineering. He has published numerous research papers and articles in the areas of transducers, automotive audio, psychoacoustics, and computer modeling. Roger also holds US and International patents related to the design of advanced acoustic systems and applications particularly in the field of automotive audio. Roger is Co-Chair of the AES Automotive Audio Technical Committee.

In addition to sound field management, sound field synthesis, and adding voice control to the car audio experience, digital signal processing also is being used to help dealing with the challenges of noise within the car cabin.

Even quieter cars still have the noise of the road to contend with. Rough driving surfaces, tire treads, and chassis vibrations. These are factors that produce noise in any type of vehicle, but they are especially noticeable in quiet EVs that do not have the conventional rumble of an engine to mask those annoyances. Road, wind, and engine noise can reduce an audio system's dynamic range and mask low frequencies. For the typical mid-sized automobile, the noise is most dominant below 500Hz. The average noise SPL is greater than 80dB for frequencies below 100Hz and velocities above 35mph (**Figures 2–4**).

Noise Control

Bose recently introduced QuietComfort Road Noise Control as an active and adaptive technology, which is also part of its Active Sound Management technologies. Bose is widely recognized as the leader in noise control, so bringing its automotive-specific solution to Audio Weaver from DSP Concepts is a huge value add for the carmakers building on that platform. Harman is expanding its ANC technology, allowing each person in the car to listen to their own audio

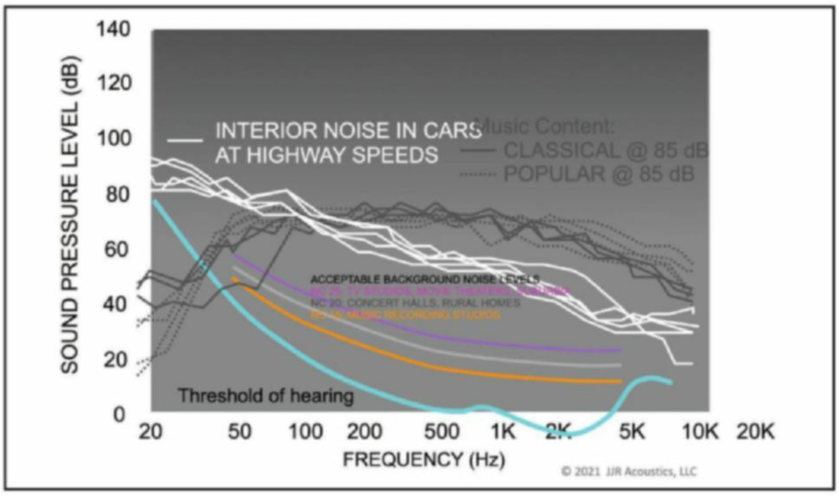

Figure 2: Even under ideal conditions, low bass is hard to hear in the car. The ear is less sensitive to low frequency even without a noisy background (threshold of hearing). Acceptable background levels are shown with the Noise Criteria curves for average overall levels of 16dB, 20dB, or 25dB. The levels of classical or pop music content are to be quite high at ~85dB. But with the typical interior noise in car at highways speeds, the noise inside a car masks most of the low bass.

(**Photo 2**), while minimizing the impact on other passengers. And, other players in Automotive ANC (e.g., Silentium) are also meeting the needs of our quieter cars.

The other perspective is the noise identity of the car—and the brand—normally associated with the engine, With an internal combustion engine, the sound of the car itself as it idles and as it is driven creates an emotional attachment to the driver. With electric vehicles, the goal of automotive audio engineers is also to help create that comfort of emotional attachment through internal sound synthesis of engine and road sounds played through the audio system. This can help car owners hear when to shift, but it can also allow car makers to create unique, unconventional powertrain

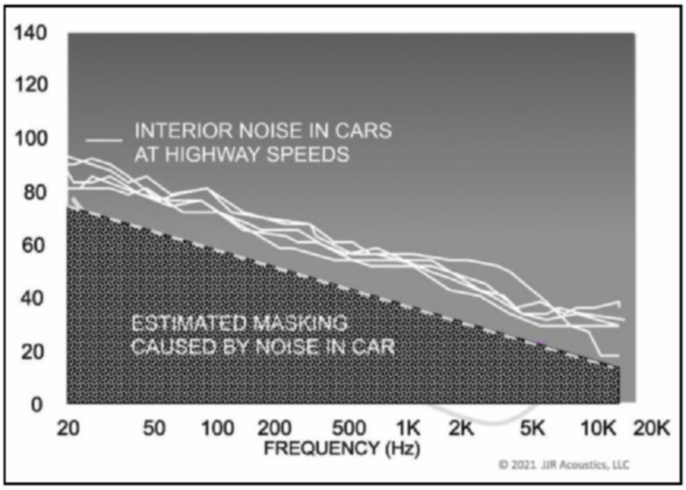

Figure 3: In cars, everything suffers when in motion. The estimated masking of the content from 20kHz to 10kHz, not just the low frequencies, is shown.

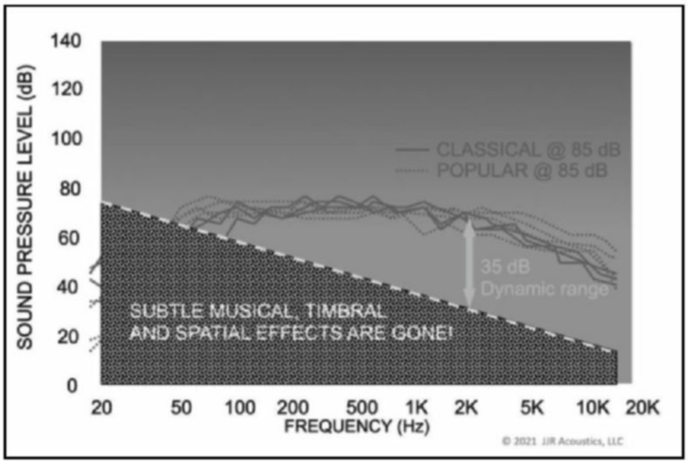

Figure 4: And with that level of masking reduces the estimated dynamic range in a car 35dB. A typical 16-bit, 44.1kHz sampled recording would normally have 100dB of dynamic range available. The car reduces that to almost one third.

sounds inside the cabin—delivered through the sound system's speakers. This is part of the Electric Vehicle Sound Enhancement Bose provides, which was developed with the help of DSP Concepts (**Photo 3**).

Granular synthesis is becoming a mainstay in research and development labs. Sound samples are split into very small pieces (1ms to 50ms) called grains. In playback, multiple grains may be layered on top of each other, and may play at different speeds, phases, volume, and frequency, among other parameters. At low speeds of playback, the result is a soundscape; at high speeds, the result is heard as a note or notes of a unique timbre.

Conclusions

The power of digital signal processing has given us control of the raw world of the car, and provided the tools to create a world where the new opportunities that lie on the horizon are seemingly endless.

The existing platforms for DSP are evolving in automotive as in other industries and the features we have come to expect, with varying degrees of comfort, are now a part of how we interact with the car and the audio experience it provides.

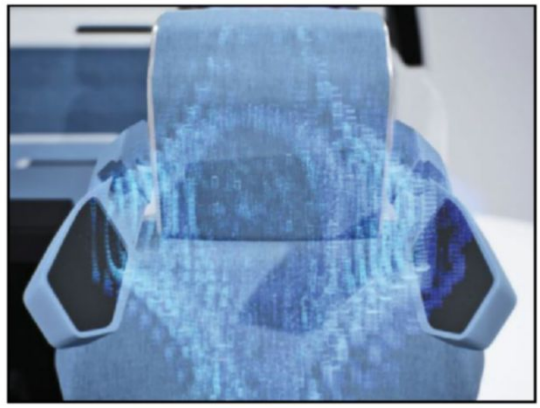

Photo 2: Harman is able to create Individual Sound Zones (ISZ) as an integrated audio entertainment solution that gives every passenger the freedom to choose his or her own audio entertainment while maintaining a harmonious in-cabin experience for everybody.

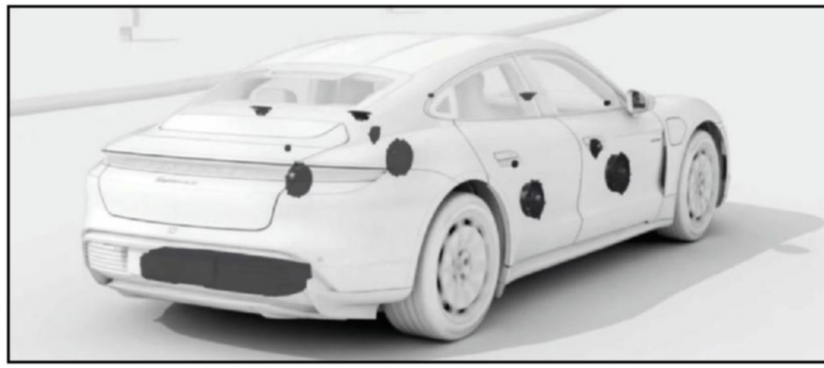

Photo 3: Bose supplied a full suite of Active Sound Management (ASM) technologies to design and engineer the sound system for the unique cabin acoustics of the Porsche Taycan. For this unique project, Bose implemented his Bose QuietComfort Road Noise Control (RNC) processing solution into Audio Weaver with the support of DSP Concepts.

A DSP Audio Processor at the Heart

By
Al Clark
(CEO, Danville Signal Processing)

In this article, Al Clark, the co-founder and CEO of Danville Signal Processing, shares the journey that led his company to develop and successfully bring to market a dedicated DSP audio processor. An interesting inside story about how the author's experience led him to maximize the possibilities of digital audio.

Have you ever had the opportunity to create something that is truly new? As a designer of hundreds of different products over many years, I can honestly say not very often. Most of our products are tweaks to existing ideas that are built on the knowledge that we gained from our own earlier designs and more importantly, from others who perhaps share our passion.

This is the story of Danville Signal's dspNexus DSP Audio Processor. Simply put, the dspNexus combines the traditional role of a modern preamplifier (control center) with a powerful DSP crossover. As much as many of you may want to delve into the engineering details, this article will focus on the journey and the genesis of the dspNexus. I hope that many of you will have a chance to further investigate the dspNexus and draw your own conclusions of whether we accomplished our goals.

The dspNexus DSP Audio Processor serves as the control center, the preamplifier, and the DSP engine for powered loudspeaker applications.

The Journey

I don't think I can tell the dspNexus story without sharing a little of my personal background and later the founding of Danville Signal.

In my youth in the mid-1970s, I joined a small team of gifted audio enthusiasts not so different than the audience and contributors of *audioXpress*. I was a sophomore with an electrical engineering major. We started by modifying Dynaco preamplifiers and power amplifiers, and we had some commercial success. By the end of my junior year, we started designing more original circuits. At this time, I invented the split passive phono preamp, which was later patented. Modifying existing circuits is like covering songs, creating new circuits is more like writing your own material. Besides learning a bit of circuit theory, I think the other skill I learned from this collaboration was the ability to listen critically, which is essential for any good audio designer.

In the early 1990s, I designed noise reduction products using Analog Devices DSPs. I think this was a fairly natural transition for an analog engineer since the math is not so different and the applications often overlapped. In 1998, I founded Danville Signal Processing with the intent of building audio signal processing products using Analog Devices DSPs. I designed my first SHARC-based DSP design in 2000 with a first-generation Analog Devices' floating point DSP. In my view, the SHARC is the gold standard of DSPs for audio processing.

Of course, a village with one hut is not much of a village, so now the story moves to our team, since

Danville Signal is no longer a solo engineer and a few subcontractors. Danville now has an engineering department so our products evolve from a collective effort. We also manufacture all of our products on our own internal automated line in Cannon Falls, MN. This flexibility really helps when it comes to creating new products.

Danville Signal's customer base has traditionally been aimed at other manufacturers, so you may not be well acquainted with our products. Since our customers are the brand, we are more like "Intel inside." Nevertheless, we have likely infiltrated your audio world. Our DSP crossovers have been inside popular studio monitors for more than 10 years, which means some of your music content was mastered with our products. We have also provided special mastering electronics for vinyl pressing, DSP modules for beam forming in large acoustic spaces, DSP modules for a SOA CD player, and DSP crossovers for high-performance consumer loudspeaker systems.

An early iteration of the dspNexus shown at AXPONA 2019. (Photo courtesy of Marc Phillips, *Part-Time Audiophile*)

DSP and the Digital Revolution

It should be clear to just about everyone that digital audio is the new normal. Most of the content that we routinely playback is either stored on a local music server or streamed from an Internet source. We might also grab content from a TV or phone, which usually is available in a digital format. This means that your playback system will need at least one DAC.

Now I will make another assertion: With the possible exception to cost, there is no advantage to a passive loudspeaker crossover. This has been known for a very long time. I have a video on the Danville website that discusses this issue and you can find many other presentations that present similar arguments.

So if you accept the previous two assertions, I will claim that a good DSP-based active crossover is superior to an analog active crossover for the following reasons:

- Content is probably already in digital format
- Filter complexity is rarely limited given powerful DSPs
- Optional FIR filtering
- Time alignment is easy and usually desirable
- Parameters are software driven and easily adjustable
- Stability of performance
- Look ahead limiters are possible

Danville Signal and the large-space professional acoustics market has embraced this idea for quite some time. From their perspective, the ability to steer arrays, align systems in time, and so forth are compelling arguments for DSP.

The small-space professional market (recording studio monitors) has now also embraced DSP crossovers as a desirable feature in most high-quality studio monitors. In this market, most speakers are already powered, so adding a DSP internally does not change the form factor.

The high-end consumer market has been the slowest to adopt, but this is changing. I remember inadvertently meeting a reviewer from *Stereophile* in an elevator at CES around 2014 or 2015. He asked me what we did and I said: "DSP Audio." His response was, "You must be very busy." This was clearly a change from early years where DSP was more suspect. I think the tipping point started around 2016 when consumers started asking speaker manufacturers if their products use DSP crossovers.

The dspNexus Idea Takes Form

Danville had been manufacturing a basic DSP box for several years called the dspMusik 2/8. It started out as very simple box that had selectable stereo inputs: balanced analog, USB Audio Class 2, and S/PDIF. It had eight balanced audio outputs, making it well suitable for DSP crossover applications.

audioXpress highlighted it in an issue a year ago. Internally, it had a SHARC DSP that was running DSP Concepts' Audio Weaver, a powerful GUI-based design tool that lets you drag and drop optimized

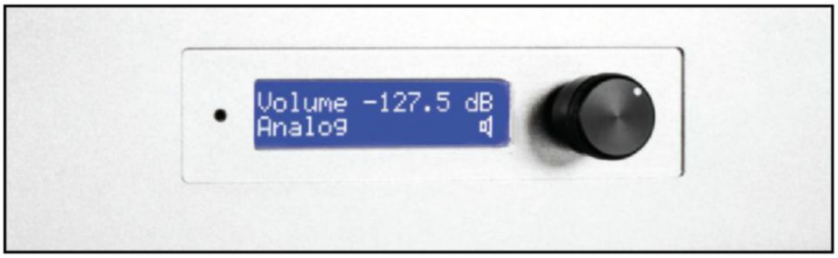

With the front display, users can select input sources, including Balanced Analog Stereo, USB Audio Class 2, S/PDIF, Ethernet, and Wi-Fi. Don't want to wake up the neighbors? A headphone amplifier is included for an immersive personal listening experience.

DSP modules to create the signal processing needed for the DSP crossover and other signal processing functions. The strength of the dspMusik was that it was probably the most powerful and best sounding DSP crossover available at the time. It used top performing AKM ADC and DACs.

One day I was visiting Rich Hollis of Hollis Audio Labs (HAL). Rich makes some very good planar magnetic line arrays that used a dspMusik as the crossover. We were invited to a demo at one of Rich's local customers. This system sounded great! Not surprisingly, it had a stack of excellent power amplifiers, fat expensive cables, and a music server all residing in a very nice room. In the middle of everything was this humble dspMusik 2/8, running the whole system. It was clearly time to rethink the dspMusik with a fresher approach.

Rich and I started defining what would ultimately become the dspNexus platform. The dspNexus target market is primarily speaker manufacturers and serious loudspeaker DIYers who want to create the highest performing loudspeaker systems. You really can't do this without a modern active DSP crossover. This means that a loudspeaker manufacturer using this approach needs a solution that is easily tailored for the specific requirements of the target loudspeaker. Most dspNexus systems sold to end users will be preconfigured for a specific brand and model. We would supply DSP Concepts' Audio Weaver.

We envisioned a product that would never apologize for its audio performance. This suggested that most key components would be constructed as replaceable modules. Today's best DSPs, ADCs, and DACs are tomorrow's second best. To this end, it also meant that each module would have read-back capability so that as new modules are created, we could automatically determine what was plugged into the motherboard and adapt accordingly. This also required the product to be software upgradeable without special programming tools.

It was clear that we needed to upgrade the enclosure to be much larger and more attractive. In the high-end audio market, there has been a bit of an arms race to make enclosures with incredibly thick front panels and shiny mirrored surfaces. This greatly increases manufacturing cost, without any benefit to performance. We thought the dspNexus needed to look good and professional, but we weren't aiming for bling.

In many ways, this part of the conceptualization process and later the actual implementation was the easy part. We already knew how to make the DSPs, ADCs, DACs, and supporting circuits work well. This area is our company's core competency after all. If this was our entire goal, we would have just repackaged a dspMusik, but we really wanted to rethink how to make a preamplifier/control center that would be relevant to modern stereo systems.

In today's world, audio sources come from many places. We settled on the following input sources as a reasonable set to accommodate most needs:

- Balanced Audio Inputs
- USB Audio Class 2
- S/PDIF
- Bluetooth with aptX HD support
- Wi-Fi
- Ethernet

We decided to add a headphone amplifier because it was clear that sometimes people want to listen to music when others may not want to share their experience (perhaps they're sleeping?). The headphone amplifier needed to compare well with other high-quality headphone DACs. Since most of the supporting circuits would already exist in the dspNexus, this was really just a matter of creating a small power amplifier with appropriate signal routing.

