

audioxpress

ADVANCING THE EVOLUTION OF AUDIO TECHNOLOGY

Market Update

MEMS Microphones and Voice Pickup Sensors What Do You Want to Hear?

By Mike Klasco and J. Martins

R&D Stories

Unified Coverage Map A Radically Better Technology for Today's Hybrid Environments

By Erica Yorga (Nureva)

On the Subject of Standards Filling in the Gaps Left by AES67

By Ekin Binal (Crestron)

Fresh From the Bench

Arylic Audio Up2Stream Amp V4 Wireless Multiroom Easy to DIY

By Stuart Yaniger

Development and Design

Modern PC Board Design Practices — Part 1

By Ethan Winer

Sound Control


Suitability of Venues for Various Events

By Richard Honeycutt

You Can DIY!

High-Performance Tube Hybrid Preamp

By James Skov



ANATOMY OF A 300B

PRESENTED BY *Western Electric*

New, patent-pending graphene-coated anode improves thermal emissivity, electron affinity, and eliminates secondary emission

Molybdenum grid creates negative coefficient of resistance

Ultra-high vacuum finishing at 2×10^{-7} Torr, made possible with state-of-the-art turbomolecular pumping systems

Proprietary mica coating increases dielectric strength, optimizes higher voltage operation

Filament core's proprietary oxide coating allows for 40,000 hour lifetime and superb impulse response

Quality guaranteed with a five-year extended warranty



There's more to it than cathode, anode, and grid. Each Western Electric 300B is assembled at our new facility by the hands of our highly trained staff using the highest quality materials and modern technology—all while honoring the original 1938 specifications and tooling.

What's the result? What's in it for the listener? Unmatched manufacturing quality and the most powerful, full-bodied sound there is. When you're through compromising, order a matched pair online or from Authorized Dealers worldwide.

AVAILABLE NOW

westernelectric.com/300b



MANUFACTURING / DESIGNING TRANSFORMERS & ENCLOSURES SINCE 1917





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, suggestions, and contributed articles to:
audioXpress, Editorial Department
P.O. Box 417
Chase City, VA 23924, US
E-mail: editor@audioxpress.com

Legal Notice:
Each design published in audioXpress is the intellectual property of its author and is offered to readers for their personal use only.
Any commercial use of such ideas or designs without prior written permission is an infringement of the copyright protection of the work of each author.



Where Do We Want to Go Next?

Among the biggest technology stories of 2022, it’s impossible not to reference OpenAI. This artificial intelligence (AI) research company from San Francisco, CA, founded in December 2015, has focused on artificial general intelligence (AGI) applied to augmenting and outperform human skills. OpenAI cleverly made its efforts accessible to the general public as a way to simultaneously promote AI and improve it with actual human interaction—which in turn helps to reveal boundaries and usage restriction needs.

Among OpenAI’s most interesting efforts is a trained model for human interaction through text input called ChatGPT. Its proven uses span from helping us to debug code to answering questions, to writing essays about any topic and in any given suggested style. While in research preview, usage of ChatGPT remains free (chat.openai.com).

Another notable technology trend in 2022, from an early fascination with voice recognition and virtual assistants in always-connected smart speakers, the industry is now transitioning to AI/ML on-device (edge AI)—particularly for hearables. As I intended to write about the possibilities of AI/ML in audio applications, I asked ChatGPT about this topic

Very quickly, ChatGPT identified the uses of AI/ML in speech recognition for transcription, in voice assistants and text-to-speech systems, and also for music classification (metadata indexing and recommendation engines), and music generation—curiously one of the earlier uses to create copyright-free music beds and even full original soundtracks in homemade videos. An interesting and somewhat surprising suggestion was the use of AI/ML for audio compression, dynamically preserving the quality of the original audio signal for storing or transmission. Which could mean an adaptive audio codec, but also choosing codecs according to conditions and requirements. I assume ChatGPT was referencing actual applications.

Finally, and as expected, ChatGPT mentions audio enhancement applications. “AI can be used to improve the quality of audio signals, such as by removing noise or by increasing the volume of quiet sounds.”

In this issue of *audioXpress*, we address the fast-evolving world of MEMS microphones and voice pickup sensors, increasingly combined in true wireless earbuds, headphones, headsets, smartphones, and other personal systems. And one of the current trends is to add on-device AI for audio enhancement. To run machine learning tasks directly on edge devices, we can take advantage of the latest ARM Cortex-M and RISC-V architectures, as well as deep neural network processors and advanced processors specialized for both machine learning and audio processing such as those from Syntiant, BrainChip, FemtoSense, Synaptics, memBrain, and many others.

Much like OpenAI is showing us unheard of possibilities, in 2023 we can expect to see audio applications and performance levels we have never imagined.

AI/ML shows infinite possibilities for audio product design when combined with modeling and simulation, which will enable us to reach the most efficient audio amplifiers with the lowest distortion, or near-perfect speaker diaphragms for a given set of parameters.

AI/ML tasks for audio DSP are actually very limited in scope compared to other tasks that are being required from current available engines. Models for audio applications are light enough to be loaded into miniature, ultra-low power designs. Of course, we also want to be able to use full-featured voice assistants, which require connectivity and are the reason why we will continue to connect wearable devices to smartphones (the most powerful connected computers we ever carried in our pockets).

The topic of connectivity is another major one for 2023. One of the unpredictable effects of the pandemic was that network operators had an unexpected reprieve that allowed them to accelerate the transition and coverage deployment of 5G. While we were all confined at home, the priority was home fixed broadband and Wi-Fi—an area that saw important progress with 6/6E and will soon evolve to Wi-Fi 7. As we traveled out of the home and upgraded to the latest smartphone models, we are now discovering the possibilities of being connected to 5G networks. A completely new world for audio streaming that is the new normal in 2023, while the industry is already working toward a 6G future.

J. Martins
Editor-in-Chief

The Team

Publisher:	KC Prescott	Associate Editor:	Shannon Becker
Controller:	Chuck Fellows	Graphics:	Grace Chen, Tori Zienka
Editor-in-Chief:	João Martins	Advertising Coordinator:	Heather Childrey
Technical Editor:	Jan Didden		

Regular Contributors: Bruce Brown, Bill Christie, Bob Cordell, Joseph D’Appolito, Dimitri Danyuk, Vance Dickason, Jan Didden, Scott Dorsey, Gary Galo, Chuck Hansen, Richard Honeycutt, Charlie Hughes, Mike Klasco, David Logvin, Ward Maas, Oliver Masciarotte, Nelson Pass, Christopher Paul, Bill Reeve, Michael Steffes, Nora Wong, Stuart Yaniger, Ethan Winer

OUR NETWORK



SUPPORTING COMPANIES

Cine Gear Expo	53	Nureva	17
COMSOL, Inc.	9	Parts Express International, Inc.	68
DSP Concepts	11	Redco Audio	65
Hammond Manufacturing, Ltd.	3	SB Acoustics	31
HEAD acoustics	21	Schurter, Inc.	33
KAB Electro-Acoustics	47	Triad Magnetics	67
Menlo Scientific, Ltd.	25, 57	Western Electric	2

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.

The Audio Voice

Sign up now for our **FREE** weekly e-newsletter



Weekly Newsletter for the Audio Industry, Professional Designers and Audio Enthusiasts
Product Design | Audio Electronics | Acoustics | DIY | Innovations in Audio





Part of the Loudspeaker Industry?

Designed to help you locate the resources that are essential to loudspeaker manufacturing operations, the *Loudspeaker Industry Sourcebook* keeps you informed and connected with the companies making it happen.

Don't miss out! Sign up for your FREE listing!

www.loudspeakerindustrysourcebook.com



● Features

8 **Arylic Audio Up2Stream Amp V4**

By Stuart Yaniger

The Up2Stream Amp stereo streaming amplifier board from Arylic is a complete, affordable solution that invites integration by DIYers and integrators for multiroom and wireless audio setups. *audioXpress* explores how it measures up and shares the experience.

34 **On the Subject of Standards Filling in the Gaps Left by AES67**

By Ekin Binal

Is the AES67 interoperability standard a viable alternative as the core solution for audio-over-IP distribution? Crestron proved that it is with its DM NAX technology filling the existing gaps of network discovery and routing.

42 **Modern PC Board Design Practices Part 1 — An Overview of the Process**

By Ethan Winer

In this two-part series, Ethan Winer examines current practices and methods for designing four-layer printed circuit boards using surface-mounted components. Part 1 presents an overview of the process, while Part 2 will explain specific strategies. Many related issues will be described along the way.

48 **Power Transformer Parameters, Selection, and Testing Part 4 — Parameters and Their Losses**



By Chuck Hansen

The initial group of articles focused on the history of transformer cores, their construction methods, their materials, and various testing methods. Part 4 of this series discusses the parameters of transformers and their losses.

54 **The Pre-5 A High-Performance Tube Hybrid Preamp**

By James Skov

In a never-ending search to improve the performance of tube-based audio circuits, this project, which the author calls the Pre-5, proposes pushing the envelope to achieve even lower distortion from an already great-sounding concept.



● Microphones and Voice Pickup Sensors

MARKET UPDATE

16 What Do You Want to Hear?

By Mike Klasco and J. Martins

The latest developments and new technologies in MEMS microphones promise to rival the performance of electro-dynamic and electret condenser microphones. Now, voice pickup sensors, DSP integration, and edge AI promise to take voice capture to uncharted territory.

28 Unified Coverage Map

A Radically Better Technology for Today's Hybrid Environments

By Erica Yorga

This article discusses a new patent-pending feature of Nureva's next-generation technology for audio presentations and conferencing that provides a new perspective on how microphone coverage works in large meeting rooms.

● Departments

4 From the Editor's Desk

5 Supporting Companies

● Websites

www.audioxpress.com

www.voicecoilmagazine.com

www.cc-webshop.com

www.loudspeakerindustrysourcebook.com

www.circuitcellar.com

www.linuxgizmos.com

● Columns

SOUND CONTROL

38 Suitability of Venues for Various Events

By Richard Honeycutt

Understanding the optimal reverberation time vs. frequency performance of a venue, particularly one used mainly for amplified music, is important for acoustical consultants and architects. Recent research has provided valuable design guidelines for such venues.



@audioxp_editor



audioxpresscommunity



linkedin.com/company/audioxpress



audioxpress

Arylic Audio Up2Stream Amp V4



By **Stuart Yaniger**

Photographs by Ysmay Walsh

The Up2Stream Amp stereo streaming amplifier board from Arylic is a complete, affordable solution that invites integration by DIYers and integrators for multiroom and wireless audio setups. *audioXpress* explores how it measures up and shares the experience.

I really should stop being surprised when electronics get smaller, cheaper, and more versatile. Every time it happens, I think, “How can this possibly get even more commoditized?” and then the next one comes along to show my failure of imagination.

That slightly depressing and highly elating thought came to be when I first saw the Arylic Up2Stream Amp Version 4 (V4). Arylic is a Chinese company that makes an exceptionally broad range of inexpensive audio electronics designed for lifestyle and smart house setups. They also sell a line of loudspeakers meant for in-wall and in-ceiling installations.