Perhaps the biggest departure from a traditional preamplifier is that we decided to integrate a Raspberry Pi 4 (RPi4) small board computer (SBC). The RPi4 has very broad support and is likely to be available for many years. This made it the SBC of choice over similar devices.

By adding the RPi4, we opened the door to running applications in the dspNexus. This could be for speaker measurement programs, music server, or room correction/bass management. It also could function as a bridge for external Wi-Fi or

About the Author

Al Clark is the co-founder and CEO of Danville Signal Processing. As one of the four founders of Van Alstine Audio Systems, Al launched his career 45 years ago as an analog engineer. This company made high-end preamplifiers and power amplifiers in the mid-1970s. He started designing DSP-based products with first-generation DSPs in the late 1980s and has since created more Analog Devices' SHARC-based DSP boards than anyone else in the world. He has a Bachelor of Electrical Engineering degree from the University of Minnesota and is a life member of the Audio Engineering Society.

Ethernet sources. In the dspNexus, a local SHARC DSP does the main signal processing, so the RPi4 is just another peripheral.

Many customers are going to make speaker measurements, so we added a phantom-powered microphone input. It is also possible to measure delay from each DAC to the microphone. This allows relative time delay to be measured in a customer environment for aligning subwoofers, since the placement of a subwoofer would be unknown until installation.

Frankly, the only feature we added that I think might be heresy for some users, is that the ADC module can be configured as a phono preamplifier front end. The differential input impedance is 47kΩ. It also has load capacitors that can be switched for cartridge loading. This sets the stage for a DSP phono preamplifier for moving magnet cartridges or moving coil cartridges with an external transformer or pre-preamplifier. In principle, this opens the door to pop removal since we can look at the output of the phono cartridge before equalization and delay it a small amount. This would allow the pop to be removed before it was too late to snip out. Regardless, it will be fun to play with a DSP phono preamplifier implementation.

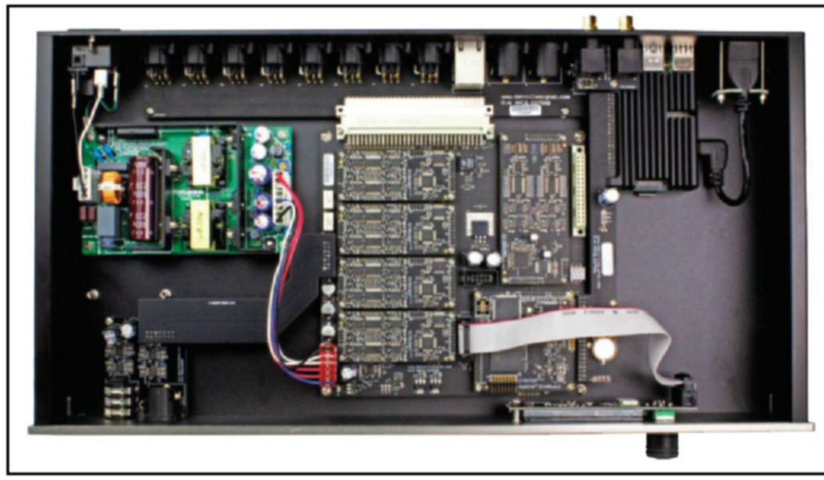

ADI SHARC DSPs and Velvet Sound ADCs and DACs provide cutting-edge performance, while an integrated Raspberry Pi SBC allows for local network streaming and an interface for various applications. The dspNexus even comes equipped with a phantom-powered microphone input and measurement ADC.

The first dspNexus units are shipping to earlier adopters. We are excited about the dspNexus and its unique set of capabilities. Since it uses modular construction and is software upgradeable, the dspNexus will continue to improve as we move forward. Get your hands on one and let us know what you think as well as what you'd like to see in future upgrades.

From Open Office to A/V Recording Complex

By
Richard Honeycutt

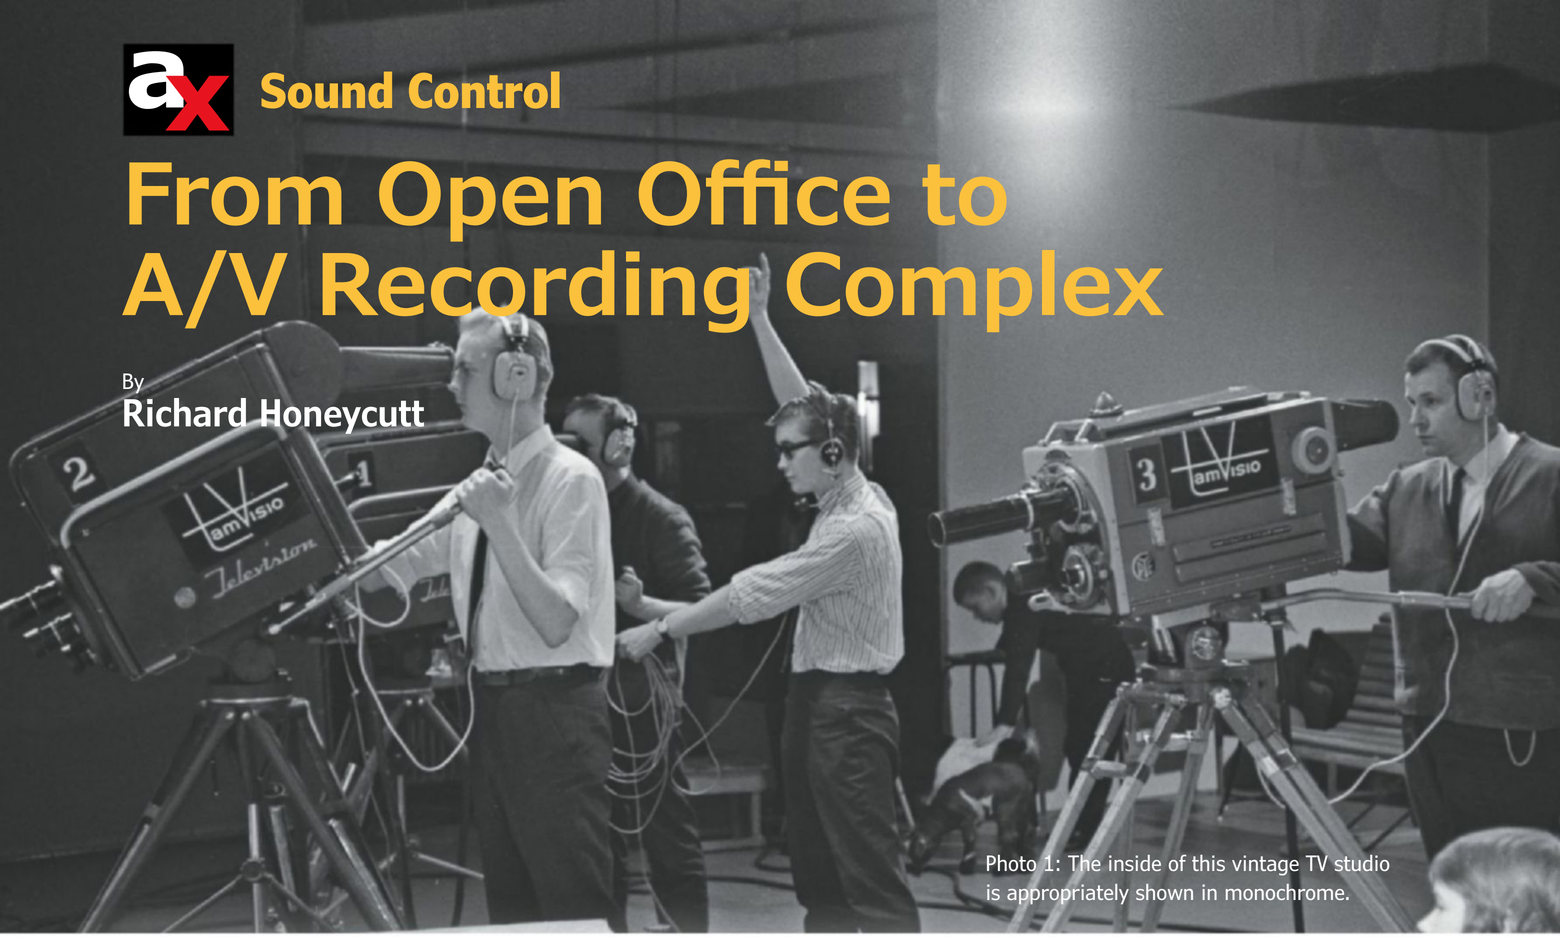

Photo 1: The inside of this vintage TV studio is appropriately shown in monochrome.

Although classic TV studios have changed dramatically as they morphed into modern AV studios, the acoustical design challenges remain largely the same.

When we think of a TV studio, we old-timers probably conjure up an image of a large brick or concrete building housing offices and one or more sizeable rooms having trolley-mounted vidicon or image orthicon cameras (**Photo 1**), and boom-mounted ribbon or condenser mics similar to the one shown in **Photo 2**. Earlier television studios were mostly improvised structures designed with more concern for avoiding outside noise being captured. Acoustics were far from ideal.

Photo 2: Outside of the studio, the major concern is ambient noise. This outdoor production scene shows the use of a microphone boom for proximity to the actor, with the microphone protected with a windscreen to reduce wind noise.

Advances in technology have vastly improved the technical quality of TV and other forms of AV media while reducing the cost, so that today it is not uncommon to find excellent TV studios in schools, commercial office buildings, and even churches. In fact, the capabilities of a modern video teleconferencing (VTC) room, home, church, or commercial streaming room far outstrip the wildest dreams of the TV engineer of the late-20th century. At the same time, the outer appearance of the building housing the studio may have changed dramatically.

An Acoustical Challenge

A commercial client recently engaged me to assist in the design of a new A/V recording complex, built into the renovated shell of a former open office building. Although the steel building (**Photo 3**) is very different from most people's concept of a TV studio, the acoustical requirements remain the same:

- Noise isolation—This includes isolation from interior noise sources (e.g., conversation or music being played on cell phones) and from external sources (e.g., nearby highways, big trucks, and heavy equipment).

- HVAC noise control—In open offices, a somewhat high level of HVAC noise is sometimes intentionally allowed in order to mask nearby conversations so as to provide speech privacy. Typical conversation involves sound pressure levels (SPLs) of about 65dB at 1 meter from a talker. The SPL in an adjacent cubicle would be reduced by distance and the cubicle walls, albeit they will likely be partial-height walls. So it may be about 15dB lower, or 45dB. Humans can understand speech when the background noise level is up to 10dB higher than the speech level. So to mask speech, we might need a background noise level of at least 55dB. This can be provided by the HVAC system or by a purpose-built noise masking system.
- Reverberation and echoes—Unlike some sound-recording studios, TV studios are virtually never designed to contribute any acoustic signature: by achieving a very low reverberation time (RT) and freedom from slap and flutter echoes, we allow the acoustical environment of the videos to be manipulated by Foley artists if desired, or alternatively to be as neutral as possible.

Internal Noise Isolation

Let's examine the noise isolation design process for this particular studio. **Figure 1** shows a partial layout of the building. We'll refer to this as we discuss the noise isolation analysis.

Table 1 lists the Noise Criterion (NC) targets for the studio and related rooms. The effects of internal, external, and HVAC noise are included in these targets. Information on NC ratings can be found at NC—Noise Criterion (see Resources).

As you can see from Figure 1, the corridors provide some isolation from adjacent-room noise, so avoiding noise intrusion from conversations in the corridors becomes the main internal noise issue. Toilets can be very noisy both during flushing and when high-velocity hand dryers are used, but the distance (22′) between the closest points on the studio and the toilet help significantly with isolating toilet noise.

Besides noise intrusion through walls, noise can enter a space by passing up through the ceiling of the noisy room, then across and down through the ceiling of the sensitive room. Fortunately, the architects designed this building so that all partition walls extend to the roof deck, minimizing noise intrusion from this path. I recommended sealing the walls to the roof deck using acoustical caulk to prevent sound leaks.

Photo 3: Using a commercial steel building can introduce acoustical design challenges.

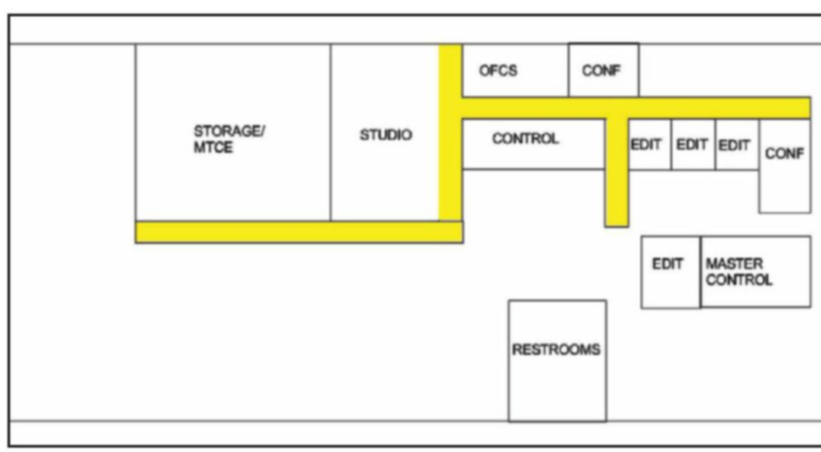

Figure 1: This partial plan of the building shows the studio and related rooms. Yellow areas are corridors, and blank areas are offices, storage, and so forth.

Room Types	NC	Aproximate Overall Sound Pressure Level	
A/V Recording Studio	15 to 20	30dBA	60dBC*
Control Room and Master Control Room	20 to 25	35dBA	60dBC*
Editing Rooms	20 to 25	35dBA	60dBC*
Conference Rooms	30	40dBA	60dBC*

Table 1: Noise Criterion Targets (*Corresponding dBC levels vary widely due to frequency content of the noise, so these should be taken with many grains of salt!)

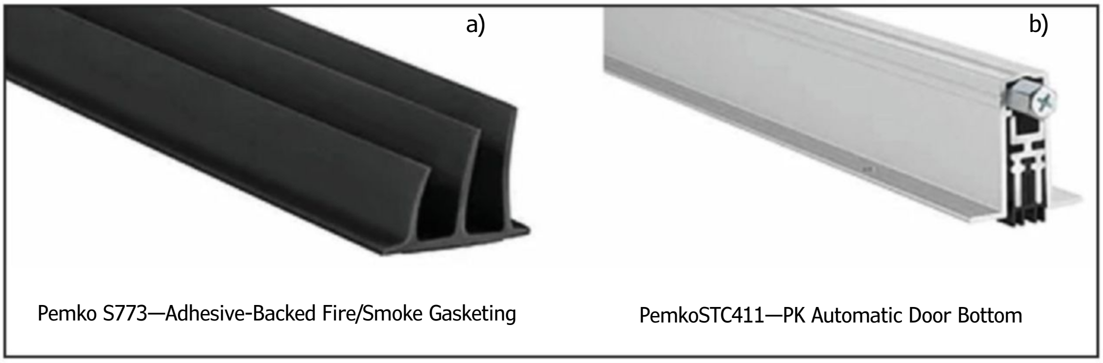

Figure 2: The side/top gasketing (a) and the door bottom seal (b) improve the STC by about 18. (Image courtesy of ASSA ABLOY Accessories)

Conversational noise averages about 65dBA, and since talkers in a corridor could be very close to a wall, door, or other barrier, a level reduction of 35dB would be required for the studio, if the conversational noise were the only concern. If external noise and HVAC contribute equally with conversational noise, an additional 5dB of attenuation in the studio boundaries will be needed. This leads to a required Sound Transmission Coefficient (STC) of about 40.

The architect had designed the partition walls to be framed with 25-gauge steel studs and one layer of 5/8″gypsum wallboard on each side. The cavity will contain R-13 fiberglass batts. This wall construction yields an STC of 47, which is quite satisfactory.

HVAC ducting from a common plenum can provide an additional noise path. Fortunately, the owner had already planned to install a separate HVAC system for the studio, eliminating this concern for the most sensitive space.

Noise intrusion can occur through any wall, window, door, or ceiling. Often, windows and doors only have an STC of 17 (single-pane window) to 27 (hollow-core interior door). Part of the reason for the low STC is the low mass of the barrier: STC is directly related to the mass of the barrier. Thus, a concrete wall will have a higher STC than a wood-paneled wall, and a double-pane window will have a higher STC than a single-pane one.

Another issue with doors is leakage around the door. A 10% leakage area around a door can decrease the STC significantly. I recommended solid wood doors with acoustical seals (**Figure 2**) for the studio, edit rooms, and control rooms in this project. These door assemblies achieved an STC of 38. (I previously consulted on a project in which standard glass-and-steel doors had been used for edit rooms in a university Communication program building, and the sound leakage had made the rooms virtually unusable.)

Isolation from External Noise

Investigation revealed that airplane overflights would not be a problem for the studio under consideration, and although there are occasional deliveries by eighteen-wheelers, the main external noise source is the diesel backup generator located just north of the building, on the other side of a wall from the studio, but offset to the west by about 20′. The generator housing is only about 6′ from the outer wall of the building. Isolation from this serious noise source required careful planning. Site measurements indicated that when the generator is running, it produces a C-weighted noise level of about 100dB at the door of the generator (**Photo 4**).

Complicating the analysis are the facts that the building wall is located in the near field of the generator, creating a complex coupled-wall effect, and that the loudest noise from diesel engines occurs at the exhaust stack. However, the top of the exhaust stack is about 10′ from the roof of the building housing the studio, and it acts as a point source, so the noise level will decrease at a rate of 6dB per doubling of distance.

Photo 4: This diesel generator is the primary exterior noise source.

Sound Level Differences

The sound level difference between a location outside a room and inside the room is never quite as low as you might suspect by subtracting the STC of the intervening partition from the outside level, but for rooms with high sound absorption (low RT), and/or small partition area between the source and the receiving room (compared to total wall area) that approach gives a pretty good approximation.