Most of the company’s devices are ready-to-use packaged electronics, including amplifiers, preamplifiers, and DACs. They have, however, incorporated a line of unpackaged circuit boards and accessories under the Up2Stream brand aimed at the DIY and custom installer markets. Both its packaged electronics and boards are all oriented

toward streaming in recognition of the preferred mass market sources and the ubiquity of home Wi-Fi.

The Product

One of Arylic’s most versatile and popular boards in the Up2Stream line is the aforementioned V4 (**Photo 1**). Packed into this small (11cm x 8cm) board is an analog input (3.5mm jack), a mini USB for connection to computer sources, a standard USB to accommodate thumb drives, an optical input, Wi-Fi, LAN, and Bluetooth 5.0 (SBS and AAC). For those who like handheld IR remote controls, the V4 has you covered. Oh, and there’s also DSP and a pair of 50W power amplifiers, which can be configured as a single 100W mono amp (**Figure 1**). All of this for under \$100.

The board requires external power, which can be 12V to 24V, although the achievable output power is reduced at lower than 24V. Arylic sells a suitable power brick rated at 24V and 5A to obtain the maximum power the V4 is capable of.

To assist customized integration, the V4 also has accessory sockets for optional expansion boards. These include a “button board,” with manual selector switches for source, volume, and program control; a volume potentiometer board;


Arylic Up2Stream Amp

Multiroom Wireless Streaming Stereo Amplifier Board

Regular price: \$119.00

www.arlyic.com

www.arlyic.com/products/up2stream-amp-amplifier-board



Take the Lead in Acoustics

with COMSOL Multiphysics®

Multiphysics simulation drives acoustics innovation by providing insight into every design aspect that influences product performance. The ability to account for coupled physics phenomena lets you predict, optimize, and virtually test a design under real-world conditions – even before a first prototype is built.

» comsol.com/feature/acoustics-innovation

SPDIF I/O; an extension for the IR detector so that the board can be enclosed but still support remote operation; and (for me) most interestingly, a substitute DAC featuring the ESS 9023 Sabre, seemingly an acknowledgment that the stock DAC used might be a performance choke point. An IR remote is also available as an option.

Other headers on the board allow the integrator to access clocks for multi-unit master/slave

arrangements for IIS connection. Note that data on IIS is sent in a 16-bit format only. There's also status registers that can be used for a display panel.

One odd layout feature worth noting is that, although the inputs and outputs are all helpfully arranged on the same edge of the board, the IR sensor for remote control is on the opposite side. This will complicate packaging unless you spring for the optional IR extension board. For me, it was not an issue because I prefer using my iPhone for control, but if you like dedicated remotes, this is something of which you should be aware. Likewise, the LED used for indicating board status by changing color and blinking patterns is also on the opposite side.

There are two 7" wire antennas attached to the board that are used for Wi-Fi and Bluetooth. They are supplied with double-stick tape at the end for placement inside an enclosure.

Packaging is, befitting the low price, a rather basic polypropylene storage box. Oddly, a screwdriver is supplied with the board, but despite this being a common tool, I won't look a gift horse in the mouth. The board is packed with a paper copy of the manual with print tiny enough that I needed magnifying lenses to read it. You'd do better getting a PDF copy from its website (see Resources) and enlarging it on your computer or phone screen.

Speaking of the website, it is truly one of the most horrible I've used, bristling with constant popups and banners, making reading a challenge. It's also weirdly over-responsive and badly needs a rewrite from a native English speaker. Nonetheless,




Photo 1: The Arylic Up2Stream V4 board is neat, compact, and packs a lot of functionality into a small area.




Figure 1: The V4 has a variety of I/O options, sockets for expansion boards, and a built-in pair of 50W power amplifiers.



DSP
Concepts

Creator of **Audio Weaver:**
the toolbox answer to the black box problem

Automotive Voice

proven to meet the rigorous standards of the board room
now available for automotive

Voice Communications

TalkTogether

The full-duplex solution

Voice UI

TalkTo

The advanced audio-front-end

Two design levels

Essential supporting front-seating

Advanced for the whole cabin

Learn more at <https://w.dspconcepts.com/automotive>



most of the needed content is there, if sometimes challenging to access. There's also a user forum where technical questions can be asked, which are then answered by company representatives. Some confusion on the website comes from the product name, which is variously referred to as Up2Stream AMP 2.0, V4, and Amp Mono.

The App

The interesting feature tying together Arylic's product line is its 4STREAM smartphone/tablet app. 4STREAM allows multiple devices to be controlled—each device can be individually selected, customized for source and volume, and grouped for

synchronizing the source (although each device in a group can have its own volume and EQ settings), enabling multi-room sound even across different devices (**Figure 2**). The 4STREAM app is simple and intuitive and enables the user to access most streaming services as well as control volume, signal source, and streaming service (**Figure 3**). Note that the list of supported services is quite long—the screenshot I took could only capture the end of the list. It should be noted that the V4 works exclusively in 16-bit/44.1kHz mode. It will accept other formats up to 192kHz, but resamples them to its native 16-bit/44.1kHz.

One very nice aspect of how the app functions for streaming is that, instead of the phone doing the streaming and then transmitting to the V4 board via Bluetooth, the V4 board connects to the Wi-Fi network and streams the sources directly. This can, in theory, reduce lag and compression, but that assumes Bluetooth is the weak link in the signal chain, which may not be the case. Still, direct streaming is more elegant, and if a potential compromise can be removed, why not?

The Up2Stream V4 in Use

For my initial tests, I set the board's jumpers to provide stereo, then connected a pair of small, inexpensive two-way passive speakers as an example of a realistic use-case. With the 4STREAM app installed on my iPhone, connecting the app to the board, then connecting the board to Wi-Fi following the instruction sheet was simple and painless. After the struggles I've had with other systems, this was a great relief.

Once connected, I selected my preferred service (in my case, Amazon HD), and let 'er rip. On subsequent power-ups, after a short delay, a female voice would announce, "Connected to Wi-Fi," which is the default connection. From time to time, I also used the analog Line input fed by my RME ADI-2 FSB DAC. If there was any sonic difference between streaming and analog, I certainly couldn't hear it without a detailed direct comparison.

Inexpensive Class-D amplifiers built into active speakers are notorious for being noisy. This was not the case here. I could barely hear any hiss from my speakers with my ears just a few inches away. Set up as a desktop system with just a few feet of distance, noise was inaudible to me.

Not unexpectedly, the sound was generally fine, but I was able to get the amp a bit distressed on loud and complex music (e.g., heavy metal), with a sense of congestion and edge in the loudest sections. But let's be clear—30W of power is not a lot unless you have much more efficient speakers




Figure 2: Although only one device is shown here, the 4STREAM app can control multiple Arylic devices.





Figure 3: The 4STREAM app allows the V4 board to stream services directly from Wi-Fi.

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.



than I do. With folk, Americana, small jazz ensemble, and singer/songwriter music more typical of my playlists, the amp played seemingly effortlessly.

No matter what, the amp section ran cool; the hottest part of the board were the Wi-Fi chips, and even those were barely warm to the touch. Remote functions controlled by the phone worked as expected, and navigation through the app was easy and intuitive.

DSP FTW

The V4 has built-in DSP capabilities, but their ability to be used in the 4STREAM app is limited. To unlock the V4’s DSP capabilities, you need the ACPWorkbench software from Shanghai Mountain View Silicon Technology written for the MVSILICON DSP chip used in the V4. This software is a \$20 option which, for most installations, I would strongly recommend.

Once the software is installed on your computer or pad (note that the license is for a single installation on a single computer), the V4 board is connected via USB or a serial UART, the V4 card is detected, and the current DSP chip configuration is read (**Figure 4**).

The next tab shows all of the available DSP options for the V4 DSP chip, which is a lot. The most basic, and most important, is equalization (**Figure 5**). The software can add in up to 10 EQ points, with available functions that include low and high pass (second order), shelving, peaks, and notches, all with variable Q and depth. With a few exceptions (e.g., high-order high-pass or low-pass filters, and comb filtering), nearly any EQ curve can be synthesized from these building blocks.

The software also has a supplied set of predetermined EQ settings, which may save some development and adjustment time. Although names such as “Bass Booster” or “Treble Reducer” are pretty self-explanatory, I do wonder a bit how they arrived at EQ settings for “Jazz” or “Classical.”

Dynamic range tailoring is also available. This ranges from simple limiting (hard or soft) to compression/expansion with variable depth, attack, and release.

The DSP functions can also be set to be the same for both channels, or different settings for each channel, which is a very useful and flexible option for setups that may not be symmetrical from left to right, or for where the amplifiers are used to biamp a two-way setup. The software also allows programming the V4 for different gains between inputs to normalize any volume differences.

Integration into compact systems (e.g., soundbars or desktops) can be enhanced by use



Figure 4: The ACPWorkbench program allows configuration of the V4’s versatile hardware capabilities.




Figure 5: The DSP can be controlled in real time via the mini-USB port on the V4 board using the ACPWorkbench.




Photo 2: The V4 board took up very little space on my test bench.

of some of the special functions: Stereo Widen and MV3D Plus. These functions, as their names imply, use phase and EQ effects to move sounds away from the actual speaker positions. With both, you can dial in the amount of the effect applied, and the Stereo Widen gives you the option of phase alone or phase plus some EQ shaping.

One of the best features of this software is the ability to be able to change parameters in real time so that their effects can be heard as they're being adjusted. They can also be switched on and off to gauge their effectiveness against a control.

I should note that the porting over from Chinese to English versions is not perfect. There are occasional pop-ups in Chinese script asking for a response.

Fortunately, the meaning is usually obvious from context, and the occurrences are infrequent. Once DSP settings are determined, they may be flashed onto the chip via a menu command. They can also be saved as a file for later editing.




Figure 6: The V4 has low output noise and exceptionally low carrier leak-through.




Figure 8: The V4's 1kHz distortion spectrum shows mostly odd-order harmonics and a lot of spread in the fundamental line.

Measurements

As usual, measurements were taken using an Audio Precision APx525 audio analyzer along with various dummy loads for the power amp (**Photo 2**). The basic measurements showed 25dB of voltage gain from the line input to the output. Output impedance was slightly over 0.6Ω , which is large enough to have a slightly audible effect on frequency response with typical loudspeaker loads.

Once the basics were completed, the first thing I looked at was the output noise spectrum, shown in **Figure 6**. Noise was low, correlating to the lack of hiss in my system. The most remarkable part to me was the carrier noise at 400kHz—at about 6mV, it is much lower than just about every other Class-D amp I've tested. This doesn't likely have any audible consequence, but it is still a testament to some good engineering.

Checking midband power with an 8R load, the amplifier starts running into distress between 30W to 32W, where distortion reaches the 1% level (**Figure 7**). The distortion spectrum for



Figure 7: The V4 in stereo mode is capable of over 30W of power into an 8R load.




Figure 9: The distortion versus frequency at 1W into an 8R load has significant rise in the top two octaves.

1kHz into 8R at 1W (2.83V) is shown in **Figure 8**. It's a bit messy, with odd-order distortion products dominating, and their level not falling off rapidly with increasing order. As well, the broadening of the fundamental suggests that either there's a lot of jitter or that the carrier frequency is not stable.

Distortion versus frequency at 1W/8R is shown in **Figure 9**. It can be seen to rise significantly in the top two octaves, exceeding 1% at 20kHz. The spectrum of a 10kHz signal at 1W is shown in **Figure 10**. Compared to the 1kHz distortion spectrum, there's a similar broadening of the fundamental, but the third harmonic is significantly greater. Will this be audible? This harmonic is at 30kHz, well above audibility, but it still points to excessive nonlinearity.