In spite of my efforts to get octave-band noise readings for the generator, the owner was not willing to pay for them, so I had to infer from the C-weighted levels they provided. Since diesel engine noise is highest at low frequencies, C-weighted measurements are probably preferable to A-weighted measurements, given the unavailability of octave-band measurements. With the generator running, the noise level outside the building 6′ from the generator was measured as 80dBC. The noise level in the sensitive rooms was measured with the HVAC on, and both with and without the generator running. HVAC noise was highest in the studio, and a few decibels lower in the control and edit rooms. The present article does not specifically discuss HVAC noise control.

The noise levels with the generator running were 1dB to 4dB higher than with the generator off. Average HVAC level was 62.4dB. Since the decibels correspond to the log of a ratio, the noise contribution from the generator could be calculated at 56.4dBC to 64.2dBC, requiring an average area-weighted STC of the walls and ceiling to be higher than the present value. Without octave-band noise levels, I could not calculate the exact STC increase required. However, based upon typical diesel engine noise profiles, an engine producing a noise level of 80dBC at 6′ would produce about 77dBA at that distance. To achieve an indoor noise level of 30dBA (approximately corresponding to RC 15-20), an average area-weighted wall/ceiling STC of 47 is required.

As a historical note, a public school located near a naval base had a problem with noise from training jet overflights, and asked a colleague of mine to recommend wall improvements to solve the problem. However, investigation revealed that the primary culprits were poorly sealed single-pane windows, direct fresh-air venting through operable steel louvers, and a thin mobile-home-grade roof. Thus, wall treatment would not have solved the problem.

Outer Walls

In the present project, there are no windows or doors in the outside walls located near the noise-sensitive areas. The outer walls themselves are similar in construction to the partition walls described earlier, but the outer leaf is 25-gauge steel over 3/4″ plywood rather than 5/8″ gypsum board, and the 5 1/2″ steel studs are used instead of 3 1/2″ones. The result is an STC of 47, same as for the interior walls. I recommended adding a second interior wall spaced 1″ from the existing one, and consisting of a layer of 5/8″ QuietRock on each side of 3 1/2″ 25-gauge steel studs with fiberglass batts in the cavity. No fiberglass will be used in the 1″ space between the walls. The builder objected to this plan, so I compromised by accepting a double layer of 5/8″ QuietRock on the inside of the second wall only. This addition raises the STC above 60.

The bigger concern is the roof, which consists of a rubber membrane over a steel deck backed by a rigid layer of polyisocyanurate foam. This construction only provides an STC of about 22. I recommended adding a layer of 5/8″ Quiet Rock to the underside of the roof joists and using a double-layer, high-CAC suspended ceiling in the studio.

A ceiling tile having a CAC of 35 will reduce the sound level of noise from one room entering another room by 35dB. However, this is based upon the sound being attenuated by the ceiling in the room in which the sound source is located, and again by the ceiling of the "receiving" room. Thus, the sound attenuation from the "source" room to the air space above the ceiling is less than the CAC value: 25dB for the two layers of CAC 35 tiles in this case. If we add the attenuation of the added Quiet Rock and the ceiling to the STC of the roof structure, we obtain a composite STC of 72.

About the Author

Dr. Richard Honeycutt fell in love with acoustics when his father brought home a copy of Leo Beranek's landmark text on the subject when Richard was in the ninth grade. Richard is a member of the North Carolina chapter of the Acoustical Society of America. Richard has his own business involving musical instruments and sound systems. He has been an active acoustics consultant since he received his PhD in electroacoustics from the Union Institute in 2004. Richard's work includes architectural acoustics, sound system design, and community noise analysis.

Resources

ASSA Abloy Accessories,
www.assaabloydooraccessories.us

"NC - Noise Criterion," The Engineering Toolbox,
www.engineeringtoolbox.com/nc-noise-criterion-d_725.html

Calculating the area-weighted STC of the entire room consists of adding the STC contributions of all the different surfaces, taking into account the areas. The following equation is used:

$\tau C = \tau_W + \tau_{clg}$,

where τ is the transmission coefficient of the outside wall (w) and ceiling/roof (clg), given by: $\tau \approx$ antilog10(-STC/10).

For the studio, (the most critical space for low noise), this comes out to 69.

A similar calculation could be done to assess the STC applying to interior noise, using the STC of the interior walls and door.

Although the analysis and design process discussed in this article cover the needs of most studios, there is one other possible concern that applies to some projects: noise and vibration conducted through the ground and re-radiated into the studio by the floor.

During the 1980s, I made some measurements in a certain university's anechoic chamber, which had been constructed about four blocks from a railroad. During sensitive vibration measurements, significant noise components in the 0.4Hz to 4Hz range could interfere with measurement accuracy when a train passed.

If the railroad had been closer, the infrasonic noise might have been louder, and noise could have extended up into the audible frequency range. Had this problem been noticed before construction, the chamber could have been built with a floating floor to prevent conduction of vibrations into the floor and subsequent re-radiation into the chamber.

Ceiling Attentuation Class

The Ceiling Attenuation Class (CAC) quantifies the sound attenuation between two rooms whose ceilings abut a common air space. Sound from one room must pass through that room's ceiling to get into the common air space, then down through the other room's ceiling. The CAC primarily depends upon the distributed weight of the ceiling material (pounds/ square foot). Fiberglass ceiling tiles have a low CAC; mineral fiber tiles have a higher CAC. A value of 35 or above qualifies a ceiling as "high-CAC."

Ultra-Linear: Why It Works And When It Will Not (Part 2)

Some Interesting Conclusions

In Part 1, Morgan Jones reviewed the basic tenets of ultra-linear operation and looked into the detailed process with several tube types. In Part 2, he addresses several questions and comes to some interesting conclusions.

By
Morgan Jones

To begin this article, we formulate the following questions:

- Does g_2 contribute power to the load, or is it purely negative feedback?
- What does the g_2 current waveform look like?
- How do we ensure that ultra-linear works well?

A standard 300V mostly Class A EL84 ultra-linear push-pull output stage was driven directly from the balanced output of a Prism Sound dScopeIII audio test set. The output transformer was $8k_{a\text{-}a}$ with g_2 taps at 20% of turns, and the 16Ω secondary was loaded with a 16Ω non-inductive dummy load (**Figure 9**).

Unusually, the two half-primaries were independently available before the high tension (HT) connection (this fact determined transformer choice). One channel of the audio analyzer monitored the voltage across the 16Ω dummy load, while the other monitored the output of a Tektronix current probe/amplifier system on one valve's g_2 lead.

Given that current probes are not especially common, a few words of explanation are appropriate. A current probe is a no-contact device that senses the magnetic field associated with a current in a wire. Typically, two halves of a magnetic core open to allow the test wire to pass through, and are subsequently closed. The enclosed wire becomes a half-turn primary, and windings on the core form the secondary.

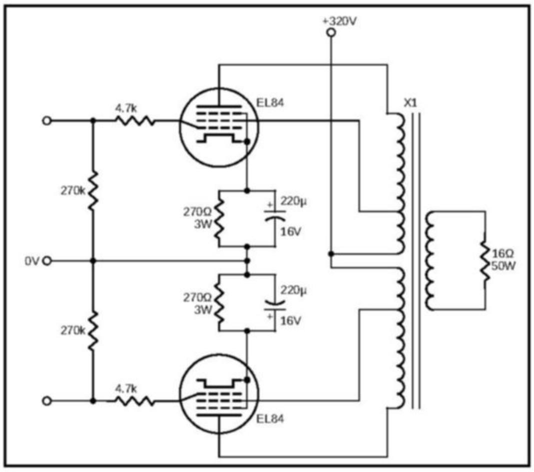

Figure 9: Standard EL84 ultra-linear output stage

The really clever bit is when the core also includes a Hall effect device. The Hall effect device measures magnetic flux down to DC and its current output is amplified such that when passed through the secondary winding, it bucks the magnetic flux in the core to zero, preventing core saturation. The principle works not just at DC, but at all frequencies for which the Hall effect device's signal can be amplified cleanly (typically <200kHz).

The combined probe's output is derived from the Hall effect device at low frequencies and the transformer at high, with the two great advantages that the core doesn't saturate (until the amplifier saturates and can no longer buck the flux), and the transformer can be optimized for high frequencies. Low current ($<50A_{pk}$) combined probes such as the Tektronix TCP312A and TCPA300 amplifier can operate from DC to 100MHz.

Load voltage was monitored and triggered by Ch1 of a Tektronix TDS3032 oscilloscope. Ch2 monitored the current probe, allowing it to indicate g_2 current polarity as well as amplitude, so the first test alternated the current probe between g_2 and anode to determine whether the signal currents had the same polarity. They had, so that answered the first question; g_2 contributes power to the load, and is not just an input for feedback.

Adjusting Output Power

The next test adjusted output power while monitoring load voltage and g_2 current. The oscilloscope made automated RMS and Mean measurements on the first cycle after trigger (CycRMS and CycMean). Making cycle measurements avoids the error that would arise from making these measurements over a non-integer number of cycles—as would be caused by measuring over the entire waveform record (**Figures 10-13**).

Quiescent g_2 current is a nominal 4mA, and as can be seen from the 1W, 2W, 4W, and 8W succession of oscilloscope traces, increasing output power increases g_2 current swing about quiescent (CycMean), but while this current can easily rise, it cannot fall below zero, causing considerable distortion that was quantified by the dScopeIII (**Table 2**).

Table 2 shows that g_2 current distortion was sufficiently high that there were no concerns regarding the current probe's (measured) 1kHz 0.2% THD residual.

Efficiency is important in output pentodes, so designers aim to increase I_a/I_{g2} ratio beyond the typical 4:1 seen in small signal pentodes, with the EL84 exceeding 9:1. Distortion at individual anodes is likely to be of the order of 3% approaching full power (largely cancelled by push-pull action in the output transformer), so although the g_2 direct current is one tenth of anode current, the magnitude of its distortion current is comparable to the anode's distortion current, but do they cancel?

If cancellation between anode and g_2 distortion occurs, it should be observable by comparing signal current harmonics between the anode only section and the common section of one half of the output transformer. The output stage was set to deliver 10W at 0.64% THD, while the current probe/analyzer was moved between the anode (a) and common sections (a + g_2) (**Figure 14**).

The evidence could scarcely be clearer; odd harmonic distortion is very nearly absent in the common section (a + g_2), whereas returning the current probe to the anode only section causes odd harmonics to rise sharply. The current probe is not distortionless, so comparison

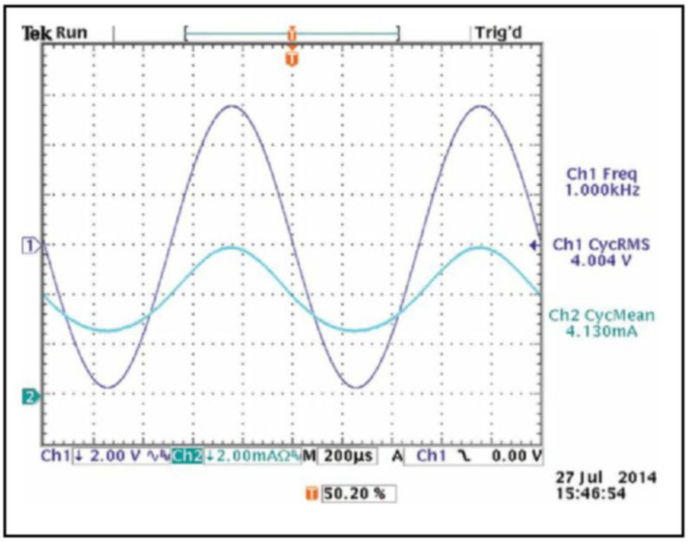

Figure 10: Ch1 (blue, 2V/div.): load voltage. Ch2 (cyan, 2mA/div.): I_{g2}. 1W into 16Ω.

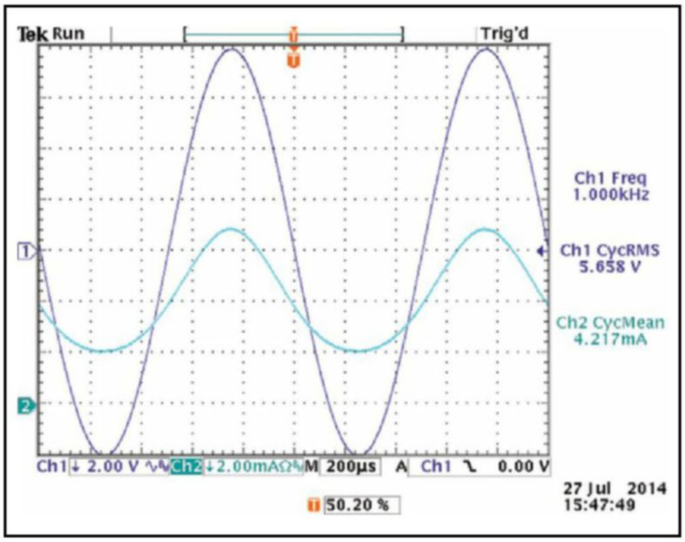

Figure 11: Ch1 (blue, 2V/div.): load voltage. Ch2 (cyan, 2mA/div.): I_{g2}. 2W into 16Ω.

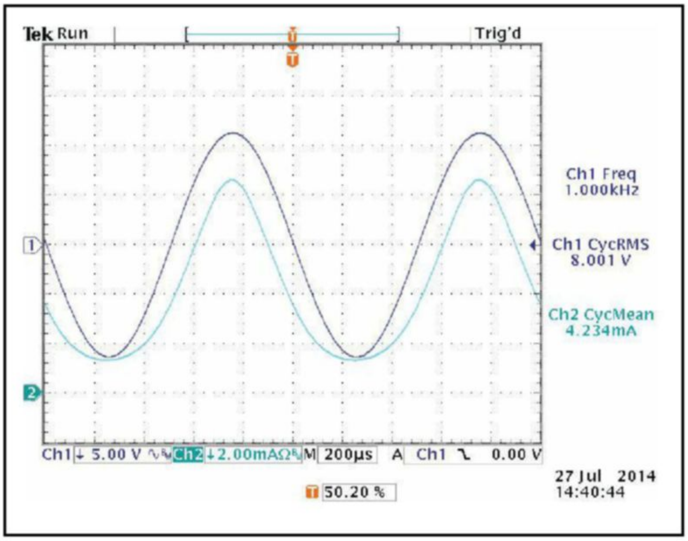

Figure 12: Ch1 (blue, 5V/div.): load voltage. Ch2 (cyan, 2mA/div.): I_{g2}. 4W into 16Ω.

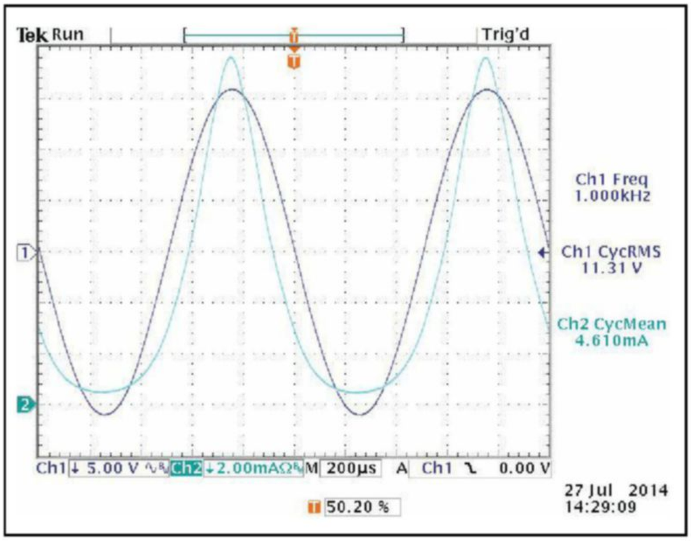

Figure 13: Ch1 (blue, 5V/div.): load voltage. Ch2 (cyan, 2mA/div.): I_{g2}. 8W into 16Ω.

was restricted to harmonics showing >10dB change between probe positions.

Referring to Figure 14 and starting with the fundamental, we see a 1dB increase in amplitude going from a to (a + g_2), which we expect because we had earlier noted that the g_2 current was of the same polarity and contributes power. At the second harmonic, we again see a 1dB increase going from a to (a + g_2). At the third harmonic, we see significant anode current, but near-perfect cancellation causes (I_a + I_{g2}) to fall by 28dB. The fourth and fifth harmonics broadly repeat the trend of second and third. Noting that the reduction of distortion was better than expected from negative feedback and gain considerations, various commentators in the 1950s had speculated that distortion cancellation was involved, but they lacked confirmatory evidence.

1kHz Output Power into a 16Ω Load	Voltage THD into 16Ω Load	g_2 Current THD
1W	0.20W	8.5%
2W	0.27W	12%
4W	0.34W	19%
8W	0.50W	33%

Table 2: EL84 g_2 current distortion as a function of output power

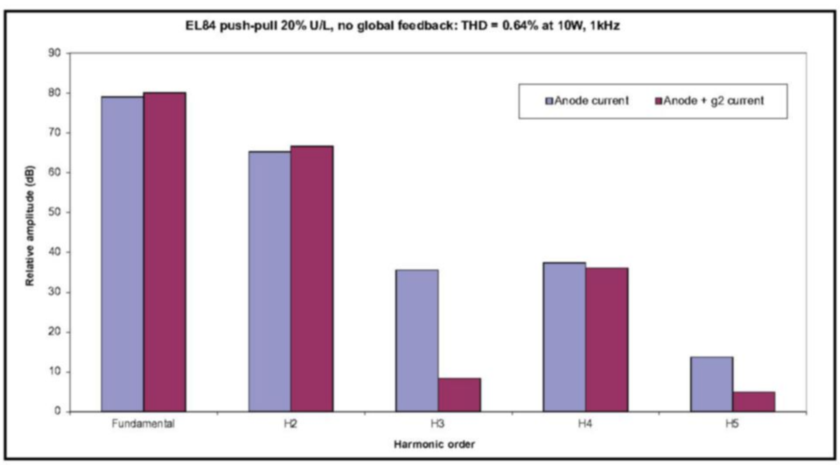

Figure 14: Harmonic amplitudes compared between anode and common (a+g_2) sections

About the Author

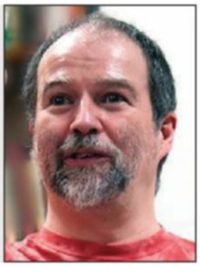

Morgan Jones began professional life as a broadcast engineer at the BBC. Upon leaving, he wrote *Valve Amplifiers*, first published in 1995 and now on its fourth edition. For a day job, he taught Broadcast Engineering, but now (semi-retired) he designs low-noise test and measurement equipment and thinks a nanoamp is a large current. When not in the lab, he enjoys making accurate swarf.

References

[8] R. Moers, "The Ultra-Linear Power Amplifier," *Linear Audio*, Volume 2, pp. 11-45, September 2011.

Distortion Cancellation

Based on the evidence, we can make two important observations regarding distortion cancellation in an ultra-linear stage:

- Cancellation only occurs when harmonic current amplitudes from anode and g_2 are of equal amplitude but opposite polarity, and this occurs for odd harmonics only.
- Cancellation occurs in the common section of the output transformer irrespective of the tapping point, but the signal delivered to the load comes from the entire primary, not just the common section that benefits from cancellation.

Taking the two observations together, we can state that while moving the transformer tapping point toward the anode reduces odd harmonic distortion, it is the equality of the particular valve's anode and g_2 current harmonics that is critical, not the output transformer's precise tapping point. There is no need to forego the flexibility of an output transformer primary composed of identical sections that can be configured in series with a 50% tap, or in parallel to match lower impedance triodes. Noting that cancellation occurs between the anode and the g_2 of a single valve, the technique would work in a single-ended amplifier.

It has been suggested [8] that a critical tapping point might deliver straight anode characteristics intermediate between the convex pentode and concave triode, and further that this would imply zero distortion (which would be true if considering I_a against V_{gk} mutual characteristics). For completeness, this author measured a bogey PCL83 pentode's anode characteristics at 10% voltage tapping intervals from 0% (pentode) to 100% (triode), with V_{gk} at 2V intervals from 0V to -12V (**Figures 15-25**).

As the inventor Alan Blumlein pointed out, the most interesting behavior occurs between 25% and 50%, with the >50% curves being barely indistinguishable from pure triode, although it takes until the 100% set before the slight convexity near the origin of the V_{gk} = 0 curve disappears.

In terms of distortion, the important feature of output characteristics (I_a against V_a for various V_{gk}) is the curve spacing when crossed by a load line, and the shape of individual curves is interesting, but ultimately irrelevant.

To conclude observations regarding the tapping point, we note that a 100% tapping point corresponds to triode operation (significantly reducing maximum anode swing and power), so a given pentode might achieve a given power at reduced distortion using a 43% tap simply because it is further away from full power/overload than when operated as a triode, but greater maximum power using a 20% tap due to increased anode swing (closer to pentode operation), and it is these considerations that provoked the Philips/Mullard 20% and 43% recommendations, not critical output transformer cancellation.

The Mullard EL84 datasheet states 15W for 20% tap and 11W for 43% tap, amounting to a 1.35dB maximum power difference that would be almost inaudible but is good for marketing. For practical amplifiers, anywhere from 20% to 50% voltage tap will do nicely, with perhaps a preference for 50%.

Remembering that the technique relies on distortion cancellation between anode and g_2, this requires a constant I_a/I_{g2} ratio, but aligned grid audio beam tetrodes (e.g., KT66 and KT88) have very variable I_a/I_{g2} ratios; not just sample variation, but also variation with I_k (**Figure 26**).

The Comparison

The comparison shows that the EL84 and the EL34 have very nearly ideal characteristics whereas the two aligned grid types show significant deviation of I_a/I_{g2} ratio and cannot be expected to cancel odd harmonic distortion effectively or consistently. The author had only one NOS Mullard EL34 and three

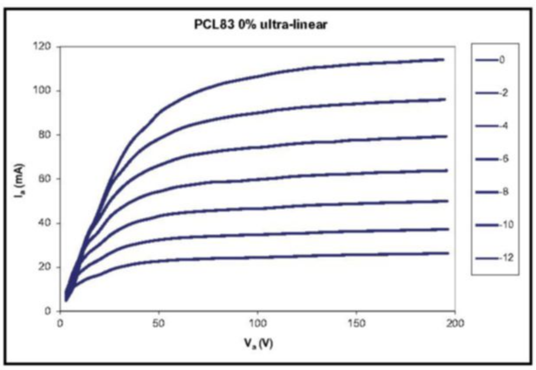

Figure 15: g_2 tap at 0% (pentode)

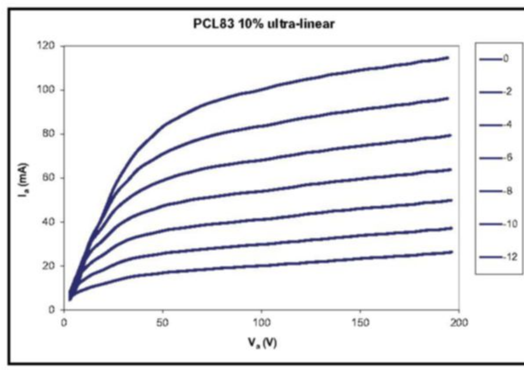

Figure 16: g_2 tap at 10%

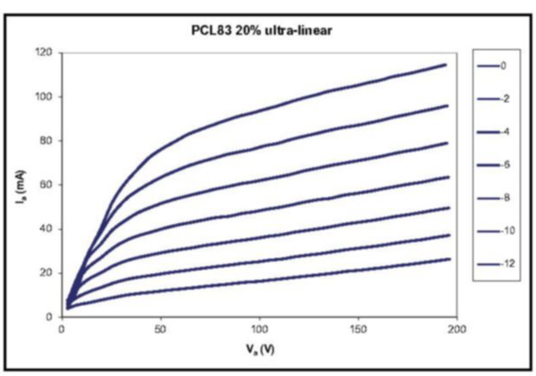

Figure 17: g_2 tap at 20%

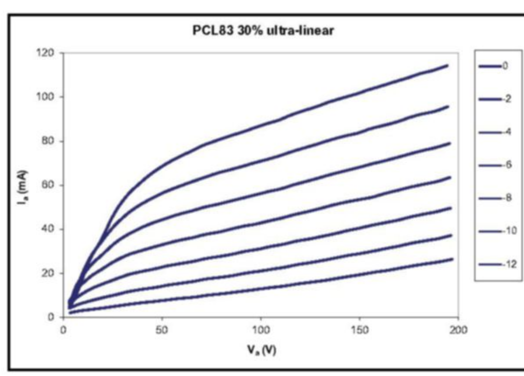

Figure 18: g_2 tap at 30%

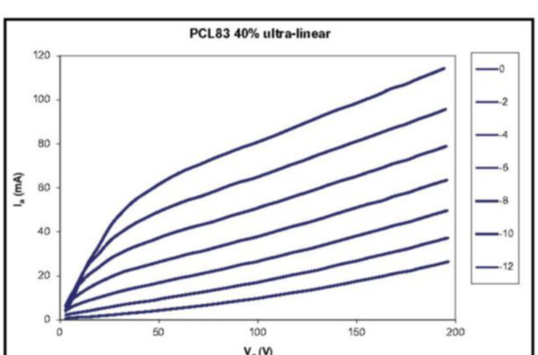

Figure 19: g_2 tap at 40%

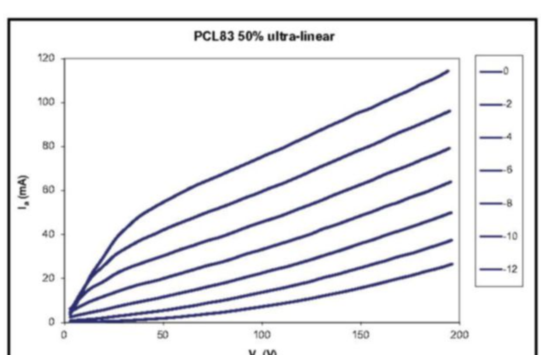

Figure 20: g_2 tap at 50%

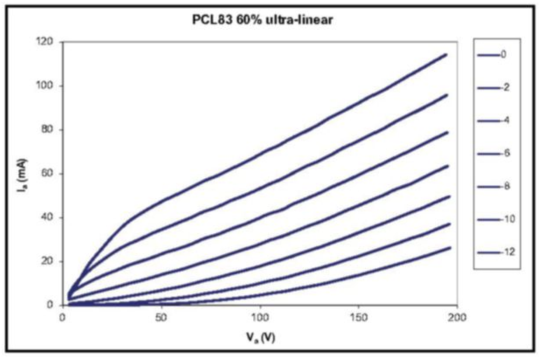

Figure 21: g_2 tap at 60%

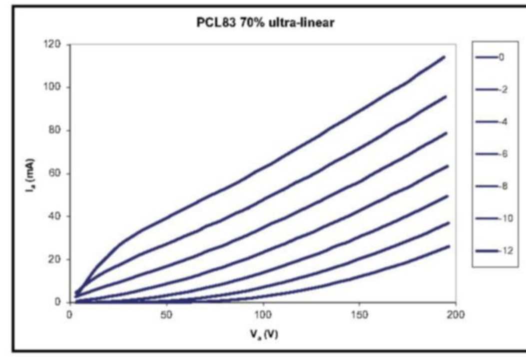

Figure 22: g_2 tap at 70%

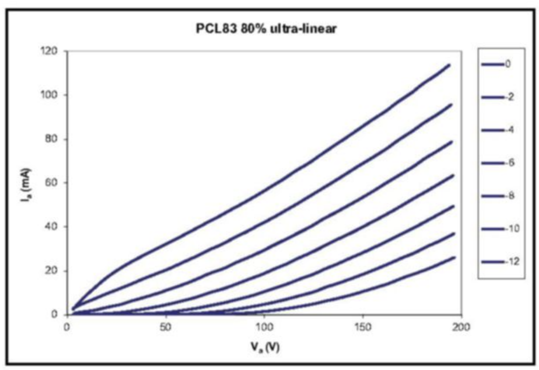

Figure 23: g_2 tap at 80%

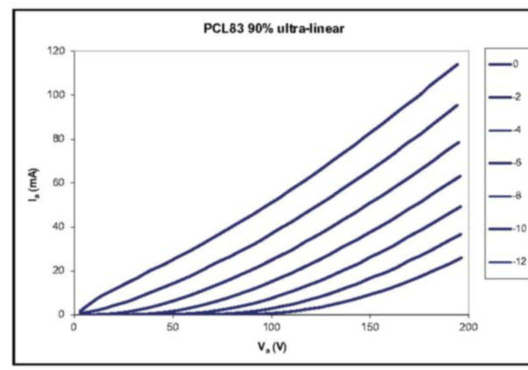

Figure 24: g_2 tap at 90%

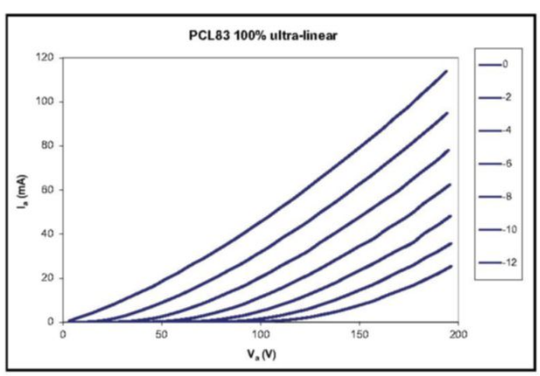

Figure 25: g_2 tap at 100% (triode)

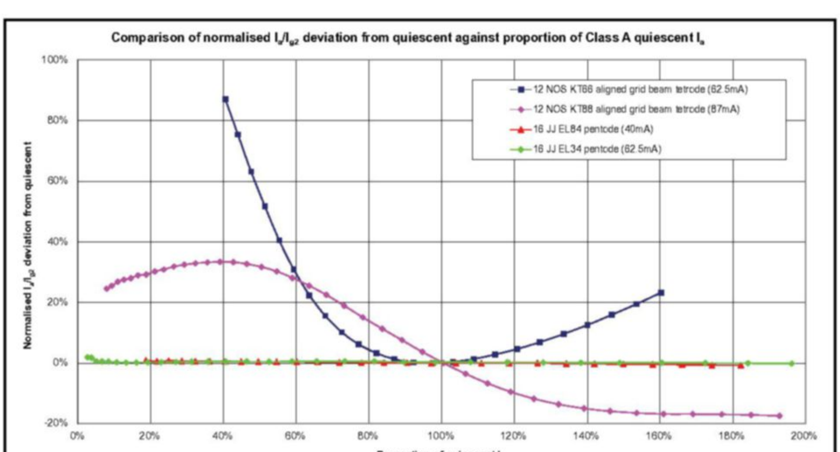

Figure 26: I_a/I_{g2} ratio compared for audio valves

NOS Mullard EL84, hence measurement of current production, which are presumed to be replicas of the original Philips/Mullard designs.

Although pentodes inherently have a reasonably constant I_a/I_{g2} ratio, measurement shows that the EL34 and the EL84 are particularly good compared to other pentodes, and this author can't help wondering whether Philips/Mullard fully understood the distortion cancellation mechanism in ultra-linear, optimized their audio valves, then kept quiet to maximize the resulting commercial advantage.

In Summary

Summarizing, ultra-linear connection cancels odd harmonic distortion within a device and relies not on critical transformer tapping but on constant I_a/I_{g2} ratio within that device.

Thus, aligned grid devices (most beam tetrodes) should be avoided because their inescapably variable I_a/I_{g2} ratio precludes consistent cancellation, whereas EL84 and EL34 pentodes are the ideal choice because they have the most nearly constant I_a/I_{g2} ratio. Because cancellation occurs within a single device, ultra-linear will work with single-ended or push-pull Class B.

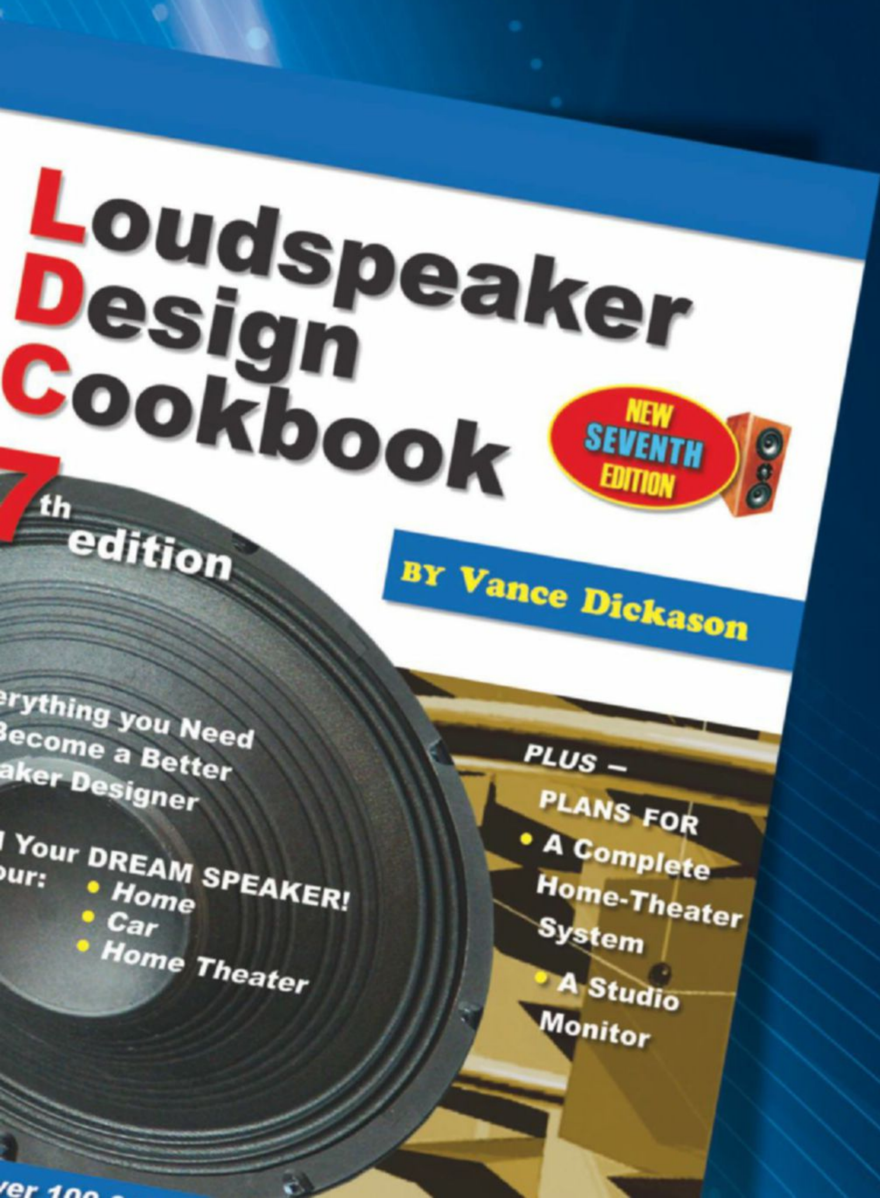

Loudspeaker
Design
Cookbook
7th edition
NEW SEVENTH EDITION
BY Vance Dickason
Home
Car
Home Theater
PLUS –
PLANS FOR
• A Complete Home-Theater System
• A Studio Monitor

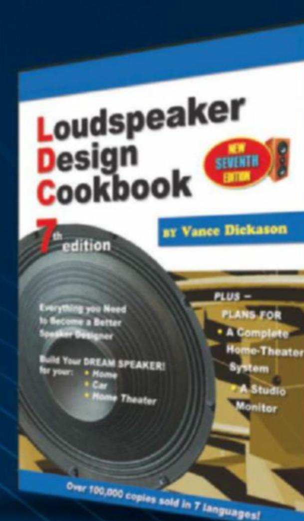

Loudspeaker
Design
Cookbook
7th edition
BY Vance Dickason
Everything you Need to Become a Better Speaker Designer
Build Your DREAM SPEAKER! for your: • Home • Car • Home Theater
PLUS – PLANS FOR • A Complete Home-Theater System • A Studio Monitor
Over 100,000 copies sold in 7 languages!