The time domain response of an impulse is shown in **Figure 11**, indicating the filtering and resampling to 44.1kHz. I'm glad to see that Arylic didn't go with the anachronistic but fashionable so-called "non-oversampling" filters, which suffer from significant artifacts.

Just out of curiosity, I looked at the Widen Stereo effect. **Figure 12** shows how it works in the time domain, with interchannel




Figure 10: The distortion spectrum from a 10kHz sine wave at 1W/8R shows third harmonic to be rather large.




Figure 12: The Stereo Widen effect operates by an interchannel time delay.

delay set at about 25ms. With the frequency response shaping selected, the time delay remained the same, but the frequency response had a midrange dip (**Figure 13**).

Wrap Up

I saw an online review of a similar product from Arylic with similar measurements. The review was highly negative because, when compared to top-quality individual gear, the absolute performance comes up short. That's fair, but I think it misses the point of the product and its likely uses. For a sound system that is meant for exclusive listening—specifically sitting down and listening to music as a primary activity as opposed to having tunes playing while doing other tasks—the performance is slightly compromised, at least to young and sharp ears.

For the more likely applications—multiroom sound as a lifestyle feature, casual listening, and background—the low cost, compact size, low heat, low noise, and versatility of inputs and customization make the V4 a great option for DIY and integrators. 




Figure 11: The impulse response shows well-chosen anti-image filtering.





Figure 13: With Shaping enabled, the Stereo Widen effect causes a dip in the midrange frequency response.



What Do You Want to Hear?

By
Mike Klasco
and **J. Martins**



The latest developments and new technologies in MEMS microphones promise to rival the performance of electro-dynamic and electret condenser microphones. Now, voice pickup sensors, DSP integration, and edge AI promise to take voice capture to uncharted territory.

We have in previous market reports discussed how traditional electret condenser microphones (ECMs), once used in massive quantities in laptops, camcorders, hearing aids, baby monitors, cellular phones, PDAs, and pre-smartphone mobile designs, have been replaced by microelectromechanical systems (MEMS) microphones. Nearly 90% of the more than two billion microphones produced annually in 2010 were still based on the principles of the foil-electret invented in 1961 by Gerhard Sessler and James West at Bell Laboratories. Ten years earlier the industry reached the landmark of one billion microphones annually thanks to ECMs. The cellular phone was the product that created that massive demand.

In 2007, the iPhone was introduced and helped create the largest consumer product category of all times—the smartphone—paving the way for a wave of technological disruption. With low-cost yet high-production yield, low-noise with minimal handling microphonics, and single element directionality, ECMs were still the best option for the first iPhone in 2007. But soon those devices would embrace the emerging category of MEMS microphones due to being smaller in size, lower power, and much more robust. The massive industrial scale that was applied to MEMS microphones transpired to many other consumer audio applications and fostered many new voice capture applications, including playing a vital role in the voice recognition field and smart speakers, by enabling innovative microphone arrays.

Originally researched for applications that required DSP and IC integration, MEMS microphones have quickly surpassed their recognized technical limitations that resulted in large measure of its very small size. When MEMS microphones typically reached

60dB SNR and the threshold of 70dB seemed to be feasible, MEMS microphone production volumes quickly surpassed all other microphone technologies. Which doesn't mean that life was easy for the MEMS microphone pioneers.


The Tough Business of MEMS

The technology itself took more than a decade to mature, and MEMS microphones have remained restricted to working in close proximity to the source, which nevertheless made the devices a natural choice in smartphones and true wireless earbuds.

The massive industrial scales and price constraints required by the mobile industry have continuously forced some of the original innovators and designers of MEMS microphones to exit the segment or negotiate manufacturing rights. Analog Devices sold its MEMS microphones business to InvenSense in 2013, Knowles acquired the MEMS microphone ASIC design business from ams AG (now ams Osram), while other MEMS companies were acquired by larger ones: Wolfson by Cirrus Logic, Akustica by Bosch, InvenSense by TDK, and so forth.

Other leading companies in the space include Infineon, Knowles, Sonion, Goertek, and AAC Technologies, all of which are finding healthy business in the growing true wireless stereo (TWS) and wearables category, gaming, and other consumer product categories, not from the competitive smartphone segment.

As a result of the lower margins in an incredibly competitive market where any single cost reduction option was explored to the limits, all these companies were basically forced to work with the industrial giants that own the ASIC and MEMS element



INTRODUCING

Nureva[®] HDL410 system

A revolution in pro AV performance

Getting pro AV performance in your large spaces doesn't need to be complicated. The HDL410 offers patented technology and powerful features – including our groundbreaking unified coverage map (read about it in this issue). And it's all in a USB plug and play, wall-mounted system.

Discover the next generation of large-room audio – only available from Nureva.


nureva.com/audioexpress




foundries and large OEM manufacturing operations. A smartphone typically requires a minimum of three to four MEMS microphones, making the scale of the market compelling for the larger players. In 2014, the large players leading the mass production of MEMS microphones were already giant-scale MEMS foundries with automatic production lines, including MEMS assemblers buying MEMS microphone elements elsewhere. The MEMS microphones found in most products consist of the actual MEMS mic element, the associated ASIC and the package. Many manufacturers are just packagers assembling these elements together, while others mix and match their own products with elements from other manufacturers or brands.

According to the latest MEMS microphone market report from Yole Intelligence, which analyzes the actual components used in shipping smartphones, Knowles is the main player in MEMS microphones representing 46% of the total components, followed by Goermicro with 29%, and Omron with 12%. Lesser known companies such as MEMSensing, Zilltek, and Novosense are relatively new players in the market, but together already represent more than 25% of the market.

As Yole notes, the interest in the consumer microphone market is massively increasing and the main drivers are wireless earbuds, where Goermicro and Knowles are already leading players. Other companies supplying the emerging space are Omron, MEMSensing, Zilltek, BSE, Novosense, TDK-Invensense, and Hosiden.



According to SAR Insight & Consulting, the market for MEMS and ECM microphones continues to see strong growth, with 10 billion units shipped in 2020. This was driven by two main factors: increased use of voice interfaces and multiple microphones per device. With the emergence of TWS, earbuds and headphones are now the second largest market. (Image Source: "MEMS Microphones Focus," published July 2021, SAR Insight and Consulting).



2017 VocalZoom's concept of optic voice activity sensor in conjunction with acoustic pickup. VocalZoom has stepped back from pursuing consumer applications for their optical sensors to focus on the industrial market.

Much like smartphones, true wireless earbuds are now typically using four MEMS microphones due to the popularity of active noise cancellation (ANC), transparency modes, and hearing augmentation (paving the way to hearables). Those massive scales will naturally benefit other lower volume application markets, which still generate a significant business. According to recent reports, the global MEMS microphone market size is expected to continue to grow, owing to the widespread usage of microphones in automotive electronics, smart speakers, speech recognition systems, home appliances, and even speakerphones and conference soundbars, all of which use between two to four, with soundbars using as much as a dozen or more.

For smartphones, TWS earbuds, latest-generation hearing aids, hearables, MEMS microphones are integrated with other sensors and built-in edge AI processing. This allows designs to overcome the known limitations of MEMS microphones, combined in arrays and allowing the use of beamforming for enhanced directionality and increased immunity to noise. And these increased-value systems are also allowing far-field applications in the home, the office, and inside the car.

Integration with other sensors, including voice pick-up (VPU) bone conduction transducers allow improving the performance of headphones and in-earphones. Another opportunity arises for introducing new technologies in MEMS microphones, including the use of optical sensors.

Clean Voice Pickup in Noisy Environments

An office worker jogging outside on a call, a field reporter on an Electronic News Gathering (ENG) assignment in a storm, or recording dialog in an onsite movie scene, are some of the more challenging tasks for microphones. Back when consumer applications were a shoebox cassette recorder using a lavalier mic worn around the talker's neck, we seemed to have the requirements "aced." Mobile phones certainly upped the game with rainy weather, wind handling, noise, and size challenges. Now voice recognition, voice command, and speech-to-text transcription has put us on another long road to reach the precision required to achieve accuracy.

Today's voice assistant products still have only marginal accuracy when commanded across even in a quiet living room, let alone in a factory setting or while driving home with a window open and listening to the news. MEMS' microphone diaphragms don't like wind noise pressure. Yes, there are some MEMS mics that handle these conditions better than others and proper application of wind screens can achieve more robust operation with less diaphragm saturation and latch-up. Talker separation and ambient noise pickup can be mitigated by mic arrays, but all these techniques have yet to tame most consumer frustrations to voice command. Ultra-low noise floor mics have not arrived as scheduled, with 70dB+ SNR MEMS microphone devices hitting a glass ceiling for years now—even with some mic vendors using two mic elements within the MEMS package to wring out the last couple of decibels.

Yet there is another dimension for achieving cleaner voice pickup than just airborne techniques, and this is bone conduction speech pickup. Instead of using a mic, or in conjunction with a mic, the time has come for product designers to consider using bone conduction pickups as industry leaders such as Apple already have.

Bone Conduction?

Vocal tract vibration can be picked up through bone, subcutaneous (under the skin) cartilage or the skin surface using a bone conduction pickup. There is also the combo approach with the vibration pickup sensor used as a speech activity detector in conjunction with a conventional mic.

Bone conduction is not new and has been in use for military, first responders, and industrial communications. More specifically, vibration pickups have been used on the throat (or even tooth contact in helmets for fighter pilots), and also as an electro-larynx to aid voice production by those who have lost their voice box due to cancer.

For aesthetics, an ideal location for the bone conduction pickup would be on the pinnae (the outer ear) although this is not the most optimum voice pickup location. The cheek and the back of neck are good spots for speech pickup—but we have to balance our developments keeping in mind that we are designing for a headworn device that users will consider aesthetically acceptable.

Bone conduction within the talker’s skull and resonances in their sinus cavities are one reason why a person’s voice sounds different to them when it is recorded and played back. Because the skull conducts lower frequencies better than air, people perceive their own voices to be lower and fuller than others do. Sound quality and sound levels, and more specifically dynamic range, are all influenced by the differences in listening conventionally vs. bone conduction; both due to psychoacoustic and physiological reasons.

The first mainstream consumer bone conduction voice pickup application was the Jawbone, developed 20 years ago by Aliph (the original corporate name for Jawbone). This technology was developed from Aliph’s contracts with the Defense Advanced Research Projects Agency (DARPA), the US military’s research arm, to explore ways for combat soldiers to communicate with each other in difficult conditions.

In 2007 Aliph revealed its wireless Jawbone headset at the Consumer Electronics Show (CES), just in time for the state of California to pass legislation to ban the use of handheld phones for drivers. The Jawbone offered a single earpiece Bluetooth with integrated bone mic pickup used in conjunction with a conventional mic and conventional earpiece.

More specifically, for contact voice pickup, an ECM mic was compliantly in contact with the skull and integrated into the on-the-ear Bluetooth device with a second mic used conventionally. Voice range sound energy that was coherent from both the mic and vibration sensors being passed onward to more signal processing. The addition of this “voice activity detector” has proven quite effective.

The voice intensity, where the vibration pickup is typically positioned, is higher than the background noise. If the structural energy pickup is negligible at the voice activity detector, then it is assumed the wearer is not talking and the airborne sound mic can be muted. The use of both structure-borne pickup and airborne speech used with appropriate signal processing promises higher noise rejection along with the wider bandwidth of airborne sound pickup microphones alone.

TWS earbuds are now the best app for bone pickup sensors because the in-ear fitting provides a viable coupling to the skull vibration from the vocal tract. Bone conduction sensors are becoming common in the higher quality TWS devices (e.g., Apple AirPods Pro, Jabra Elite Pro, Huawei FreeBuds Pro, and Edifier DreamPods).

Recent generations of AirPods are based on the combination of single and double microphones, bone conduction pickup, and sophisticated algorithms. While AirPods and the Huawei FreeBuds-3 both use sensors for voice detection, the AirPods’ dual microphone plus bone vibration sensor starts with digital, while the Huawei’s FreeBuds-3 is obtained as an analog signal. Specifically, Huawei’s approach is to use Sonion’s VPU14AA01, a single-axis vibration low power consumption sensor that outputs an analog signal.





Diagram of the different vibration paths of sound to the inner ear by bone conduction with their rating for pickup performance. (Image Source: one of the Jawbone patents)



Vesper OnePlusOne is the combination of one microphone plus one voice accelerometer and both are combined through the software sensor Fusion Library, removing ambient noise. Earlier in 2022, Vesper updated its VA1210 high bandwidth, low power accelerometer, complementing its piezoelectric MEMS microphones.

Apple's AirPods use the STMicroelectronics LIS25BA, a triaxial vibration sensor that outputs a digital signal. In any case there is significant proprietary signal processing. These applications truly require deep learning and neural network signal processing strategies for effective characterization of more complex signals with adaptive capabilities. These deep learning-based algorithms



The Jabra Elite 7 Pro true wireless earbuds improve call performance and battery life with adjustable ANC, featuring Jabra's MultiSensor Voice technology, which combines a bone conduction voice pick-up sensor, four microphones, and proprietary algorithms. The VPU sensor in both earbuds is activated when detecting wind noise and the intelligent algorithm then uses the best combination of the bone conduction sensor and microphones to transmit the best call clarity.




Illustration from Alango's OnlyVoice DSP voice capture technology for TWS earbuds. A vibration sensor captures the user's speech via bone and tissue vibrations that is free of acoustic noise but limited to a frequency of 1kHz. Acoustic echoes in the microphone signals and potential "vibration echoes" in the bone conduction signal are cancelled. Acoustic and bone-conducted speech signals are cleaned up using noise reduction technology.

for noise suppression and speech enhancement must be agile and automatically adapt to various scenarios, resulting in high processing workload, increased latency, and expanded convergence of parameter adjustment.

Bone Conduction Transducer Types

All bone conductor pickups convert mechanical vibrations into electrical voltages, but depending upon their sensing element and output characteristics, these pickups are categorized as accelerometers, velocity pickups, or displacement pickups. Applications include not only sound/vibration pickups, but also shock sensors such as burglar alarms, air bag sensors, hard disc drive (HDD) protectors, and more.

An accelerometer consists of a mass that is supported with respect to a surrounding case by means of a spring compliance and guided to prevent motions other than those along the seismic direction of support. The mass exerts a force on the spring's support, which is directly proportional to the acceleration being measured. This, in turn, is converted into an electrical voltage. Depending on the design approach, you may be looking for a single- or three-axis accelerometers.

A velocity pickup generates a voltage proportional to the relative velocity between two principal elements of the pickup, the two elements such as a coil of wire and a source of magnetic field (electrodynamic).

A displacement pickup generates an output voltage, which is directly proportional to the relative displacement between two elements of the instrument. These pickups are similar in construction and behavior to velocity pickups. The only essential difference is the use of a frequency-weighting network, required to make them direct-reading.

Non-contact sensors acquire speech signals by detecting the vibration of a human body and include optical; infrared rays, light waves, and lasers. Optical pickups are non-contact with a significant advantage of avoiding damping of the signal source (e.g., your cheek) thereby altering the source of the vibration signal.

Most contact pickups limit usable response up to about 1kHz. The limitation can be offset by combing a conventional mic with a bone-conduction pickup to detect speech activity. Optical sensing is non-contact but some loss of consonants due to marginal coupling from the talker's articulators (i.e., lip, teeth, and tongue) will occur. Until now, optical sensors also have been expensive, have complex interfaces, and have been limited to lab, industrial, military (surveillance,) and medical applications.

The laser microphone is in the same optical pickup class and is often used for surveillance devices, which utilize a laser beam to detect sound vibrations (eavesdrop) at distant objects.

Microphones, Sensors, and More

The Vesper VA1210, the second-generation analog piezoelectric smart voice accelerometer, remains a reference for the industry for voice pick-up. Vesper anticipated the need to combine new approaches for voice pickup when it launched the original VA1200. The VA1210 was advertised as a Smart Accelerometer with ZeroPower Sensing to boost the potential of a vibration and

ViBRIDGE

In-ear device testing – faithful to the bone

Bone conduction simulation – the game changer




REAL-LIFE-PROOF.

www.head-acoustics.de   

bone conduction sensor in TWS earbuds. In addition the device could analyze the acoustic environment and power the system only when needed, adding features of Wake-on-Voice (WoV) and hardware voice activity detection (HVAD). Other potential uses included Automute, Tap Detection, and Voice Authentication. Availability and official confirmation about the future of Vesper's products has not been disclosed at the time of writing.

Even though we do not intend to specifically discuss DSP solutions (a topic for a dedicated feature), another mandatory reference in this field is the OnlyVoice technology from Alango Technologies. The company offers an extensive range of processing solutions, including its Extended Voice Communication Package (eVCP) for voice applications, supporting all major DSP and MCU platforms. Among Alango's technologies, OnlyVoice is an advanced voice acquisition solution relying on sensor fusion for devices incorporating acoustic microphones and a vibration/bone conduction sensor.

OnlyVoice improves communication intelligibility in noisy environments by employing an in-ear signal source, which due to its positioning, is isolated from ambient noise. The user's speech is acquired with either an in-ear microphone via the eustachian tube or an in-ear accelerometer via bone conduction. The in-ear signal



With Alango's OnlyVoice solution, the microphone signals are enhanced with adaptive beamforming technology, which is most efficient at high frequencies, and mixed with the bone conduction signal at low frequencies. Intelligent mixing is controlled by the levels and spectra of speech and noise in the acoustic and bone-conducted signals to achieve the best sound quality.



Sonion's hermetically sealed Voice Pick Up Bone Sensor was designed to improve speech intelligibility, device wake-up, high-quality hands-free phone calls, and voice commands directly inside the ear.

is combined with the external microphone's beamforming signal and intelligently mixed, resulting in full-spectrum speech. OnlyVoice also includes Voice Activity Detection (VAD) functionality, allowing low-power, always-listening operation of voice-activated devices.

Alango has been perfecting this technology combining microphone and bone conducted signals for several years, with several ongoing projects about to arrive in the market. Alango has been testing this technology with standard headphones, and reports effective results in the low part of the spectrum.

Sonion offers the Sonion Voice Pick Up Bone Sensor, which, as with Alango's OnlyVoice, is supported on Qualcomm's QCC audio SoCs. The combination allows audio manufacturers to implement reliable voice quality and voice recognition applications that are not hindered by surrounding noise. As Sonion highlights, the introduction of this type of sensor in TWS earbuds delivers the equivalent performance to a high SNR MEMS microphone, combining high bandwidth and low noise.

The Smart Directional Approach

When designing a voice capture solution that is not in the ear or worn by the user, there's a different approach to explore, which is to design a highly directional microphone and eventually combine the signal with sophisticated filters. TDK InvenSense is a company that specializes in MEMS sensor platforms and offers a complete range of MEMS microphones and all types of sensors, including accelerometers, combining proprietary algorithms and firmware to maximize the application's performance.

InvenSense became part of the TDK MEMS Sensors Business Group in May 2017 and the company's pace in responding to the fast-evolving market requirements has been significant. In the attempt to expand the applicability of MEMS microphones to areas that remain the domain of ECMs, one of the approaches was the design and implementation of a highly directional MEMS microphone.

The majority of MEMS microphones are inherently omnidirectional by virtue of having a single sound port (typically located on either the top or bottom of the package), which allows sound pressure to push and pull the diaphragm. TDK InvenSense developed the ICS-40800 dual-port MEMS microphone, which inherently has a bi-directional polar plot. This type of microphone is known as a pressure gradient microphone because its electrical output is proportional to the difference in pressure between the front and back of the diaphragm.

With this principle, the company managed to successfully demonstrate the potential of the ICS-40800 for directional microphone applications and create an effective alternative to beamforming arrays with multiple microphones for voice capture applications, all in a single package. Of course, the implementation of directional sound pickup using the ICS-40800 proved suitable when used with a gooseneck or microphone arm but far-field applications demonstrated important trade-offs. Because the microphone has both sides of the diaphragm exposed, the output sensitivity at low frequencies suffers and requires additional response compensation, while the application of the microphone in assembly processes requires extra precaution measures.

Soundskrit, a startup from Montreal, Canada, is exploring the same directional concept in two fronts. First by revisiting the MEMS microphone technology to create a high-performance directional device that is able to replace traditional microphone arrays. Second, by coupling this device with AI/ML trained models to optimize the capture and recognition of voice (and other) signals.

The company introduced the SKR0400 MEMS microphone with a dipole, or figure-8 pickup pattern, and claims to be able to maintain consistent directionality across the entire audio spectrum with very high SNR performance. The bold claims attracted the attention of industry heavyweights since its first demonstration at CES 2018 and the company says the MEMS device and integrated IC is now ready to be manufactured at scale. The formula for Soundskrit to solve the known trade-offs of the two-port directional approach is a combination of hardware and software.

The Soundskrit SKR0400 microphone is a miniaturized MEMS device with a new type of sensing diaphragm, combined with a DSP chipset and proprietary software that optimizes the audio response according to the application and the surrounding environment. The SKR0400 microphone transducer was designed from the ground to improve performance in directional microphones, delivering 63.5dBA SNR. According to the company, the MEMS element offers improved sensibility using a new capacitive readout architecture in a very small package, optimized to preserve directionality across all frequencies.


Soundskrit also developed an AI engine, based on the latest advancements in neural networks to unlock new features to developers. Soundskrit's SimplyClear software includes proprietary building blocks, compatible with both single and multi-microphone arrays.

The combination is so effective that Soundskrit demonstrates seemingly "impossible" applications, such as isolating a single voice in the middle of multiple surrounding conversations, and even separation of sound sources with a single microphone, only using software processing.


Smart Microphones

With the latest generation of edge AI solutions, which are small enough to be integrated into earbuds or hearing aids and powerful enough to run sophisticated AI/ML, it becomes clear that both the microphone polar response/sensitivity and the signal captured by the microphone can be adjusted to real-time conditions, which suits consumer use cases, particularly for personal and augmented audio. But of course, there are many applications where only an excellent microphone can provide the intended results, and that is why professional shotgun (super-cardioid) microphones from specialty manufacturers will continue to remain undisputed in TV and film sound production, in the same way as large diaphragms will remain the champions in studio vocal recording and certain broadcast voice commentary.