Cascodes, Folded Cascodes, and Current Mirrors (Part 2)

Current Mirrors and Line Stages

By
Morty Tarr

In Part 1 of this three-part article, Morty Tarr discussed cascodes and folded cascodes. Cascodes are useful to improve high frequency response and device linearity. Folded cascodes add the ability to move the signal level back toward ground. In Part 2, he focuses on current mirrors and line stages.

In addition to cascodes and folded cascodes, another method to improve the performance of a basic amplifier is to use a current mirror. A basic current mirror is shown in **Figure 9**. The inputs and outputs of a current mirror are currents, not voltages. In Figure 9, an input current is applied to PNP transistor Q1.

Q1 functions as a diode, so the current I_{in} is applied across the base-emitter junction of Q1. It is assumed that Q2 is matched to Q1, so Q2 has the same base emitter voltage and therefore the same collector current (however, this is not exactly true, we'll get to that later). For current mirrors, the input is a current sink, and the output is a current source.

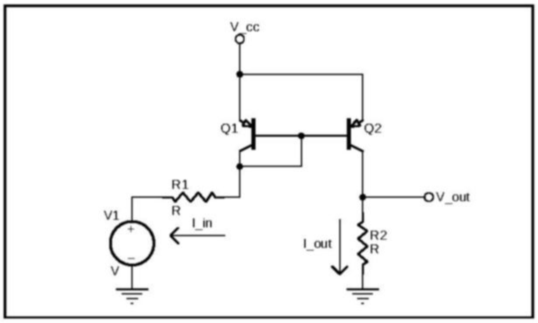

Figure 9: In this basic current mirror, the inputs and outputs are currents, not voltages. Here an input current is applied to PNP transistor Q1.

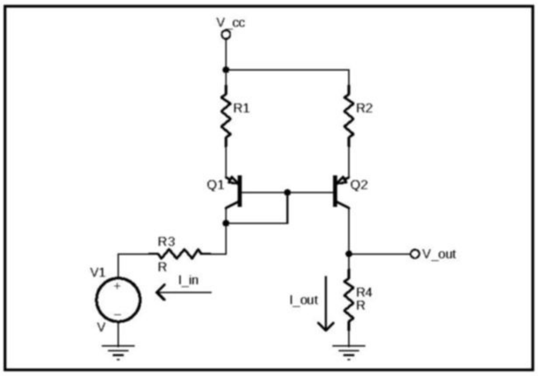

Figure 10: This schematic shows a current mirror with emitter resistors R1 and R2 added.

Current Mirrors

Current mirrors are used extensively in op-amp design, as matched transistors are easy to create and the supply voltages are limited. This is not the case with discrete designs.

So far, current mirrors seem very simple, where can it go wrong? In discrete designs, the matching of the transistors may not be ideal, and the temperature tracking will not be ideal. To minimize the impact of imperfectly matched transistors, we add emitter resistors. These resistors also serve to reduce the noise generated by the current mirror. Remember, the devices do not know they are a current mirror; they have a high impedance collector load (a current) and a small emitter resistance (approximately 26Ω at 1mA with no additional resistor). This is the formula for a very high gain.

Figure 10 shows a current mirror with emitter resistors R1 and R2 added. We also add base resistors to suppress the tendency to oscillate as shown in **Figure 11**. Figure 11 shows a current mirror adapted to discrete design. The input to the current mirror is a current, which can be applied using a voltage source and a resistor, or an active device such as a BJT (Q3), JFET (J1), or MOSFET (not shown).

Note that an active device of the proper polarity can be the input to any current mirror. In the schematics that follow, only a voltage source and resistor will be shown, but in all cases, an active device could be used instead.

The value of the emitter resistors should be as high as you can stand it. If it is a current mirror operating at 1mA, a 1KΩ resistor will drop about 1V and dissipate 1mW. If the current mirror is operating at 10mA, a 1KΩ resistor will drop about 10V, and that's going to be too much. So something between 100Ω and 500Ω might be appropriate. Just remember, as the resistor gets smaller, the ability to minimize any mismatch decreases and the noise gain increases. As in most designs, there are many trade-offs!

The circuit in Figure 11 will not make a very good current mirror for audio. A good current mirror should have a very high output impedance since the output is a current. The output impedance of the circuit in Figure 11 is in the tens of kilohms. We can increase this with, guess what, a cascode.

Also, we want Q1 and Q2 to match, but their operating conditions do not match. Q1 has a collector-base voltage 0V and Q2 has a much greater collector-base voltage. We can add a cascode to just the output and make a Wilson Current Mirror (there are both three and four transistor versions of the Wilson Current Mirror), or we can add a cascode on both sides as shown in **Figure 12** and make a cascode current mirror.

The cascode raises the output impedance by about two orders of magnitude, greatly improving performance. It also enables Q1 and Q2 to have very similar operating conditions, which also improves performance.

Note that the current gain is less than 1.00 by a factor of four base currents (4% if Beta = 100). This would not be problem if Beta was constant (but Beta changes with collector current). Also, Q1 and Q3 are operating at a collector-base voltage of 0V, which is not a good operating point for a transistor in a linear circuit. We can add LEDs to the design to increase the collector-base voltage of Q1 and Q3.

Since the transistors draw only microamperes of base current, LEDs are a good (low noise) choice for increasing the collector-base voltage. Green LEDs drop about 1.7V each in this configuration (1.9V to 2.1V with a 5mA current).

Bipolar junction transistor (BJT) current mirrors can be made with NPN or PNP transistors. **Figure 13** shows a PNP current mirror that will reflect against the positive rail. An NPN current mirror will reflect against the negative rail. BJT current mirrors suffer from Beta changing with collector current. While the change is small, it creates second harmonic (2H) distortion. A metal-oxide semiconductor (MOS) current mirror can mitigate this problem.

There are other forms of current mirrors, but due to limited space we cannot discuss them here.

Figure 14 shows the current mirror of Figure 13 with the BJTs replaced by metal-oxide semiconductor field-effect transistors (MOSFETs). The MOSFETs contribute much lower levels of 2H distortion compared to BJTs.

When I use current mirrors in audio circuits, the best matched devices become Q1 and Q2. The match

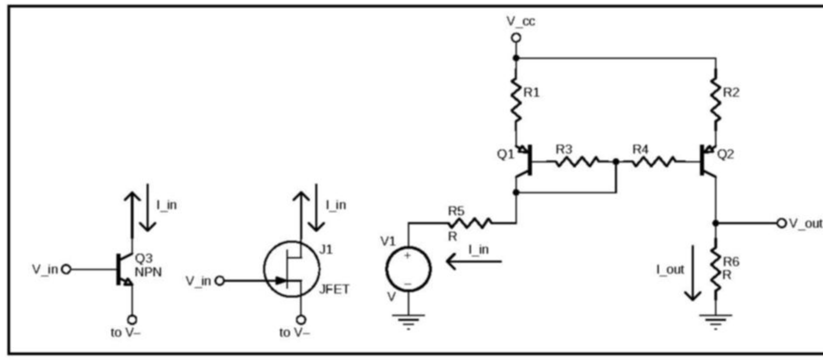

Figure 11: Base resistors have been added to suppress the tendency to oscillate and shows a current mirror adapted to discrete design.

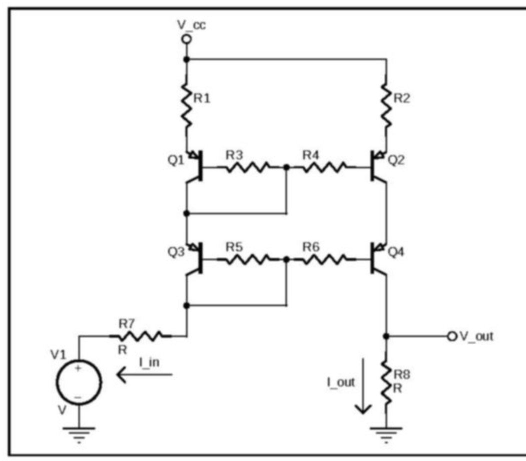

Figure 12: A cascode on both sides makes a cascode current mirror.

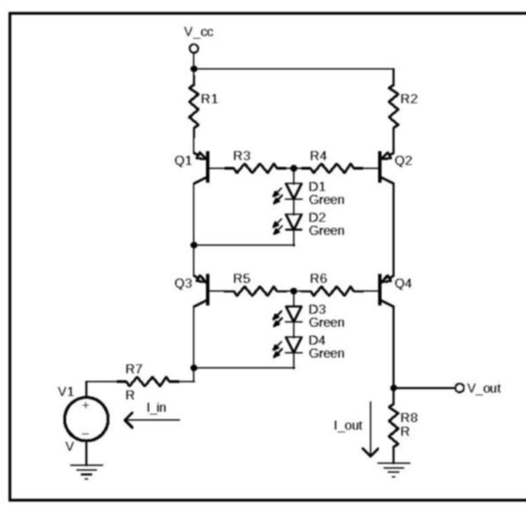

Figure 13: A PNP current mirror like the one shown here will reflect against the positive rail.

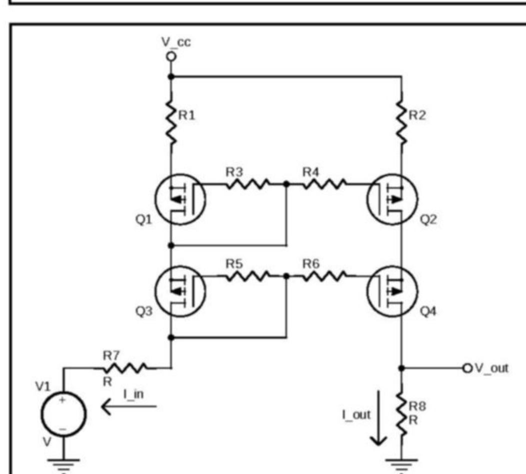

Figure 14: Here is a current mirror of Figure 13 with the BJTs replaced by MOSFETs.

requirements for Q3 and Q4 do not have to be as good. It is better to use high voltage devices, especially for Q3 and Q4 due to channel modulation effects.

There are significant differences between current mirrors and a folded cascode. Current mirrors are non-inverting and current mirrors can have gain.

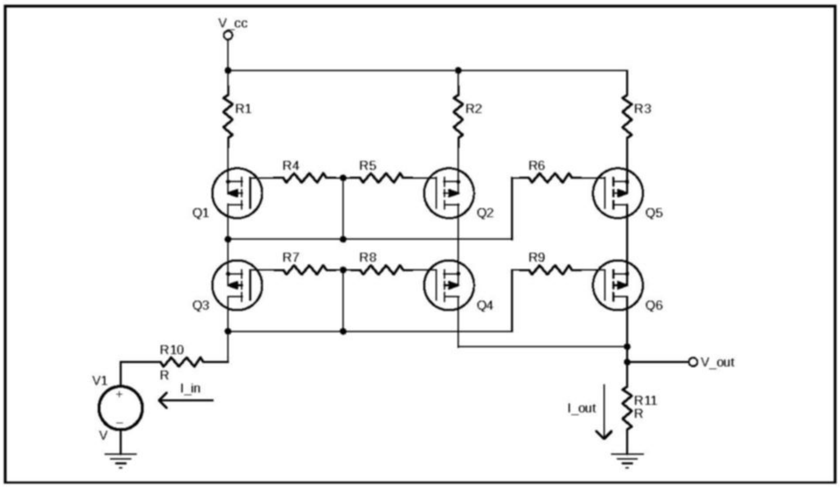

Figure 15: This is the same current mirror as shown in Figure 14, but with Q5, Q6, and R3 added.

A current mirror with gain is shown in **Figure 15**. Figure 15 shows the same current mirror as Figure 14, but with Q5, Q6, and R3 added. Q5 is matched to Q1 and Q2, Q6 is matched to Q3 and Q4. This configuration has double the output current of the circuit shown in Figure 14. The concept can be extended to include additional output mirrors so long as reasonable matching can be maintained.

Folded Cascode Line Stage

An amplifier can be built using cascodes, folded cascodes, and/or current mirrors. An example of a folded cascode line stage is shown in **Figure 16**.

Note that the current sources are biased from a string of forward-biased LEDs. I chose green LEDs, which are about 2V each. LEDs provide a low noise, low impedance bias point for the current sources. Zener diodes above 5.6V are actually avalanche diodes and contribute significant noise, which is hard to filter out.

Figure 16 shows the "block diagram" of a line stage using cascodes and folded cascodes. The block diagram is easier to follow than the actual schematic (**Figure 17**).

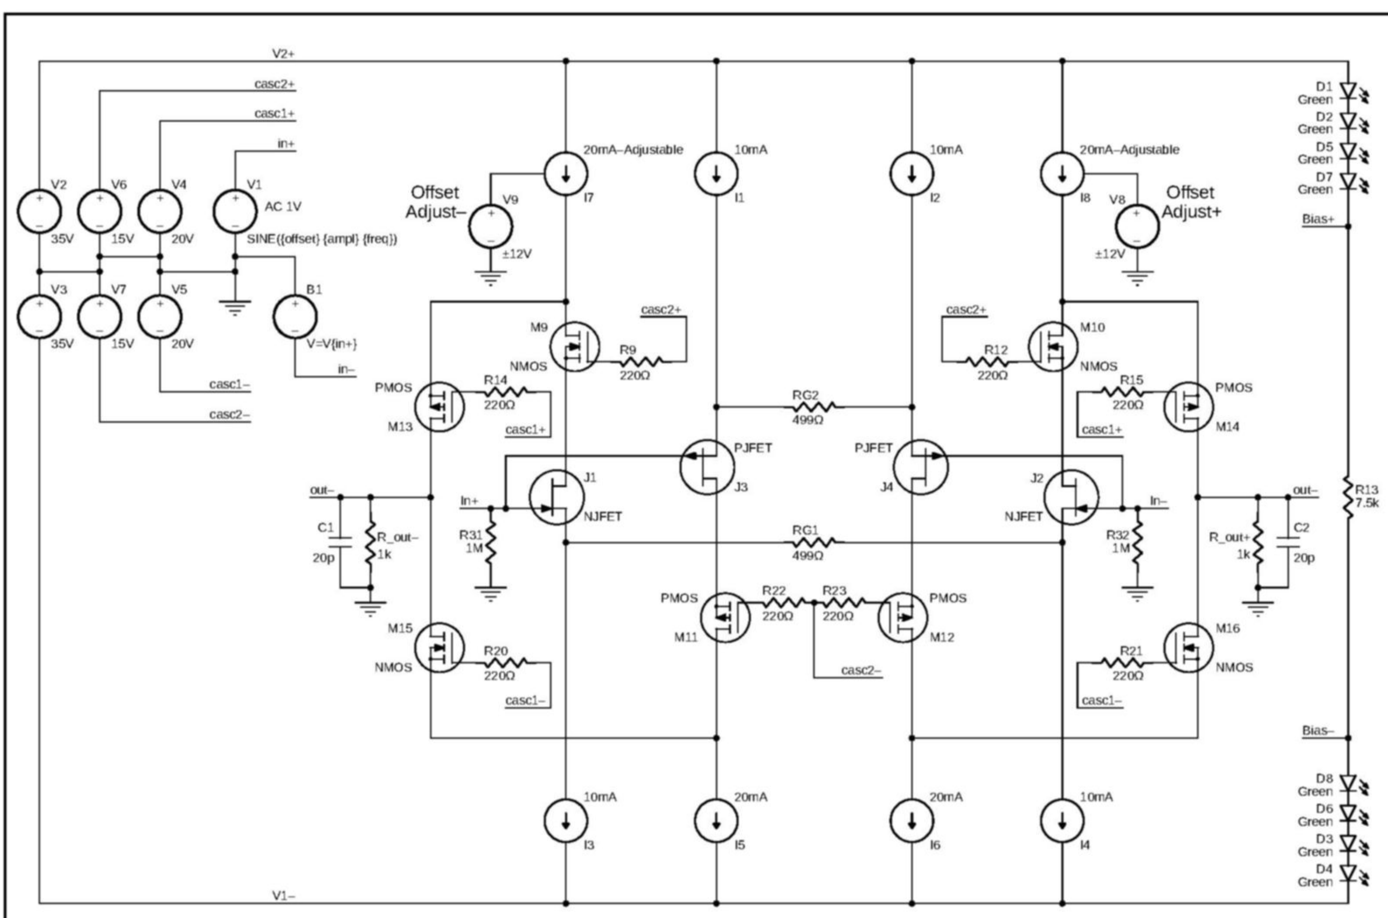

Figure 16: This block diagram shows a line stage using cascodes and folded cascodes.

A note about terminology: The gate of a device will be designated with a G suffix, the source an S suffix, and the drain a D suffix. The gate of J1 will be J1G, the source, J1S, and the drain J1D.

The input stage includes J1, J2, J3, J4, the input JFETs. J1 and J2 are biased by current sources I1 and I2. J3 and J4 are biased by I3 and I4. The current sources are set to 10mA, which is good operating current for the JFETs. The drains of the JFETs are cascoded by M9, M10, M11, and M12. The cascodes keep the JFET drain voltage constant and also limit the dissipation in the JFETs.

Since any errors in this circuit show up as distortion, we want to treat the JFETs as “nicely” as possible. J1 and J2 are Linear Systems LSK389Ds, dual, high transconductance N-Channel JFETs. J3 and J4 are LSJ74Ds. Since the JFETs are operated at 10mA, we want I(DSS) of the JFETs to be at least 12mA to account for signal currents. Since Linear Systems does not offer a dual, high-transconductance P-Channel JFET, the LSJ74s should be matched for I(DSS).