But a simple pass filter is not always the solution during a live transmission or committing a source signal to recording. Even the best professional microphones can be improved with AI algorithms running on device, such as to remove unwanted handling, the hum from air-conditioning systems or the rumble from wind. Of course,



In 2022, TDK Invensense introduced new digital MEMS microphones as part of its SmartSound family of products for mobile, TWS, and other consumer devices. The T5837 and T5838 PDM microphones push the boundaries of microphone acoustic performance, providing advanced feature sets in small package footprints. The T5848 offers the highest dynamic range with I²S direct audio connectivity to a wide variety of SoCs, MCUs, and IoT processors. The T5828 MEMS microphone with MIPI SoundWire functionality adds 68dBA SNR and Acoustic Activity Detect features with always-on ultra-low power mode.



TDK Invensense ICS-40800 dual port MEMS microphone inherently has a bi-directional polar plot allowing for highly directional applications.



Soundskrit has developed high-performance directional MEMS microphone technology combining hardware and software. The company's first product is the SKR0400, a dipole pickup device with a new transducer design, combined with a DSP chipset for acoustic background noise rejection and superior signal isolation.

this requires a set of skills and out-of-the-box thinking that not all specialty microphone manufacturers possess.

Maybe that is why we continue to see innovation on that front appearing from unsuspected innovators, particularly many startups, but also from established companies with a broader perspective of innovation. That is the case with Audeze, a company mainly recognized as a premium headphone manufacturer and pioneer of planar magnetic transducers. The California manufacturer has invested continuously in not only refining its manufacturing capabilities but also investing in audio signal processing, allowing it to enter new markets and new applications. Following interesting projects in the gaming space with binaural rendering of spatial audio sources, and advanced noise-canceling for medical applications, in November 2021 Audeze announced its first ever non-headphone product, the Filter personal conference speaker.

The Audeze Filter is a portable speakerphone equipped with intelligent voice processing that connects to virtually any digital source device over Bluetooth or USB and can be used

with all conferencing software. Even though it was launched with a crowdfunding campaign to test the market and had only a limited first production run, the product was named to *Time's* List of Best Inventions of 2022. The reason why the Filter personal conference speaker was found so innovative resulted from its elegant thin design using a 70mm × 105mm planar magnetic driver, combined with noise cancellation that is effective to filter the noise and leave only the users' voice. The small form factor with rechargeable battery also comes with an adjustable touch control panel, making a very practical speakerphone for work from anywhere. Its effectiveness results from the variable microphone beamforming capture, picking up all voices within 10' (3 meters), and edge AI processing to deliver perfect calls in any environment. It uses an advanced Intelligent Zero-Noise processing algorithm—trained on more than 500,000 noise samples—so it can accurately and effectively detect background noise and eliminate it.

Improving the Microphone

No matter how efficient and powerful the signal processing we use is, the results can always be improved by perfecting the performance of the microphone itself, and MEMS microphones are also evolving and getting better on all fronts. Most of the developments in this area are purely incremental, but still the pace of improvement is significant. Unfortunately, it's very hard to obtain specific information from many of the leading MEMS microphone manufacturers, which basically deliver OEM solutions optimized for specific projects and large clients, such as Apple. If we do know something about the technology improvements from Knowles, one of the leading names in the category and an Apple supplier, in contrast it's almost impossible to know what Goermicro (Goertek), and AAC Technologies are doing for Apple and other major clients.

Even a major supplier and one of the largest manufacturers of MEMS microphones, such as STMicroelectronics (ST), rarely makes any announcements in the product category. Its only announcement in 2022 was an update about its third generation of MEMS sensors combining ST's machine-learning core (MLC) integration and electrostatic sensing.

Examples of MEMS microphones products that we can highlight come from companies such as Infineon, a major innovator in the MEMS microphone space. The company's latest range of XENSIV MEMS microphones combine selectable power modes with its latest Sealed Dual Membrane (SDM) MEMS technology, which delivers high-ingress protection (IP57). The sealed MEMS design prevents water and dust from entering between the membrane and the backplate, avoiding mechanical blockage or electric leakage issues commonly observed in MEMS microphones. Furthermore, microphones built with the SDM technology can be used to create IP68 devices with the highest protection class requiring only minimal mesh protection.

The IM73A135 microphone achieves 73dBA SNR and 135dB AOP, while its digital counterpart, the IM72D128, offers an SNR of 72dBA and an AOP of 128dB SPL. Similar performance levels are available in further models that are smaller in size, lower power, or simply lower cost, meeting all requirements



The Soundkrit SKR0400 directional MEMS microphone is able to isolate a person's voice while eliminating any other sounds and noise surrounding it.



Audeze's Filter personal conference speaker uses a Deep Neural Network (DNN)—trained noise cancellation model combined with an effective far-field variable microphone design and the company's proprietary planar magnetic speaker.

TDK Invensense has explored the combination of improving the MEMS microphone designs for integration in digital signal processing in its SmartSound family for consumer-connected devices. The family offers a range of high-performance, ultra-low power digital MEMS microphones with 68dB SNR and 133dB acoustic overload point (AOP) in small packages and maintaining consistent performance even at very low power levels. This is achieved thanks to its Acoustic Activity Detect (AAD) features, which enable the microphone to monitor the acoustic environment at ultra-low power and then wake up a connected SoC or application processor when activity is detected. After introducing the first PDM digital output microphone with built-in Particle Ingress Filter (PIF) in 2021, TDK Invensense expanded its SmartSound family of performance MEMS microphones in 2022 with new I²S MEMS microphones that offer direct audio connectivity to a wide variety of SoCs, MCUs, and IoT processors, while the new T5837/38 PDM MEMS microphones improved performance levels for environments that shift from very quiet to very loud (e.g., far-field voice pickup for Smart Speaker to ANC TWS applications).

We also see continuous innovation from Knowles Corp., which recently celebrated its 75th anniversary and pioneered MEMS microphones under the SiSonic brand since 2003. The SiSonic MEMS microphones produced in two state-of-the-art manufacturing facilities in Suzhou, China, and Penang, Malaysia, continue to set

industry benchmarks, now combined with an “always-on” Acoustic Processor featuring Voice Wake and Voice ID keyword detectors. The latest low-power Knowles AISonic family of SmartMics was designed for mobile and TWS low-power developments, with a small footprint. Processing the captured sound source at the edge, SmartMic MEMS are integrated with a Tensilica audio-optimized DSP in a single, miniature, microphone package. The integrated programmable DSP is made available to third-party algorithms, enabling full customization and is offered with software, firmware, code and tools package.



sensiBel developed a new generation of MEMS microphones using its patented optical technology, which was based on years of research at SINTEF, one of Europe’s largest independent research organizations.



Audio Materials Science

WHEN SCIENCE MEETS INNOVATION

Your product idea could be great if you talk to the audio materials specialists early. Save yourself tons of headaches. Talk to someone who understands material properties, manufacturing processes, acoustics, and actual product design and development. Because our search for solutions started decades ago. Audio Materials Science is here!

Audio Materials Science is part of Menlo Scientific, a specialist audio consultancy firm, offering product development, design and technical support, since 1983.

Want to step-up your green credentials? Interested in sustainable product design? Talk to us.

MENLO SCIENTIFIC


Richmond, California, USA

- tel 510.758.9014
- mike@menloscientific.com
- www.audiomaterialsscience.com

Engineering Success since 1983

Knowles is also increasingly cooperating with specialized software and AI companies for these integrations, offering packages for voice recognition from Fluent.ai and the Voice Communication Package (VCP) from Alango Technologies, as just two examples. Another collaboration with AISTorm, a pioneer of AI-in-Sensor SoCs for edge devices, allows analog audio signal pre-processing for acoustic activity detection (AAD) or voice activity detection.

AISTorm’s analog processing of incoming audio data, prior to digitization, helps to perform advanced signal processing tasks such as noise cancellation, beamforming, keyword spotting (KWS), speaker identification, sound classification, and other algorithmic sound enhancements at very low power levels. An important factor



The sensibel optical MEMS microphone offers a variety of operating modes to integrate with complex headphone systems, to optimize for performance, latency and battery life. sensibel’s high SNR microphone is also able to handle large low-frequency sounds up to 150dB, which can easily occur in close-to-mouth applications. sensibel does not overload before the sound pressure level reaches 140dB—even 150dB at the lowest frequencies.

for differentiation in true wireless earbuds. Of course, Knowles’ close integration between the MEMS device and signal processing is key for this approach.

Different MEMS Microphone Technologies

Around 2010, the MEMS microphone’s SNR was historically inferior to that of ECMs. Five years later, most MEMS could still not reach beyond 58dB SNR and MEMS remained vulnerable to high sound levels and wind saturation distortion.

One of the issues was that Knowles had several packaging patents on MEMS, which it effectively enforced. But the competitive market evolved dramatically and MEMS microphones quickly took over ECMs on the volume application markets. Commodity capacitive analog MEMS microphones reached at least 60dB SNR, and 64dB SNR is available on higher-performance MEMS microphones.

When Vesper MEMS appeared in the market in 2017, the company’s approach introduced a fundamental change in the MEMS microphone space by proposing a new transducer technology. Since 2009, Vesper has built a broad and deep patent portfolio around its proprietary piezoelectric technology, which worked quite differently from conventional capacitive MEMS microphones. Vesper’s piezoelectric diaphragm is divided into four triangular cantilevers and is water, particle, and chemical resistant so it can be deployed in extreme harsh environments.

Yet, manufacturing and yield challenges kept Vesper from gaining market traction and its 70dB roadmap was unfortunately never realized. Unable to find a strong semiconductor foundry for the critical fabrication of the piezo-MEMS microphone element, it appears that Vesper’s future is tied to a rumored acquisition. Yet, no official notice has been released at the time of writing.

Now ranking at the top as the most promising company in the MEMS microphone space is sensibel, a Norwegian company with patented optical technology that it says is able to offer 80dB SNR, 140dB SPL AOP, and 10Hz low-frequency roll-off, rivaling

Resources

M. Klasco, “Optical Transducers and the Laser Microphone, Extract Clean Sound in Noisy Environments,” Voice Coil, December 2015
<https://audioxpress.com/article/Spotlight-Optical-Transducers-and-the-Laser-Microphone>

“Knowles, Goermicro, Omron, shoutout to the consumer MEMS microphones,” Yole Group, October 2022
www.yolegroup.com/press-release/knowles-goermicro-omron-shoutout-to-the-consumer-mems-microphones

“MEMS Microphones Market Worth \$2.99 Billion By 2025 | CAGR: 14.4%,” Grand View Research, August 2019,
www.grandviewresearch.com/press-release/global-mems-microphones-market

Research in China, “Micro Electronic-Acoustics Component Market (Global & China),” audioXpress, <https://audioxpress.com/news/Micro-Electronic-Acoustics-Component-Market-Global-China->

A. van Rhijn, “Integrated Circuits for High Performance Electret Microphones,” Audio Engineering Society Convention Paper, March 2003, www.aes.org/e-lib/browse.cfm?elib=12557

“Using MEMS Microphones in Directional Applications,” Application Note AN-000172, TDK Corp., 2020
<https://invensense.tdk.com/download-pdf/an-000172-using-mems-microphones-in-directional-applications-v1-0>

Sources

Alango Technologies, Ltd. | www.alango.com

Audeze | www.audeze.com

Infineon Technologies AG | www.infineon.com/mems

Knowles Corp. | www.knowles.com

sensibel | www.sensibel.com

Soundskrit | <https://soundskrit.ca>

STMicroelectronics | www.st.com/mems-sensors

TDK Corp. | InvenSense | <https://invensense.tdk.com>

Vesper Technologies, Inc. | www.vespermems.com


conventional studio microphone capsules, with all the benefits of a fully encapsulated solution.

The company has already received the necessary funding to ramp up production of its SBM100 series, which is rated for 94dB sensitivity at 1kHz (-46 dBFS) and 126dB dynamic range. The device uses a top port and offers a variety of data interfaces, including I²S, PDM, and TDM. A very impressive set of specifications, for which many in the audio industry have been hoping, breaking the signal-to-noise barrier that would take MEMS technology closer to conventional condenser capsules.

At the core of this performance is sensiBel's patented optical technology, based on the optical phenomena of interference and diffraction observed in light waves. In simple terms, the sensiBel microphone uses a tiny energy-efficient laser in combination with a miniaturized and integrated interferometer to measure the motion of a silicon-based membrane with extreme accuracy. The entire optical system is in fact smaller than a cubic millimeter!

This ultra-low self-noise readout principle offers the higher SNR and AOP characteristics that are challenging for existing MEMS microphones, while keeping the same small MEMS package and low power consumption. The potential for this technology would represent a breakthrough of 10 times better performance than today's state-of-the-art MEMS microphones.

We hope the next update on MEMS microphones from *audioXpress* will include mention of availability of optical MEMS



The Audio Group of CUI Devices released two analog MEMS microphones purposely designed for active noise cancellation applications. The new models with high SNR and AOP ratios deliver all the features required by the market today, including low power in compact packages to meet the voice capture and audio recording needs of any design, available with omnidirectional, unidirectional, and noise canceling directivity.

microphones, even though actual products in the market remain likely a couple of years away. Meanwhile, we will certainly see some further innovation powered with a bit of help from AI and—who knows—maybe a few more MEMS microphones players emerging. 

The advertisement features a large, stylized background with a yellow and green grid pattern. On the left, a red USB drive is shown with a black cap that has the 'audioXpress' logo and the tagline 'ADVANCING THE EVOLUTION OF AUDIO TECHNOLOGY'. To the right, the text 'aX Cache' is displayed in large, bold letters, with 'aX' in red and 'Cache' in black. Below this, the text 'Every issue of audioXpress published since 2001, on USB.' is written in bold black font. At the bottom, the website 'cc-webshop.com' is shown in large red letters. In the bottom right corner, a stack of audioXpress magazines is visible, with the top issue showing the title 'audioXpress' and the tagline 'ADVANCING THE EVOLUTION OF AUDIO TECHNOLOGY'. The top issue also features the headline 'Fresh From The Bench Audiomatica's CLIO Pocket Electro-Acoustical Portable Measurement System By Joe D'Appolito' and the section 'R&D Stories'.



Unified Coverage Map

A Radically Better Technology for Today's Hybrid Environments

This article discusses a new patent-pending feature of Nureva's next-generation technology for audio presentations and conferencing that provides a new perspective on how microphone coverage works in large meeting rooms.

By **Erica Yorga**

Photo 1: This is an image of a Nureva tiered classroom. Note the Nureva HDL410 wall-mounted microphone and speaker system on the walls, on each side.

The rise of hybrid working and learning has brought challenges for anyone who equips and maintains shared spaces. The demands in meeting rooms, boardrooms, open spaces, collaboration spaces, and classrooms are constantly evolving.

In many cases, technology and how it is embedded in the systems to support these spaces has become more complex, requiring greater technical know-how to work with it. The need to manually configure and change out equipment is outdated. Reducing system design, installation, and commissioning overhead along with technical management and complexity is paramount, especially for organizations that are trying to simplify and reduce costs as they are deploying tens, hundreds, and even thousands of systems.

How can applying innovative thinking to the technology and how it is embodied in products provide a fresh, adaptive approach to evolving requirements and spaces? Audio conferencing systems should be easy to install and then automatically configure and optimize themselves for the room.

Nureva's unified coverage map, a feature of patented Microphone Mist™ technology, takes a radically different approach by providing automated configuration capabilities that results in consistently optimized audio performance and reduced total cost of ownership. Technology has evolved, and organizations now need to support increasing numbers of remote participants who demand

reliable and consistent far-end audio performance. The unified coverage map was designed to meet this need (**Photo 1**).

The Large-Room Conferencing Challenge

The challenges that come with using legacy audio-conferencing technology aren't always apparent in small spaces that only require one audio device.

The installation, configuration and optimization of a single conferencing device can be straightforward (**Figure 1a**). In small spaces, where one device is sufficient, room coverage and setup may not be an issue. **Figure 1b** shows how Microphone Mist technology also provides full-room coverage in smaller spaces, using thousands of virtual microphones.

But for legacy microphone array systems, their inherent complications become readily apparent as room size increases. These systems require decisions to be made about the number of array devices and their placement relative to the type of room coverage needed, along with any inter-device spacing requirements.

The Old Approach

In legacy systems, each array device operates independently and is responsible for a specific area and coverage pattern, resulting in overlap that needs to be minimized (**Figure 2a**). This requires technical

specialists with detailed knowledge about acoustics and each device’s specific capabilities. Parameters such as number of active zones and beam width and shape all need to be configured for each device.

To manage the audio pickup and processing decisions, a third device is required—a digital signal processor (DSP). The DSP determines how to mix and match audio outputs of the array devices to provide room coverage and audio processing functions, such as Acoustic Echo Cancellation (AEC), gain structure (GS), noise reduction (NR), equalizer (EQ) and coverage zone configurations. These functions may be configurable on the individual devices or in the DSP unit.

As incremental devices are added to the system, there is typically no awareness among the devices to automate device adoption and the configuration process. Each device is treated as a separate, stand-alone device with capabilities that need to be manually integrated into the DSP.

This time-consuming and technically involved process of configuring system components is how it has been done historically; but in modern hybrid spaces customers need/demand more. Systems that require manual configuration of separate DSPs, microphones, and speaker devices for even the simplest of installations don’t fit with the challenges of modern hybrid working and learning spaces. It’s time for a change.

A Fresh Approach

We believe that large room audio systems should have the same minimal setup and configuration, and great performance as small rooms. We leveraged our Microphone Mist technology and augmented it with a new unified coverage map to deliver an automatically configured conference system that scales from small to larger spaces. This approach blows up legacy paradigms and sets the stage for successful and cost-efficient deployments.

The unified coverage map maintains the advantages of a uniform virtual microphone coverage map solution (already available with Nureva’s single array devices) and scales it to seamlessly adapt in small to large rooms (**Figure 2b**). It enables the physical microphones from two devices to work together, creating a combined microphone array that covers the entire room. The overall installation and system commissioning process remains the same—it’s minimal.

With Nureva’s unified coverage map, patent-pending calibration signals are used to derive a real-time system configuration and an optimized room coverage map without the need for technician level know-how. And this happens automatically when




Figure 1: The disparate coverage zone for legacy systems (a) and a unified coverage map using the Nureva systems (b) are shown here.

the system powers up. The result is a dynamically created unified coverage map that is adapted to the device locations. Unlike legacy systems (Figure 2a), the coverage map looks identical in shape and distribution in both small and large rooms. There is no device and coverage zone overlap to consider and manage.

Why Use a Single Unified Coverage Map?

The optimal way to provide the best room coverage and audio performance would be to install a large 3D microphone array with hundreds or even thousands of physical microphones spread throughout the room. Of course, this is not practical given the cost and complexity. But Nureva’s unified coverage map is the next best thing. There are many benefits to our approach that overcome the limitations of legacy systems.

Targeting conflicts between disparate array devices and zones are eliminated. Each virtual microphone has a unique location with no coverage




Figure 2: The disparate coverage zones for legacy systems have an inevitable overlap (a); while the Nureva systems have a unified coverage map (b).

overlap that requires blending (mixing) and sharing decisions, so the targeting and audio signal path is clean and optimized for all locations in the room.

All physical microphones are available to provide gain to any virtual microphone in the room, meaning consistent talker and background noise levels with significantly reduced noise artifacts and system distortions. This differs from legacy systems that limit gain to a single array device or zone, creating uneven gain and background noise levels and inter-device distortion artifacts.

Remote participants hear a consistent and natural tonal quality for each in-room talker because there are no handoffs between devices or zones that may be configured differently. For example, differing zone beam widths (wide, medium, or narrow) can have inconsistent and detrimental effects on the sound quality at the far-end of the call.

Because each virtual microphone location uses the same audio processing chain, the chain is simplified, optimized, and stable. It eliminates unexpected configurations and behaviors that can occur with a hierarchical or cascaded audio process chain where each device does its own audio processing.

Individual virtual microphones do not overlap and therefore can have unique targeting parameters applied to ignore or enhance fixed positional-based sound sources in the room, such as HVAC vents

and fans. This makes it much easier to effectively manage undesired sound sources. If an HVAC vent is situated between two coverage zones in a legacy system, each device in the system will need to deal with the unwanted noise source. If the undesired sound source is in the middle of a broad zone, it negatively impacts all desired sounds in the zone.

People moving and talking at the same time is easily managed with the unified coverage map. One virtual microphone with a set position never conflicts or overlaps with another virtual microphone with its own unique position. The talker easily moves between virtual microphones as only the focusing parameters need to adjust, resulting in consistent audio performance. With legacy systems, zone and device coverage overlap are often difficult problems to deal with, necessitating complex handoff and decision logic and combining different audio processing chains.

Placement of the physical array devices is not as critical. With legacy systems, each array device becomes responsible for a very specific area of the room with its own coverage map and configuration. With the unified coverage map, the additional device is incorporated into the physical array, providing gain and increased coverage map support. It is like adding more horsepower instead of adding a second engine to drive a second wheel.




Figure 3: The configuration for disparate coverage zones with legacy systems (a) is compared to the configuration of a unified coverage map with Nureva systems (b).


Commissioning Time Makes All the Difference

These benefits make a significant difference in required commissioning time and cost. With legacy systems, room size and complexity often require a lengthy commissioning process. But with the automated, unified coverage map, the whole process takes only a few minutes. As the number of rooms increases, the savings in technician costs add up fast.

With a legacy audio-conferencing system, the complexity of configuration and commissioning increases when new devices are added (**Figure 3a**), leading to many hours of work for a specialized technician. In step 1, a manual configuration of the network and device parameters may be required, with some systems needing their own dedicated local room network. Once configured on the network, each virtual audio path connection from device to DSP is completed by the technician (step 2), which can be a complex and detailed technical task. In step 3, the system coverage patterns and shapes need to be configured for each device, taking into account microphone array device type and location, available coverage patterns, number of zones, zone

overlap and beam widths. In step 4, the technician must review the audio parameters for each device to determine if specialized parameters, such as AEC, EQ, NR, delay and gain, are needed for each device or if global parameters are sufficient. Then thorough testing, step 5, is required to determine if the correct setup choices were made. If needed, adjustments and calibrations are implemented and then the whole process may have to be repeated until the system achieves the required performance and the commissioning is complete.