JFETs J1–J4 form a four-quadrant voltage to current converter. J1 and J2 are matched N-Channel JFETs, biased by 10mA current sources connected to J1S and J2S. If the voltage at in+ (the gate of J1) goes positive, the voltage at J1 source (J1S) rises and current flows in RG1 from J1S to J2S. Since each JEFT is biased by current sources, the current in the JFETs has to change as the current in I1 and I2 has to be constant. The current in J1 will increase with a corresponding decrease in the current in J2. This change in current is reflected in the JFET drains, J1D conducting more current and J2D conducting less current. The difference between these currents is flowing in RG1.

For example, if in+ is at +250mV and in- at -250mV, there is a 500mV difference between the

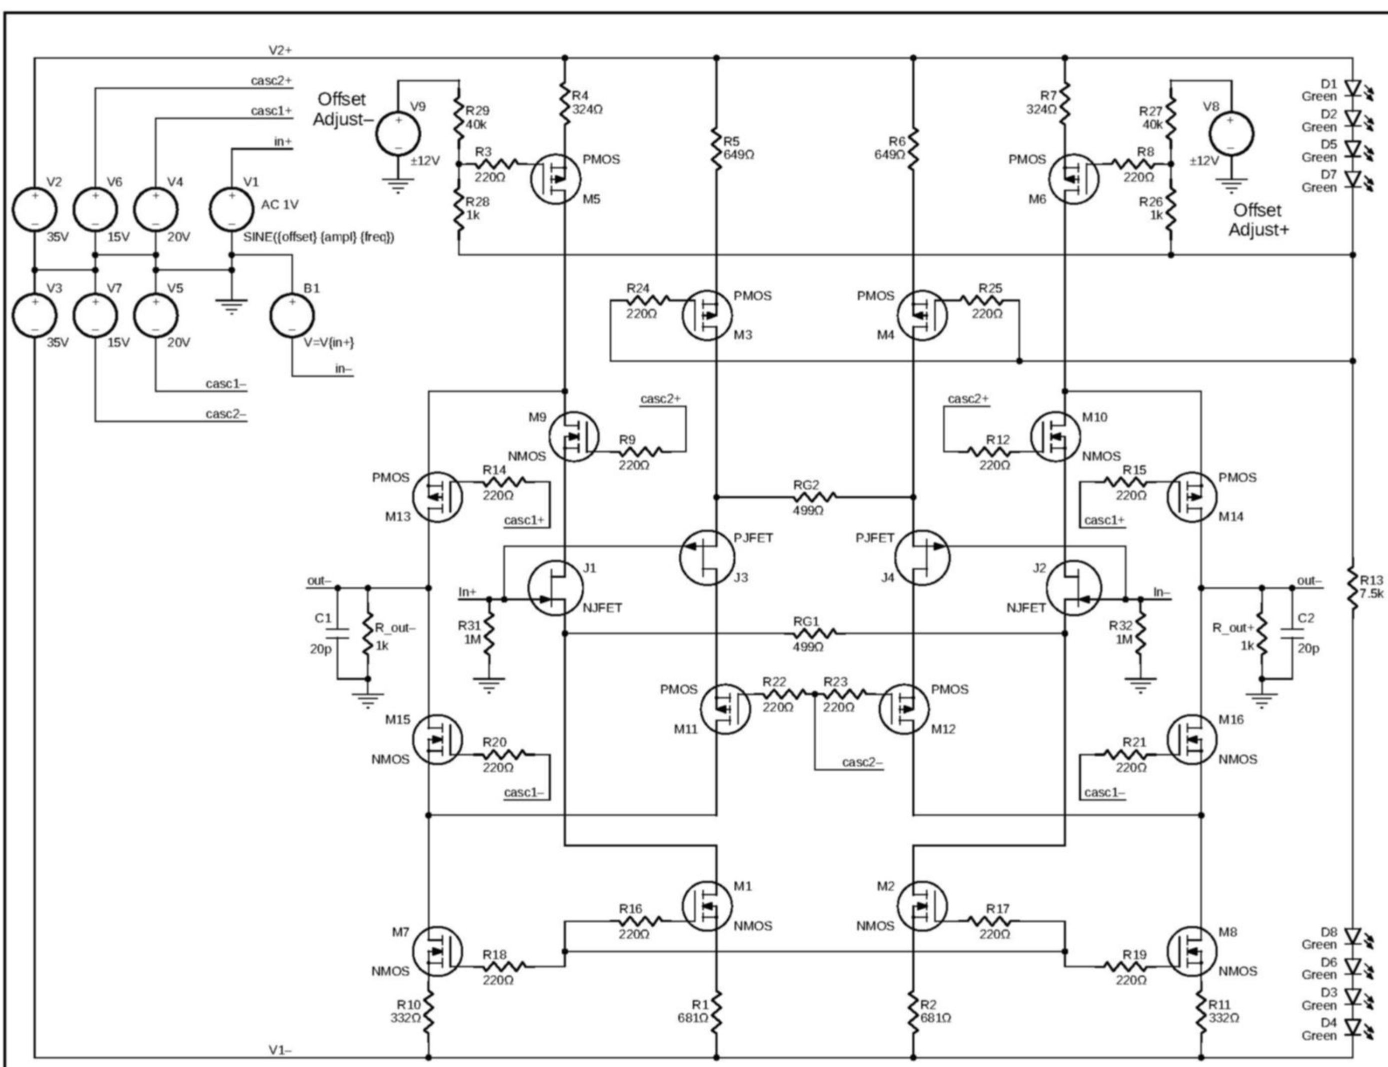

Figure 17: An actual schematic of a line stage using cascodes and folded cascodes.

gates of J1 and J2. Since J1 and J2 are matched, there is a 500mV difference between J1S and J2S, causing 1mA to flow through RG1 from J1S to J2S. Therefore, the current at J1D must be 11mA and the current at J2D must be 9mA, as I1 and I2 conduct 10mA each.

The current in output cascode M13 is 20mA – 11mA = 9mA, and in M15 it is 20mA – 9mA = 11mA. The same situation happens with J3 and J4, except the polarities are reversed. This results in the current in M14 equal to 11mA and the current in M16 equal to 9mA. The 2mA difference between M13 and M14, and between M15 and M16 is impressed across the 1KΩ output resistors resulting in out+ = 2V and out- = -2V, or a 4V balanced output.

When the voltage at in- increases, the current at J2D increases and the current at J1D decreases. In a similar way, J3 and J4 (P-Channel JFETs) convert the voltage difference between in+ and in- to a current difference in J3D and J4D.

The output (folded) cascodes invert the current and apply the current difference to the output resistors (R_{out+} and R_{out-}), creating an output voltage.

The drains of all four JFETs are biased to 20mA (10mA for the JFET and 10mA for the output (folded) cascode. I5 and I6 are fixed at 20mA, while I7 and I8 are adjustable to set the output offsets to 0.00V.

Note that the change in drain-source current in the JFETs is equal and opposite to the change in drain-source current in the MOSFET (folded) cascodes. There is no current gain in a folded cascode.

This is a current mode amplifier, as the only nodes where the voltage changes significantly with input signal are the input nodes and the output nodes. It is also a transconductance amplifier as the output of the amplifier is a current. We impress this current across a resistor to create an output voltage. These are resistors R_{out+} and R_{out-}.

Capacitance at the output of this amplifier will limit the frequency response but it will not cause oscillation. Note capacitors C1 and C2, which provide frequency response compensation to eliminate peaking. These are the only capacitors in the design.

With R_{out} = 1kΩ and a load capacitance of 500pf (which is generally more than 10′ of shielded cable), the -1dB point is greater than 150kHz and the -3dB point is over 300kHz. Reasonable load capacitance should not be a problem.

For applications that need a lower output impedance, a diamond buffer, or other type of emitter or source follower can be added.

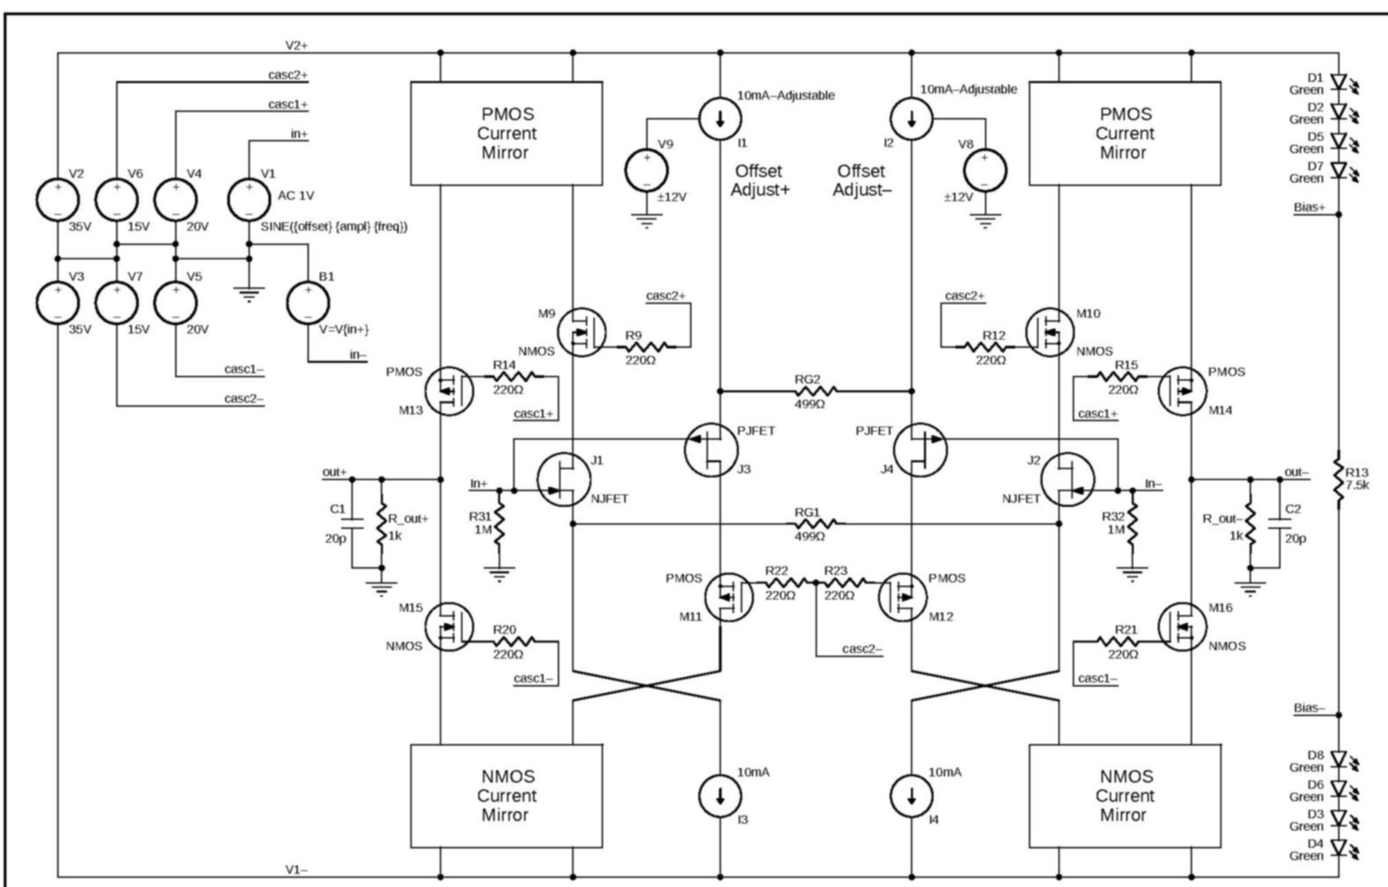

Figure 18: This is a Current Mirror Line Stage block diagram.

Here are some details based on the folded cascode line stage schematic shown in Figure 17:

- When an input voltage is present, the current in J1 and J2 becomes unbalanced as does the current in J3 and J4. Since J1 and J2 are matched, and J3 and J4 are matched, the difference between J1S and J2S is the same voltage as the difference between J3S and J4S. We now have current flowing in RG1 and RG2, The sum of the currents in RG1 and RG2 is impressed across the output resistors, R_{out+} and R_{out-}. The gain of the amplifier is approximately $2 \times R_{out}/RG1$, assuming RG1 = RG2 and $R_{out+} = R_{out-}$ (0.1% matching recommended). In this case, the gain is 2×1000/250 = 8 V/V or 18dB. We could increase the gain by making the output resistors (R_{out+} and R_{out-}) larger, say, 2KΩ, or making the gain resistors (RG1 and RG2) smaller.
- With 499Ω gain resistors, an input of $1V_{PP}$ will cause 2mA to flow in each of the gain resistors, causing a ±2mA imbalance at the outputs. With 1KΩ output resistors this is ± 2V or $4V_{PP}$ at each of out+ and out- or a differential output voltage of $8V_{PP}$, which is equivalent to $5.65V_{RMS}$. This is more than enough to drive most power amplifiers to clipping. The actual gain is about 85% of the theoretical gain due to the limited transconductance of the JFETs and MOSFETs. So, instead of 8V/V, we get about 6.65V/V or 16.5dB gain. I have found that deviations from the standing DC current of ±10% maintain very low distortion levels even without global feedback.

Selecting Devices

If you were to build this circuit, use JFETs where I(DSS) is at least 12mA for our 10mA design. It is best if the pairs of JFETs are matched. Linear Systems LSK389C and LSK389D are candidates for J1 and J2. Linear Systems LSJ74C and LSJ74D are candidates for J3 and J4. In this case, the JFETS J3 and J4 should be matched for I(DSS) on the bench.

I have used Supertex (now Microchip) VP2450/VN2450 or Zetex (now Diodes, Inc.) ZVP4424/ZVN4424 for the MOSFETs. Both seem to work equally well. I am certain there are other choices that will also work well. Note that these are high voltage devices. It is best to buy about twice the required number so that they can be matched and used in matched pairs.

Current Mirror Line Stage

We could also design a Current Mirror Line Stage using a similar technique. The block diagram is shown in **Figure 18**. In the Current Mirror Line Stage, the input JFETs function the same as they do in the Cascode Line Stage. The difference is that the folded cascodes are replaced by current mirrors. We use cascode current mirrors for accuracy and low distortion. The current mirrors are non-inverting where the folded cascodes are inverting, so the out+ and out- nodes switched sides.

As in the Cascode Line Stage, the input JFETs are cascoded to limit their power dissipation and to keep their drains at a constant voltage for best performance. The output side of the current mirrors are also cascoded for the same reasons.

Here we have to trim the JFET current sources to balance the outputs at 0.00V. This is done by having one set of current sources adjustable while the other set is fixed. The adjustment can be continuous using an op-amp integrator, or manual using a potentiometer. An op-amp integrator will limit the low-frequency response, but we can set the roll-off to well below 1Hz so it has minimal effect on the audio signal path.

This amplifier, like the Cascode Line Stage is fully differential (balanced) in and out. There are four current mirrors in the Current Mirror Line Stage schematic (**Figure 19**), a P-Channel Current Mirror for each of out+ and out-, and an N-Channel Current Mirror for each of out+ and out-.

Here again, the input JFETs act as voltage to current converters. The JFET drain currents are loaded by the current mirrors, which reflect the current toward the opposite rail. A P-Channel and N-Channel Current Mirror are combined through cascode transistors to create an output current. The cascode transistors stabilize the voltage at the output of the current mirrors, improving the linearity. The current differences are impressed across the output resistor to create an output voltage.

To control the offset, we have to adjust the current sources. The N-Channel current sources are fixed

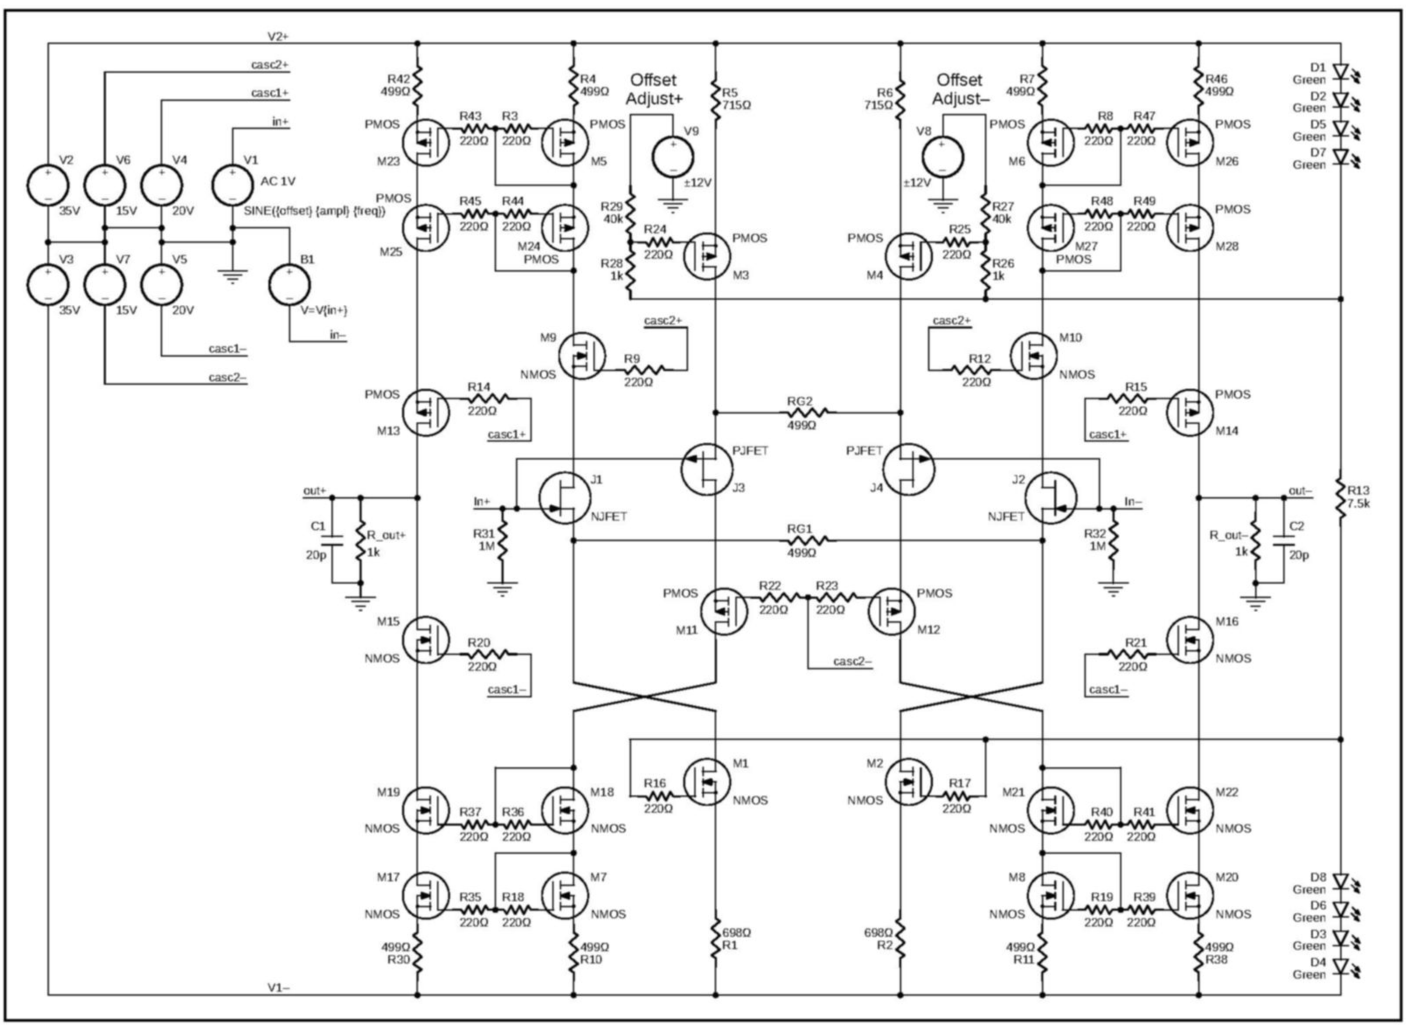

Figure 19: There are four current mirrors shown in the Current Mirror Line Stage schematic.

(M1 and M2), and the P-Channel current sources are adjustable. The adjustment can be made via a potentiometer or an op-amp configured as a non-inverting integrator.

Each has its own advantages, the op-amp continuously trims the offset to 0.000 V, but gives the amplifier a low-frequency roll-off (below 1Hz). The potentiometer setting cannot be adjusted continuously, so there will be a small residual offset at the output, but no low-frequency roll-off.

The resistor values for the current sources in Figure 17 and Figure 19 are dependent on the V(GS) of the MOSFET devices used. These values may have to be modified if different devices are used.

Observations Thus Far

For both of these line stages, I have measured distortion levels of better than -90dB (0.0032%) with $2V_{PP}$ output, single-ended in and out. At $2V_{PP}$, there are no distortion components higher than the third harmonic (3H). At $8V_{PP}$ out, the distortion is still better than -80dB (0.01%).

My measurement system consists of an HP339A Distortion Analyzer and an HP3580A Spectrum Analyzer. Since most of my consulting work is not related to audio, I cannot justify the expense of an Audio Precision measurement system. I believe the distortion specs previously quoted would be significantly better if measured on an Audio Precision system.

Next Month

In Part 3, we will discuss some of the details involved in building these line stages including matching transistors, DC servos, and capacitance multiplier power supplies.

About the Author

Morty Tarr has more than 30 years' experience in Electrical Engineering. Morty graduated from Cornell University in 1972 with a Bachelor of Arts in Physics, and a minor in Electrical Engineering. He continued his studies at Tufts University in Electrical Engineering.

Morty began his career at recording studios in NYC and Boston. Then he became interested in the design of electronic equipment and made the transition from audio to electronic design. He has extensive experience in the fields of video, audio, test and measurement, networking, and wireless (primarily Wi-Fi and Bluetooth). More recently, from 2006 to 2015, Morty led a team at Bose responsible for advanced development for a $2 billion business. He set the direction for the group and for technology strategy spanning wireless, digital, analog, and system architecture across multiple product lines. Prior to Bose, he worked for LTX developing GHz test systems, at Octave Communications developing conference call bridges, and at Avid Technology. Prior to Avid, at Data Translation, he created the first 24-bit resolution A/D for a PC for Chromatography applications, and then led the team that developed the first Media 100 video editing system.

Morty currently has 20 patents and has spent most of his professional career developing first-in-class products. Innovative ideas and new concepts come naturally to him.

Audio Memory

The History of the American Magnetic Head Industry

By
Scott Dorsey

An article several years in the making, it could as easily have resulted in an interesting reference book, given the extensive scope of the topic. Scott Dorsey shares the essentials of his research during this journey down memory lane into tape recording, centered on the American magnetic head industry.

The history of magnetic head manufacturing is a fascinating one because it touches on so many different aspects of the electronics industries. Not just for audio recorders, by the peak of the industry in the 1970s, there were hundreds of companies manufacturing heads for computer tapes, audio and video tapes, instrumentation, subway cards, news film, teaching machines, toys, programmed typewriters, and a thousand other applications that are mostly gone today.

It is not just the history of audio, it is the history of American industry as a whole, and watching how companies evolved and how people left one company to take the knowledge they learned to a new one is a view into how industry changes and grows. Today, there are a small handful of companies still making heads for professional audio recorders, credit card readers, and bill validation, but it is a tiny industry. This is how we got there.

Note that two of the major figures in this story are John R. French and his father John J. French. For many years I could not keep them straight but doing so is critical to understanding the history.

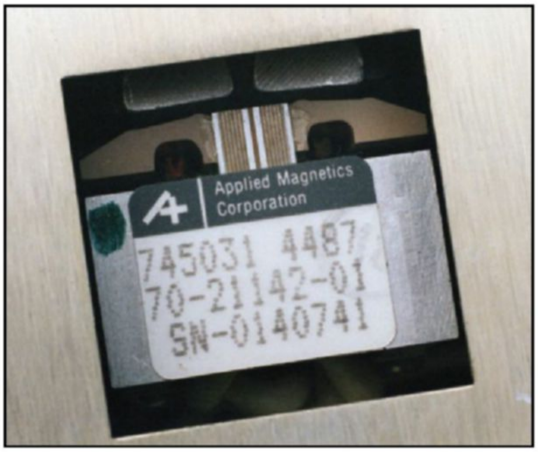

Applied Magnetics Corp. started a number of manufacturing plants abroad in the 1960s and 1970s. Today AM Belgium, a company established in 1969, operates independently and continues to design and manufacture professional magnetic heads for various applications.

The 1930s

Although wire recorders had existed, they used a very different design of head assembly to establish a magnetic field perpendicular to the direction of movement. What made recording heads for tape new and unique is that the field was established parallel to the tape surface. Doing this involved a head with a winding on it that was shaped like a ring (either square or round) and a gap cut in the ring. Tape moving past the gap would become magnetized, and the thin gap meant the field would be sharp and focused for good high-frequency response. Heads like this were first developed by AEG in Germany, which patented configurations in 1928 and 1935, with the modern “ring head” design as is used today developed by Edward Schueller in 1933. He also was the first to recognize that a larger gap was needed for the record head than the playback head and the first to recognize that a mechanical support was needed to keep the gap constant.

AEG's Magnetophon K1 machine was initially developed as a DC bias machine of dictation quality in 1935, but the system evolved so that in 1940 Walter Weber and Hans Joachim von Braunmuehl redeveloped AC bias, and AEG soon started manufacturing the K4 machines of surprisingly high fidelity.

The heads from AEG resemble our modern heads in most ways... they are a loop made of stamped laminations stacked together with an open gap over which the tape passes, and a nonmagnetic material filling that gap. The laminations have become thinner over the years, the gap material has varied from brass to glass to epoxy, and the nickel alloy used has varied slightly in attempts to improve longevity, but the AEG head is very clearly the progenitor of all audio heads today.

The Brush Development Co. in Cleveland, OH, created a machine based on those German head patents, using DC bias, with dictation grade reproduction, and apparently most of the development was completed by 1940. However, because of the war this was put on hold, and the Brush Soundmirror was not actually released to the public until 1947. A few years later, Brush was purchased by the Clevite Corp., which wanted no part of the audio market, but Brush-Clevite continued making heads for other recorder manufacturers until 1973, when its head division was sold to the manufacturing company Forgflo and moved to Sunbury, PA, forming Brush Industries. Brush continues making heads today, primarily heads for magnetic stripe cards in transit systems and credit card readers.

The 1940s

AEG lost the rights to its patents as part of reparations for World War II, and a lot of German technology was brought back to the US. When Jack Mullin brought the Magnetophon recorder back from Germany and took it to Ampex in 1947, Ampex began making tape heads for its new Model 200 recorder, which were basically copies of the Magnetophon design. They were just a simple U-shaped piece of metal with a winding around them and a gap between the pole pieces over which tape passed. The story of how Bing Crosby financed Jack Mullin's effort to manufacture a recorder based on the AEG designs is well-described elsewhere.

So we can think of the Ampex 200 head as being pretty much the ur-head, the source from which all of the later tape head designs were derived. The Ampex 300 and 350 machines used more precise and refined versions of that design, with a lot of effort going into making the edges of the gap perfectly straight and parallel. As "stacked" stereo and later multitrack heads came in, additional work was needed to make sure the individual gaps were precisely lined up; not only did they need to be exactly parallel to one another so the azimuth was the same on all channels, but they needed to be positioned perfectly above one another so that no unwanted phase shift between channels existed. All tracks recorded at the same time, and they all played back at the same time.

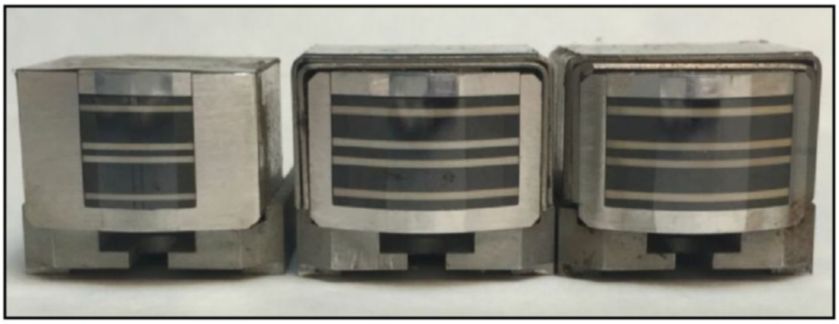

Saki Magnetics provided long-life replacement heads for the Ampex ATR-100 series and many other tape recorders. Pictured here are 1/4" 2-track Saki Magnetics ferrite heads. (Image courtesy of ATR Services Inc.)

The 1950s

In the 1950s, the market diversified from just professional recording heads into computer heads and mass-produced consumer heads. The computer industry wanted more tracks for higher data density (with the Univac Uniservo drives using eight tracks as early as 1951), but audio people soon found a need for higher track counts too. And, of course, video recording and disk drives opened up whole new markets for new head designs.

Bing Crosby wasn't the sort of person to put all his eggs in one basket, so he also financed Eddie Lipps to make heads based on the Ampex design in 1951. They hired Bob Cross to design their heads for many years, which he did until leaving to start Pacific Magnetics in 1978.

Lipps stayed in business in Los Angeles, CA, until the 1990s, after branching out into heads for IRIG instrumentation recorders and logging machines. Government sales were his main revenue source for many years, but the company also made excellent aftermarket replacements for Ampex audio machines.

Meanwhile, as Germany began to try to reconstruct after World War II, Bogen Electronics in Berlin started making recording heads for audio recorders. Although its earlier designs were crude, it advanced fast and made heads for many European open reel and cassette recorders. Bogen is still around today, making card reader heads.

3M was in the business too, and it did some very early research into making heads from ferrite materials [1]. 3M was in the materials business, making ferrite materials as well as tape and it was in its interest for people to use them.

In 1956, Nortronics opened for business with a new philosophy of building standardized heads that were all the same size, no matter what the track configuration was like. The company started out as an OEM, and within a few years it was making heads for a huge range of products from 8-track players and teaching machines to high-end studio recorders. Nortronics also did a good business in replacement heads and Quick-Kits that would allow users to put its standard heads into a wide variety of machines. Nortronics opened its facility right in Minneapolis, MN, and that started a local industry of its very own, although the company moved to the nearby suburb of Dassell, MI, at some point.

Nortronics made cheap commodity heads for consumer gear but it also made some very high-end professional audio heads. The original 3M Dynatrack machines used its brass-faced heads.

Although less well-known than Nortronics, Michigan Magnetics started at about the same time with similar products and a similar design philosophy, taking a lot of its ideas from Nortronics. Michigan Magnetics' factory in Vermontville, MI, made primarily lower-end designs for dictation machines, card readers, and so forth, and didn't do any of the higher-end business that Nortronics had.

In 1957, Harold Roy Frank started Applied Magnetics Corp. in the spare bedroom of his house in Goleta, CA. He had worked in geophysical exploration and realized that the state of instrumentation recording at the time was very poor and that instrumentation heads needed to be better. So he started out as an instrumentation head manufacturer, and he took the opposite approach to companies such as Nortronics and didn't have any standard catalog but made everything to specifications.

Applied became a huge incubator of magnetic technology. John R. French talks about his father having been a good friend of Mr. Frank. He was offered a job there but didn't want to move out into the middle of nowhere, which Goleta was at the time. Of course, today it's practically part of LA.

Applied started out mostly making heads for the burgeoning computer industry, but later moved into the audio head business, possibly at the behest of MCI. The company made all the special high-density audio heads for John Stephens' recorders. Applied grew to the point where it was a $300 million company in the 1980s, the primary manufacturer of disk and tape heads for the computer industry.

The J.B. Rea Co. of Santa Monica, CA, made computer peripherals and some small general-purpose computers based on magnetic drums in the early 1950s, so it was not a huge step for it to get into head manufacturing. John R. French remembers his father working there, and then working for a spinoff of that company in Glendale, CA, called Western Magnetics, which also made heads. J.B. Rea was started after World War II by mechanical computation people from Northrop.

John Pretto had worked at Brush-Clevite and learned to make heads there, and after a few years in the speaker business in 1959 he started International Electro-Magnetics (IEM), in a garage in Pennsylvania along with magnetic engineer Oscar Dahms. They soon moved to Illinois. They designed some early digital and multitrack heads in the 1960s and made some of the very first 8-track heads for numerous big manufacturers, as well as doing OEM 1/4" heads for Scully.

Nortronics is a widely known and respected head manufacturer. After Joe Dundovic, one of the original founders of Nortronics who served as chief engineer throughout the life of the company, retired in 2008, JRF was selected to purchase his Nortronics head inventory.

The 1960s

John J. French (the father) founded GEMCO (Gem Electro-Magentics) in Fairfield, NJ, in 1960 to build heads for airborne recorders and digital applications. Digital was always John J. French's first love and interest. He was frequently quoted as saying "Audio, Video the real money is in Digital."

His son, John R. French, was more interested in audio tape heads. After a couple of years spent working for his father at GEMCO, in 1964 he started Grandy Magnetics to build heads for the audio industry (mainly tape duplicating 8-track and cassette). His first visit to an Audio Engineering Society show in 1974 convinced him that there was a need for high-end studio heads. He wanted to

Recording head (a) and playback head (b) for the iconic Ampex ATR102

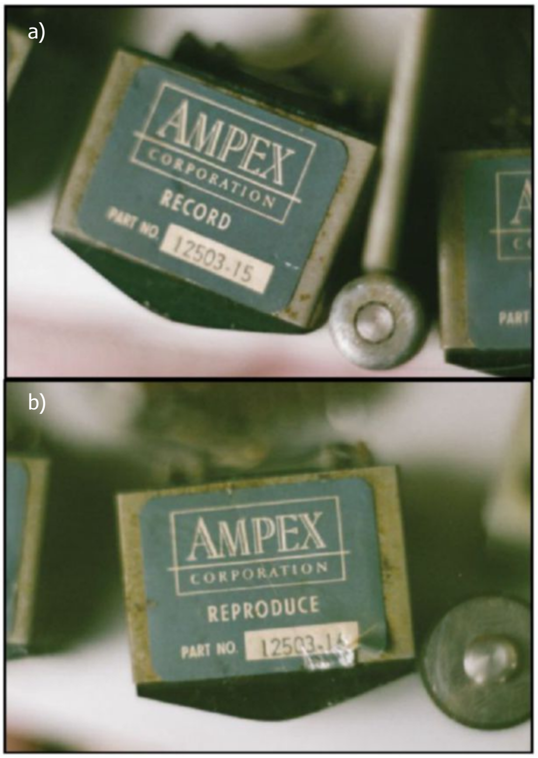

get into that business and tried with Grandy. He continued working there until leaving to form JRF Magnetics in 1980. Although Grandy was mostly known for high-speed duplication systems it also made some heads for the likes of Stephens.

When he left Grandy, his brother Jeff French went to work for GEMCO and continued there until it closed. In 1985, GEMCO and Grandy moved to Vermontville, MI, to be located with Michigan Magnetics under one roof. Jeff started Gemteck after GEMCO discontinued operation and offered ferrite tipping on heads up to 24-tk. This is a process where worn heads would have new tips of ferrite added to the pole pieces to refurbish them. He left that business to build houses.

The J.B. Rea Co. entered the 1960s with a small vacuum tube computer, the Readix II, but it was unable to transition to the solid-state world and the remains of the company was sold to Idaho Maryland Mines Corp., a holding company. Idaho Maryland's Magnetics division continued making digital heads as the "IMI Division of Idaho Maryland" but disappeared after 1963.

Ferrite heads finally started to become practical. Eugene Saksegawa (Gene Saki) had been a scout for General Joseph Stillwell during World War II, then he went to work after the war in broadcast, winding up at the USC television station in 1953. Bing Crosby Enterprises was trying to make a video recorder at the time, and the company hired Gene Saki to improve its head designs. He was the first person to come up with a head that could exceed 1MHz bandwidth with a very uniform field.