Nureva's unified coverage map dramatically simplifies this process (**Figure 3b**). The first step is to plug in all devices and turn the system on. In step 2, the system sends patent-pending proprietary calibration signals between each device. This allows the system to locate the two devices in 3D space and then build a combined physical array. The system knows the precise location of each device relative to each other. In step 3, this information is used to dynamically construct the unified coverage map, providing complete coverage matched to the room. Since the individual arrays are seen as one physical combined array, the system can calibrate and optimize for AEC, NR, and gain parameters

 **ACOUSTICS**
building your sound

10" SW26SFC38-4/8

- Shallow light weight design.
- Hard pressed vented paper cone.
- Inverted low damping rubber surround.
- Symmetric suspension.
- Stacked neodymium motor with built-in soft limiting
- Non-conducting fiber glass coil former.
- Long life lead wires.



Find us on:

-  SB Acoustics
-  sbacoustics
-  sbacoustics_
-  SB Acoustics
-  sbacoustics.com

and does not require device-specific hierarchical configurations or blending coverage maps. There is no need for a technician to manually configure and optimize the system at any step of the process.




Figure 4: Sound source targeting and processing with legacy systems (a) is shown in comparison with sound source targeting and processing with Nureva systems (b).

About the Author

Erica Yorga is the Principal Intellectual Property manager at Nureva, a Canadian technology company. Her career spans a variety of leadership roles researching, developing, and commercializing a broad spectrum of technologies in markets such as audio/visual, oil and gas, telecommunications, military defense, and Government regulatory. The challenge of solving difficult problems by constantly learning and enabling teams through the application of unique approaches to technology and process is a passion she brings to every endeavor she undertakes.



Nureva Inc. is a technology-rich, multiple award-winning private company that imagines and builds audio conferencing solutions that solve the frustrating and persistent problem of poor audio performance in meeting and learning spaces. At the core of every system is the company's patented Microphone Mist technology, which places thousands of virtual microphones throughout a space to pick up sound from anywhere in the room and deliver clear, reliable audio to remote participants. A passion for achieving simplicity through deep user understanding drives the company's product roadmap and the value it creates for its customers. For more information, visit www.nureva.com.

Resources

E. Yorga, "The Impact of Acoustic Spatial Resolution in a 3D Space," *audioXpress*, February 2021, <https://audioxpress.com/article/r-d-stories-how-virtual-product-development-improves-the-quality-of-consumer-and-pro-audio-products-1>

Optimized Audio Targeting and Processing

The second key differentiator for the unified coverage map relates to audio targeting and processing. While the specifics of the process may vary between different legacy systems, it is clear to anyone who has experience with them that they are too complex. In legacy systems the challenges and problems are well known and persistent.

The first step with a legacy system is to correctly target a person talking in the room (**Figure 4a**). With distributed devices and coverage zones, this is a daunting technical challenge. Each array device is responsible for a certain part of the room, and the audio system is required to select the correct individual or combination of devices and zones within the disparate array devices to target the talker. If the talker is located within the overlap between zones or between two array devices, this is a real challenge.

A good analogy for the target problem is what happens when a cell phone call is caught between two or more cell towers. The issues that arise—such as cross talk, messy handoffs, dropped calls and poor voice quality—are similar to what happens when using separate microphone arrays in a legacy conferencing system. They all stem from each individual array (or cell tower) covering its own area with overlapping coverage. The challenge becomes more acute if the talker is mobile and moving throughout the room (or driving between towers). Because each array device works independently, the system must decide how to switch, blend, or ignore disparate signals from each device. The legacy conferencing system needs to adapt and constantly make correct targeting and mixing decisions at the DSP to ensure that the far end audio quality stays optimal.

To add to the complexity, the audio processing function may be hierarchical or combined between devices. There may be multiple layers of audio processing depending on how the system is configured, and each array device may do its own gain, EQ, and NR processing. Once the signals are received by the DSP, additional processing on top of what the individual array devices have implemented may be completed. The processing may be device dependent, or DSP driven, and it is hard for the customer to know.

This all increases the chance that there will be inconsistent processing and incorrect device configurations and management. The UC&C client will be required to adapt and adjust, leaving the remote participant with inconsistent audio performance. When an organization manages multiple conference rooms, maintaining consistent system audio quality across all rooms may be very difficult to do.

It does not need to be like this. The unified coverage map (and the broader Microphone Mist technology platform) has greatly simplified and optimized the targeting process.

As shown in **Figure 4b**, first the system locates and focuses on the sound source within the unified coverage map using patented targeting algorithms. There is no need for complex switching between array devices and individual array coverage zones or blending between zones and separate array devices. A large room installation performs just like a small room.


Once the array is focused on the appropriate virtual microphone within the unified coverage map, Microphone Mist technology proceeds with audio processing. A single combined array focused on one virtual microphone makes for an efficient and optimized audio processing path with no unnecessary inter-device cascading or duplication of processing. The result is consistent and improved acoustic echo canceler, gain structure, background noise, and noise reduction performance. It's no longer necessary to adapt to a different array device that may have its own configuration parameters for AEC, PEQ, NR, and gain. The UC&C client will always be presented with a consistent audio signal, and audio performance at

the far end of the call will stay optimized. Nureva's novel architecture simplifies the audio processing chain, regardless of the complexity or size of the room in which the system is installed.

Summary

The unified coverage map delivers a consistent, superior audio experience while eliminating the complexity of legacy audio system setup and/or configuration. It automatically adjusts and adapts to the room, significantly reducing the system installation and management burden while optimizing audio in large spaces. It brings a radically new total customer experience into rooms where that was never possible.

As more rooms are equipped with audio conferencing products, choosing an approach that scales is becoming more important. By taking advantage of fresh approaches, such as Nureva's unified coverage map and continuous autocalibration as discussed in my February 2021 *audioXpress* article, large room conferencing can be just as simple and effective as small room conferencing. And that's what is needed to meet the demands for hybrid working and learning—now and in the future. 



The Electrostatic Speaker Design Cookbook
is a complete guide for the novice and the experienced builder to successfully build this unusual sound reproduction device.

THE ELECTROSTATIC LOUDSPEAKER DESIGN COOKBOOK
by ROGER R. SANDERS
With a Ground Cell Design by Barry McGuire

cc-webshop.com



DD12

Power Entry Module with EMC filter and ground line choke

DD12 suppresses EMI / RFI, optimizes shielding, and eliminates hum or buzz generated by ground loops

- IEC C14 appliance inlet / 1- or 2-pole shocksafe fuseholder / 2-pole switch
- Standard or low leakage filter ratings 1-10 A
- IEC / UL ratings 10 A @ 250 VAC / 8 A @ 125/250 VAC
- Top or side flange mount; Snap-in for panels 1-3 mm
- V-Lock cordsets with varying plugs and cord lengths
- ENEC and cURus approved

SCHURTER
ELECTRONIC COMPONENTS

schurter.com

On the Subject of Standards

Filling in the Gaps Left by AES67



Is the AES67 interoperability standard a viable alternative as the core solution for audio-over-IP distribution? Crestron proved that it is with its DM NAX technology filling the existing gaps of network discovery and routing.

By **Ekin Binal**

(Director, Product Management,
AV Solutions, Crestron)

Since its founding in 1948, one of the more impactful aspects of the Audio Engineering Society (AES) is the ongoing work of the nonprofit's Standards Committee. Its mission, according to the AES:

AES Standards bring the industry's engineering teams together to facilitate new technologies, improve performance and ensure interoperability. AES3 (2-channel digital audio), AES10 (MADI), AES14 (analog XLR pin-out), AES67 (networked audio)—AES Standards have contributed to your operations, making your work more successful, improved your workflow, and saved your production, more times than you realize.

Standards are an interesting concept. Ideally, the goal is to level the playing field for multiple parties to create and proliferate products utilizing a shared technology. At Crestron, we are constantly working with forums and standards bodies to ensure our products are developed to integrate seamlessly into all of our customers' complex environments. A great example of this is our contributions to the video side of things, specifically with HDMI and HDCP—by working with the standards bodies we've been able to influence improvements in interoperability and video handshakes. These improvements not only enhanced our products but allowed the industry as a whole to benefit from our testing—we've created a feedback loop in which we share our findings with the

standards bodies. That's the beauty of those bodies: They are constantly evolving and improving based on implementation.

AV has been a discipline that has seen its fair share of standards: The good (e.g., DVD and Blu-ray), and the challenged (e.g., Betamax and HD DVD). It's been a truism throughout the history of various standards that the best solution doesn't always win out in the end. AV-over-IP is now opening up another similar challenge. On the plus side, we at least have a common interface with Ethernet on Cat5e (or greater) cabling, but the software layer that runs on the network is a whole different picture. Unfortunately, there are many competing protocols running on the network to move AV data, and the industry has been challenged to make all these platforms play together nicely.

Most recently, Audio-over-IP (AoIP) has emerged as a popular method of distributing audio at a highly scalable level by leveraging the network. Enter the AES and the AES67 standard. AES67 is a suite of requirements defined to facilitate the interoperability of AoIP platforms. Popular AoIP standards that existed well before AES67 was created are Dante, Ravenna, Livewire, and QLAN. The AES looked at the common underlying technology across the various AoIP platforms noted above and built a common denominator of requirements, including a synchronization method, audio formats, and transport protocols at the highest level.

Audio Synchronization

Of these requirements, the one worth highlighting is audio synchronization. The key concern of distributing audio on a digital network is to make sure that all the audio endpoints (amplifiers, speakers, etc.) are receiving that audio at precisely the “right” time, otherwise beyond egregious issues such as clicks, pops, and echoes, you also risk the quality of the audio integrity due to phasing issues. Luckily, the Institute of Electrical and Electronics Engineers (IEEE) created a common technology called precision time protocol (PTP) that defines a clock synchronization technology (IEEE 1588) that many AVoIP platforms have adopted as the preferred method to line up their clocks to provide synchronous audio on the network. Without this core technology on all AVoIP products, AES67 would not have been possible.

THE DM NAX AoIP Technology

In 2020, Crestron launched DM NAX AoIP technology, and rather than complicate the matter with yet another competing standard, we embraced the AES67 standard and built our solution around it (Photo 1). For those who haven’t dabbled with AoIP technologies such as Dante yet, one of the most critical differentiators that Dante (created by Audinate) brings to the table is a full suite of software applications that simplify controlling and monitoring Dante AoIP audio endpoints. Dante



Photo 1: When you build a Crestron AVoIP solution today, you can plug in your various HDMI sources into DM NVX transmitters, and then output the audio to DM NAX receivers.

has a tool called Dante Controller that can run on any personal computer and, when deployed, automatically discovers all the Dante AoIP endpoints and builds a dynamic matrix of inputs and outputs. The other commercial solutions have created similar tool sets, and this is great for use in the wheelhouse of those platforms—but it’s not helpful when using other AoIP standards. This is where AES67 helps bridge the gap, as when platforms are put into AES67 compatibility mode, it allows them to interop.

While it’s certainly true that AES67 brings a great many benefits to the table, there are downsides. The platform omits a few critical pieces—namely, the enforcement of discovery and a universal method to make and break connections. This omission means that it can be much more difficult to find all the AoIP devices on a given network. Additionally, if you could find those devices, you would then struggle with sending and receiving audio between them.