When Ampex came out with its transverse video recorder in 1955, it became clear that the BCE machine was doomed and the next year BCE was sold to 3M to become the Mincom instrumentation recorder division. After nine years designing heads at 3M (where ferrite heads had been invented), Gene joined a number of people leaving the company to form Winston Research Corp. to make instrumentation recorders, where he was the manager of magnetic head engineering. 3M sued Winston for making machines too similar to 3Ms in 1965 and the end fallout was not good for Winston.

Out on his own, Gene hired Trevor Boyer, a salesman from Applied Magnetics, and opened Saki Magnetics in Calabasas, CA, in 1969. The company was sold to TDK in 1972 and Gene remained in charge of what was by far the smallest division of TDK. He said that the entire revenue of Saki was not enough to pay for TDK's corporate coffee budget, but the company was very proud of having a division that made such good products and of the fact that it never lost money any quarter it was in business. When TDK dissolved the division in the early 2000s, Saki was still profitable.

Ferrite was actually very popular in Japan, with Sony and Akai investing heavily in design of ferrite heads. Ampex even worked on ferrite heads although its adhesive-based designs were never as reliable as Saki's. Ferroxcube, which had long made ferrites for inductors, opened up its Magnetic Recording Division in 1967 to make ferrite heads for high-speed cassette duplicators, a business out of which it never expanded.

Michigan Magnetics was acquired by VSI Corp. in 1962. Applied Magnetics continued expanding, starting a manufacturing operation in Belgium (still in operation as AM Belgium) and one in South Korea. It was one of the very first companies to start any manufacturing in Korea after the war. It got heavily interested in disk drive heads as well as digital and analog tape heads.

About the Author

Scott Dorsey has a degree in electrical engineering, during the pursuit of which he worked in the broadcast and recording industries. After several years working at a major studio, he took a job with a defense contractor. This left him time to do live concert recording for acoustical music and to design and build audio devices for personal use and on contract to several audio manufacturers and importers. Scott is a regular contributor to several audio magazines. He has been publishing equipment reviews and DIY projects since the mid-1980s. He is probably best known in the general audio community for his retrofit electronics designs in inexpensive Oktava, AKG, and Feilo microphones.

Electro-Sound started up, founded by a number of people who left Ampex in Redwood City, CA, and built a plant nearby in Sunnyvale to manufacture tape duplication systems that looked an awful lot like Ampex. Electro-Sound also made an Ampex 440 clone and heads, which appeared to be copies of the Ampex designs.

And in 1969, James Lemke took his newly-minted doctorate from UC Santa Barbara and opened Spin Physics in San Diego, CA, to make high-performance magnetic heads for video and precision instrumentation applications. The company was sold to Kodak in 1972 but continued operating as an independent operation. In the audio industry it is most famous for having made the high-density heads for the early 3M digital recorders, but at one point in the 1980s it had 650 employees and claimed 50% of the television hours recorded were made on its heads.

During the 1960s, Western Magnetics remained in business, and was bought out by GJM. The company made the 10-track heads for DECtape and LINCtape drives, as well as plenty of other digital systems, with William P. Dingdale becoming president in 1969.

Another interesting company was Consolidated Electrodynamics, which started in 1937 making chemical instrumentation. Bell and Howell bought it out in 1960 and considerably expanded its instruments line, adding a number of instrumentation recorders and so of course the company started building heads in-house for those machines.

The company also sold heads to other manufacturers, taking out a big ad in 1969 *Datamation* issues advertising long-life heads. In 1972, the company bought out Astro-Science's instrumentation recorder division, but Astro-Science had been buying heads from them already.

Keith O. Johnson came up with the idea of having a focused gap head, running at a very high frequency with a mechanical structure that gave a very narrow physical field. He initially intended this for audio recorders. His first prototype was made by hand, but then his next batch was built for him by Lipps. He and Paul Gregg started Gauss Electrophysics to make machines using this technology in 1965.

One obvious application for it was in high-speed duplicators, and Gauss very quickly became a vendor of duplication machines to Capitol. In 1968, MCA bought the company because it had an optical memory technology that MCA was interested in, but MCA kept running the duplicator and recorder business as well until it was sold to Cetek in 1972.

Refurbishing existing heads is the recommended approach. Worn (but relappable) heads can be restored to factory specs for a fraction of the cost.

MCA started out making metal heads but by the time Cetek took over, it was making ferrite heads intended for use a very high frequencies. Cetek sold some to outside vendors but most were in use in its own tape recorders and duplicators.

Bud Taber, a train engineer for the Western Pacific Railroad, started working on Ampex recorders as a side job in 1959, mostly because of his skill in rebuilding precision motors. He retired from the railroad in 1967 and went full time as Taber Manufacturing and Engineering. He brought in people such as Clyde McKinney and Greg Orton and started making copies of existing Ampex head designs, later branching out into replacement heads for Scully and 3M machines.

The 1970s

In the 1970s, track counts increased higher and higher and there was an increasing demand for higher density disk drive heads. New magnetoresistive technology for disk heads became available on the high end of the computer market and on the low end of the market, cheap commodity floppy drive heads became an enormous seller by the end of the decade.

In the late 1960s, Gordon Chase was working as a metallurgist developing tape heads for IBM. He talked with Clay Whetstone, who had been developing multifilamentary superconductors at the time under the name Electro-Magnetic Components (EMC). They came up with a long-wearing magnetic material, which they patented in 1973. In order to develop this, the company obtained a five-year option to purchase Western Magnetics in Glendale, CA, but this partnership did not work. So Mr. Whetstone took the magnetic material development and moved it to his superconductor operation in San Diego, CA, in 1978.

Applied Magnetics continued spawning more companies. Dave Sutton and several other engineers and managers from Applied Magnetics left in 1970 to form a competing head company, Information Magnetics. Information Magnetics primarily made disk drive heads and lasted until the mid-1980s.

In 1971, Michigan Magnetics was bought by Computer Electronics Corp., which only owned it

The Studer A80 is still one the most popular workhorses and continues to be in high demand. Fortunately there are many options for erase, record, and playback heads.

for a short time before selling it to John J. French around 1975 or so. When Mr. French bought the company, he moved it to the Michigan area and the other family businesses, such as Gimtek, followed with him.

In 1974, James Michael Perotte left his job at Information Magnetics in Goleta. He started a company called Magretech in 1976, which was set up to make hard disk heads for DEC, CDC, and Pertec, which continued operating until 1995.

In 1976, Vikron was started by a privately owned company, Northland Aluminum Products, to manufacture tape heads for low-end digital and audio applications.

Applied Magnetics started a number of manufacturing plants abroad in the 1960s and 1970s, but in 1974 the plant the company had built to make IBM disk heads in Portugal was swept up in the Carnation Revolution. Many foreign-owned factories were nationalized but according to Paul Frank, the disk head plant just disappeared. A few years later it turned up again, having been sold to the Soviets and moved to Bulgaria. And in 1978, the company lost Bob Cross, a design engineer, who left Goleta to open up Pacific Magnetics. Pacific had a manufacturing plant for inexpensive commodity heads down in Tijuana, Mexico, with offices in Chula Vista just across the border.

As helical scan recording for video systems became more and more important, CMC started a plant in 1972 San Jose, CA, to manufacture and repair helical scan head assemblies.

According to Dale Manquen, an engineer named Ed Drexler left his job at Nortronics and started a head manufacturing operation at Coilcraft in Illinois. Nobody at Coilcraft today has any record of this ever having existed or the company ever having made heads, but Dale said they made some good 1/4" 4-track heads.

The 1970s was also the era when aftermarket head resurfacing became a commercial business. In 1975, Daryl and Dottie Sprague started Sprague Magnetics, a company in Sylmar, CA, that did no manufacturing but only relapping and rebuilding of heads. Darryl had worked for Applied Magnetics and knew the business, and their son John Austin took over at some point.

Bud Taber was having health problems by the middle of the decade, and he sold Taber Manufacturing and Engineering to Veldon Leverich, who was retiring as the engineer of KGO in 1978. The company really lasted only a short few years under Mr. Leverich's operation.

The 1980s

Ed Reichenbach, who had worked at Altec and then later Gauss, split off to form his own transformer manufacturing company, Reichenbach Engineering in the early 1970s. In 1979 he branched out into manufacturing magnetic heads with the Cinema Magnetics Division. The company primarily made heads for dubbers but also for various audio recorders, however as time went by the transformer business took over the company.

I talked with Tom Reichenbach, Ed's son in the early 1990s and wish I had spent more time pumping him for information because I don't know where anyone at Cinema Magnetics learned head design. It is important not to confuse Cinema Magnetics with Cinema Engineering, which is one of the other magnetics companies that had become part of Art Davis' Electrodyne.

Clay Whetstone at EMC actually started producing magnetic heads in the early 1980s, using his patented long-life materials. The company sold a variety of heads for 35mm dubbers, for 24-track recorders, and for Otari cassette duplicators where the long-life material was a great win. His company originated the idea of heads with replaceable tips.

At this time, John R. French left GEM Electromagnetics and was working for his father. He had gotten a contract to build 24-track heads for MCI and also a contract for 3M. His father didn't want to make audio heads or deal with the odd demands of audio users, so John broke off and started his own company under the name JRF Magnetics. His father continued selling to digital and instrumentation customers. With difficulties getting product from his father, he started reselling Saki and Applied Magnetics heads for items that he wasn't making in-house.

Bell and Howell sold all its tape operations of Kodak, where it became the Datatape division. For a few months during the transition it was owned by Transamerica Delaval, a holding company. The company primarily made production equipment while Spin Physics, also a division of Kodak, was more engaged with custom products.

Electro-Sound shut down due to the increased competition for tape duplication systems. And around 1980, John Pretto retired and his son Tony began running the business.

Magnetoresistive disk heads were getting smaller and smaller and the fabrication process was changing on a daily basis. Plenty of small companies opened up trying to capitalize on the demand for disk drive heads. For example, Su Hwang and S.C. Chang got investment capital and started Magnex in San Jose, CA, to make thin film heads in 1982. By 1987, the company was shut down and the EPA was looking over the facility. In 1990, I talked to someone there who explained "We don't make heads any more, we changed fields."

On the other hand, in 1981, Gregorio Reyes and a co-worker left Memorex to start the Read-rite Corp. in Milpitas, CA, making inductive thin-film heads as an OEM for disk drive manufacturers.

The 1990s

Greg Orton, who had worked for Taber and then for a decade for Ampex and was head of the head department there, left Ampex as it was dropping out of the audio market. In 1990, he started a company called International Magnetics. Soon changing the name to Flux Magnetics, it became the primary manufacturer of high-quality studio recorder heads and remains that today.

In 1990, Applied Magnetics was making digital heads for DEC and analog heads for MCI and a wide variety of disk heads but had just expanded beyond the size of the market. Most of its production was in magnetoresistive disk heads, which was a technology that was constantly changing. Applied Magnetics was not able to keep up and went bankrupt. However, the Belgian plant split off and became an independent operation under the name AMC Belgium. It is still around, and still making excellent heads.

In 1995, Gary Roberts and Shyam Das from Applied Magnetics left to start a thin-film head company in the Northern California Bay area called DAS Devices. It was so successful that Applied Magnetics bought it out and moved the company back to Goleta.

Read-rite expanded into magnetoresistive heads by merging with Cybernex, a thin-film head producer and by the end of the decade its plant was making giant magnetoresistive heads for many hard drive manufacturers as well as moving into making high-density tape heads for the Quantum DLT drives. By 1995, it was a billion-dollar operation but sales were falling rapidly due to massive competition in the hard drive industry.

Nortronics was still around in the 1990s, and Joe Dundovic, who was one of the company's original employees, was still working there. In 1995, DRS wanted the business as things in the industry were consolidating and bought Nortronics out, but DRS only wanted the digital and instrumentation markets. Consequently, Joe Dundovic wound up with the company inventory of audio heads after the takeover, and he sold them out of his house for a decade before retiring a second time (and selling the inventory to JRF Magnetics). Joe also got into a partnership with Brush Industries to sell its remaining audio head inventory as well.

A year later, in 1996 DRS bought out Vikron and its facility in St. Croix Falls, WI, moving what was left of the Nortronics operation to St. Croix Falls. DRS then purchased MEC and moved MEC's plant in Plymouth, MN.

Then in 1997, DRS bought out most of Magnetic Heads Co., Ltd. in Bulgaria (which only 35 years earlier had been the Applied Magnetics plant in Portugal).

All of these were combined to make the DRS Ahead division. DRS also bought CMC's plant in San Jose, which continued operating as the DRS Ahead CMC Technology Division.

With Bob Cross's death and with the general demise in demand for commodity heads, Pacific Magnetics moved into contract assembly work and custom magnetics for high-frequency electronics.

Kodak shut Spin Physics down, after years of bad management from above and after having invested huge amounts of money into production of floppy media soon before they became obsolete.

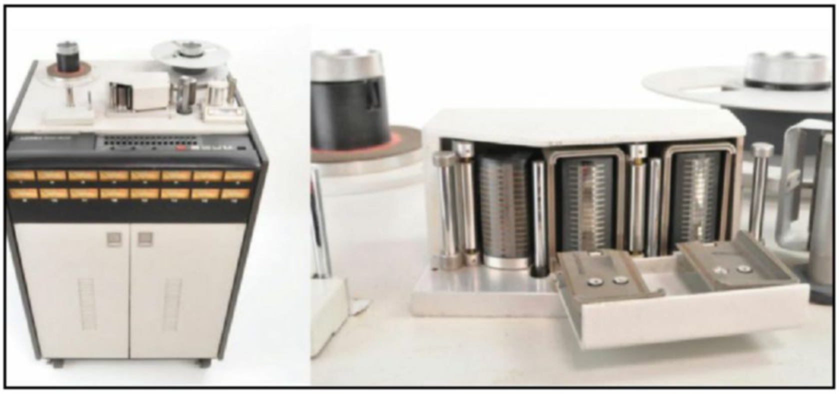

The Ampex MM-1200 16-track 2" tape machine introduced in 1976 was used in countless hit records in the 1970s and 1980s. There were 24-track (2") and 8-track models (1" tape) also manufactured.

The 2000s

In 2000, DRS Ahead sold off its tape head division, which it had cobbled together from parts of Nortronics, Vikron, and Mag-Head Engineering and so forth. Most of it was sold to Michigan Magnetics, which promptly collapsed, shutting down the St. Croix Falls plant in 2001, followed by the Michigan plant in 2002.

Its Bulgarian plant, Magnetic Head Technologies in Razlog, was still open and making products in 2006. In 2016, it reported half million dollars in profit but there is no sign of the company having actually made any products that late. The factory appears to be abandoned today.

The helical scan head plant in San Jose that DRS had purchased was split off as AheadTek, and it remains in business still rebuilding head drums for older professional video systems.

Ferroxcube shut down its line making cassette duplicator heads around 2000.

As digital recording continued to make inroads, EMC concentrated on making heads for the global ATM and point-of-sale maintenance markets, which remained very strong.

Read-rite went bankrupt in 2002, with its assets being purchased by hard drive manufacturer Western Digital, which continued operating independently for a short while but not long.

Sprague Magnetics continued operating, relapping, and rebuilding precision heads for LTO and DLT computer tape drives until it was purchased in 2014 by the NCE Computer Group in El Cajon, CA, which does drive repair work.

In the early 2000s, IEM more or less discontinued all its mass-produced stuff. IEM moved out of its big facility in 2011. However, IEM still makes custom heads and magnetic pickups, but all onesie-twosie work.

Conclusion

I've tried to lay out what I was able to determine about the head industry here and discuss how the head industry has changed the audio and computer world and how it has changed with them.

There are a huge number of companies that I know existed, mostly small job shops that have left little record behind. Magnesonics out on Long Island, NY; Mag-Head Engineering (which split off from Nortronics); Magtech America in Torrance, CA; Northern Magnetics in Minneapolis, MN; Norton Associates; Sabor and Pierce Magnetics in California; Teccon; US Magnetics; Vintage Associates; World Magnetics; IS&CS Engineering; and probably many more. If anyone knows anything about these folks, it would be very helpful if you could contact me in care of this magazine at editorial@audioxpress.com.

I talked a bit about some foreign manufacturers, but there were so many others, starting with Panasonic, which built the heads for a huge number of different Japanese manufacturers, and several of small shops in the UK, such as Phi Magnetronics. But that's a whole other set of stories for another time.

Author Acknowledgements: Many thanks to John French whom I spent several hours on the phone with collecting his reminiscences of the head business. Also thanks to the late Dale Manquen at MANCO, George LaForgia at Ampex, Clay Whetstone at EMC, as well as Tom Reichenbach and David Geron at Cinemag, and Bill Taber who is now with Ampex. Bill Ruck and Tony Pretto also supplied useful information.

References

[1] R., "Mixed Ferrites for Recording Heads," *Electronics* 24: 124-125, April 1951

Resources

C. Alden, *3M—Magnetic Media Maker: A History of the First Four Decades 1944-1985*, self-published 2012.

M. Barron, Michael Barron's Brief History of Tape," www.musicandliterature.org/features/2015/3/28/a-brief-history-of-tape

E. Daniel, C. Dennis Mee, and M. H. Clark, *Magnetic Recording, the First 100 Years*, Wiley-IEEE Press, 1998, ISBN 0-7803-4709-9

JRF Magnetic Sciences, www.jrfmagnetics.com

Museum of Magnetic Sound Recording, www.museumofmagneticsoundrecording.org

"Oral History of Paul D. Frank," Interviewed by Ron Dennison, Recorded October 25, 2017, Computer History Museum, CHM reference X8372.2018," https://archive.computerhistory.org/resources/access/text/2018/07/102740231-05-01-acc.pdf

S. Schoenherr, "The History of Magnetic Recording," presented at Institute of Electrical and Electronics Engineers (IEEE) Magnetics Society Seminar, UCSD, November 5, 2002 www.aes-media.org/historical/html/recording.technology.history/magnetic4.html

TDK Corp., www.tdk.com/en/tech-mag/ferrite02/006

TRIAD

ELECTRONIC TRANSFORMERS

- Original Equipment
- Replacement
- Amateur

TRIAD MAGNETICS, INCORPORATED

460 HARLEY KNOX BLVD., PERRIS, CA 92571

www.TriadMagnetics.com