Photo 2: The Crestron DM NAX line of networked streaming preamplifier and amplifiers diversifies the range of choices for homeowners and integrators looking for residential AoIP distribution platforms. The range natively integrates with Crestron Home and the highly scalable AES67 AoIP distribution system supports hundreds of sources to as many as 256 stereo zones (512 channels total), with the ability to provide a dedicated stream for each zone rather than just one or two streams to the whole house.

With Crestron’s DM NAX technology, we have differentiated the product by using all the greatness of AES67’s interoperability but adding in the missing parts: discovery and routing. DM NAX is new as a solution on dedicated audio devices, but the fact is we first implemented this core DM NAX technology on our AVoIP platform DM NVX.

DM NVX devices were built to carry HDMI sources to HDMI endpoints through the network, so we evolved the DM NVX platform to break away the audio from the video flow and insert or extract that to dedicated audio components by way of the DM NAX technology. The result: When you build a Crestron AVoIP solution today, you can plug in your various HDMI sources into DM NVX transmitters, and then you can output the audio to DM NAX receivers (Photo 2).

Without the additional technology that Crestron built on top of AES67, one would not only struggle to find all the transmitters/receivers on a network, but it would be very challenging to route the AV traffic. Historically, AES67’s interoperability would require you to log into the web UI of the receiving endpoints and select what audio stream to listen to. This is because AES67 transmitters are putting out network audio when enabled, and at that point, it’s the receiver’s responsibility to select the stream. Even if one were to use AES67 compatibility mode with Dante, you’d still have to juggle a secondary interface to make the “link”—again, this is because one can transmit audio out blindly, but the heavy lifting is on the receiver. It’s tasked to find and select what it wants to listen to.

Figure 1 illustrates the manner in which all this is simplified on Crestron DM NAX/NVX devices. The Matrix finds all the DM NAX sources and receivers and allows one-click routing between them. This was the primary motivator for Crestron to build DM NAX technology: There was tremendous potential in AES67, and Crestron recognized it and took it to the next level.

Loftier Aspirations

Beyond solving the “boots on the ground” issues of interoperability, there’s something loftier at work here. The general benefits of standards go well beyond making the dealer/installer’s life easier. They benefit the consumer, too—when a variety of devices are assured to “work and play well” with one another, that can inspire customer confidence across an entire industry. For example: A consumer who knows that all the equipment in a system—no matter the brand—will operate together in a manner to deliver the audio result they desire can well become a repeat customer. Upgrading single devices over time is much less overwhelming than replacing an entire system when the aspirational consumer goes hunting for a better experience. (There are issues of cost, convenience, and sustainability that come into play, too. Note what the EU has to say about standardizing devices to the USB Type-C charging port.) With AV over IP, the whole benefit of a shared network infrastructure for interoperability is reliant on these standards.

Standards also allow a framework in which to create new products instead of starting from scratch. They provide foundational baselines for further innovation. Orson Welles allegedly said, “The enemy of art is the absence of limitations.” The same might be said of technology without standards. That is, after all, the underlying principle: A universe of devices that cannot talk to one another is simply chaos.




Figure 1: The Matrix finds all the DM NAX sources and receivers and allows one-click routing between them.

About the Author

Ekin Binal is the Director of Product Management AV Solutions at Crestron Electronics, responsible for delivering innovative audio/video products across residential, commercial, and enterprise markets. Ekin leads Crestron’s initiative to modernize audio connectivity, control, and management for a truly integrated experience. Ekin is an industry veteran with more than 15 years of experience in research, development, technical management, product development, and generating new revenue streams. His diverse background ranges from chip design at IBM and product development at BICOM to serving as the VP of Business Development at Sonavox.



The “Must Have” reference for
loudspeaker engineering professionals.

Home, Car, or Home Theater!



Back and better than ever, this 7th edition provides everything you need to become a better speaker designer. If you still have a 3rd, 4th, 5th or even the 6th edition of the *Loudspeaker Design Cookbook*, you are missing out on a tremendous amount of new and important information!

Now including: Klippel analysis of drivers, a chapter on loudspeaker voicing, advice on testing and crossover changes, and so much more! Ships complete with bonus CD containing over 100 additional figures and a full set of loudspeaker design tools.

A \$99 value!

Yours today for just \$39.95.

Shop for this book, and many other audio products, at
www.cc-webshop.com.



Suitability of Venues for Various Events

Understanding the optimal reverberation time vs. frequency performance of a venue, particularly one used mainly for amplified music, is important for acoustical consultants and architects. Recent research has provided valuable design guidelines for such venues.

By
**Richard
Honeycutt**

Ever since performance venues evolved from threshing floors to amphitheaters and then to indoor rooms, architects, physicists, and acousticians have tried to improve them. The first efforts were aimed at making the sound loud enough for listeners to understand actors' speech in dramatic performances. Later, attention was focused on eliminating echoes and controlling reverberation.

When Harvard Physics Professor Wallace Clement Sabine defined reverberation time (RT) in the late 19th century, his study was inspired by the need to improve the acoustics of the Fogg Lecture Hall. His work constituted the foundation of modern scientific acoustics. Sabine and others soon recognized the importance of RT in concert halls and other buildings. In fact, when Sabine's student, Dr. Vern Knudsen, published his landmark book [1], it included chapters on schools, public and commercial buildings, churches, open-air theaters, theater buildings, music halls, residential buildings, and studios for broadcast and recording.

Auditoria in particular became a focus of attention for good acoustical design. Although the uses for an auditorium could be grouped in three categories—lecture, drama, and music—each with its own acoustical requirements, acousticians realized that the “music” category included important sub-categories according to the type or genre of music to be performed. In his acoustics textbook [2], Leo Beranek divided the “music” category into broadcast studios for music, opera, concert halls, Protestant churches, synagogues, Catholic churches, and halls for pipe organ music.

The Evolution of Acoustical Design

Until the latter half of the 20th century, acoustical design included only noise control, echo control, provision of adequate loudness, and reverberation control. Then other acoustical measures began to be examined. These included clarity of speech and music, definition, envelopment (the listener's sense of being enveloped by the sound), apparent source width, spaciousness, warmth, intimacy, bass balance (bass ratio), treble balance (treble ratio), and speech intelligibility.

Acoustical strength (sometimes called acoustic gain) was mathematically defined. These perceptual parameters are somewhat intertwined in terms of physical cause and architectural design. Yet the understanding that RT is not the only—perhaps not even the primary—factor in listener evaluation of a concert hall was slow to dawn.

In Beranek's 2016 paper on concert hall acoustics [3], he cited his 40-year effort to evaluate listener preference in concert halls, involving 58 halls, and Tapio Lokki's work at Aalto University in Finland [4], pointing out that listener preferences could be grouped into two: those preferring high definition and clarity along with adequate loudness, reverberation, and bass; and those who enjoyed a louder and more reverberant sound with strong envelopment and strong bass.

While RT is an important element for both groups of listeners, the new mathematical definitions of acoustical strength, definition, bass and treble ratios, and clarity must also be recognized as important factors.

Acoustical consultants have long been aware of the differing acoustical targets needed for different types of music: the Gregorian chant would sound sterile if performed in Johann Sebastian Bach’s relatively dry (low RT) Thomaskirche, and much of Bach’s music would sound like a muddy mess if played in a cathedral. Even romantic music such as Pyotr Ilyich Tchaikovsky’s and classical music such as Wolfgang Amadeus Mozart’s require different acoustics for optimal listening.

Yet despite the growth of jazz and country music concerts beginning in the early 20th century, rock music in the mid-century, and folk music after the mid-century, few if any studies had been done determining the optimum acoustical parameters for these genres, although they now account for most of the concerts. Consultants have primarily based acoustical design decisions upon the experiences of themselves and their mentors, coming up with guidelines that had been shown to lead to audience satisfaction.

Recent Studies

In today’s concerts involving the newer genres previously mentioned, one characteristic is shared in common: the use of electroacoustical amplification. This factor changes multiple design considerations for music performance venues. In 2010, Niels Werner Adelman-Larsen, Eric R. Thompson, and Anders G. Gade published a study objectively and subjectively assessing RT for rock and pop music halls [5], followed by a study by Adelman-Larsen, Cheol-Ho Jeong, and Bård Stofringsdal [6] examining subjective listener responses.

Research dating as far back as 1931 [7] show that the average sound energy per Hertz of various instruments’ and orchestras’ peaks anywhere from about 30Hz for a bass drum to 2kHz for a piccolo, (**Figure 1**). Many acousticians have long used the 400Hz to 600Hz range as the typical frequency range of maximum energy in music performances. Adelman et al observed that while the mean SPL in the mid/high frequencies at a classical concert may be 10dB higher than in the bass range, for rock and pop music, the bass SPLs might be 20dB to 30dB higher than the midrange SPLs.

Also, audiences at classical concerts are typically seated with a density of about two people per square meter; whereas at a rock concert, audiences are more often standing, with a density of about 2.7 people per m². These differences likely impact the optimal acoustics as much as differences in instrumentation and performance style.

The first study by Adelman-Larsen, et al revealed that the best-rated halls evidenced little

variation in RT across the octave bands from 63Hz through 4kHz. The worst-rated halls had a higher bass ratio, which would not only increase perceived bassiness, but would also muddle fast-moving bass lines of both melodic and percussion instruments. Having observed that most audiences at rock concerts stand during the performance, Adelman et al measured the absorption coefficient of a standing




Figure 1: Octave-band sound levels produced by talkers, instruments, and orchestras. (Image Source: H.F. Olson [8])

About the Author

Dr. Richard Honeycutt fell in love with acoustics after 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.

audience and found that such an audience absorbs about five times as much energy in the mid/high frequency bands as in the bass bands (**Figure 2**.) The researchers noted that when the average RT from 500Hz to 2000Hz was compared to the halls' subjective rating, the correlation was not good. However, by including the 63Hz and 125Hz bands' RT values, correlation improved greatly.

Figure 2: Standing audiences are more absorptive than seated audiences.




Figure 3: The normalized RT for best- and worst-rated halls clearly differs with frequency.




Figure 3 shows the octave-band RTs of the best- and worst-rated halls, normalized against the 500Hz to 2000Hz values. The best-rated halls showed more uniform RT vs. frequency. A question for future research is whether the approximately 80% increase in RT in the 63Hz band, as compared to the 500Hz to 2000Hz band average, is truly optimal, or whether it is an artifact of the 63Hz RT actually achieved in the halls studied. The paper by Niels Werner Adelman-Larsen, Cheol-Ho Jeong, and Bård Støfringsdal points out that the sound energy in the 125Hz band is often greater than that in the 63Hz band, and that there are indications of "a high tolerance toward a high T30 in the 63Hz band as long as the T30 in the 125Hz band is controlled." If the mid- and high-frequency RT is too low compared to that at 125Hz, listeners will hear the same deleterious effect as if the 125Hz RT were too high: it's not only absolute RT, but also the frequency balance of RT that matters.

Effect of Room Size

The human ear/brain system is capable of estimating the size of a room, and it accomplishes this feat primarily by noting the time delay between the onset of direct sound and the arrival of the first reflected sound, a room parameter known as the Initial Time Delay Gap (ITDG). Further, studies have determined that audiences' preferred RT in a hall depends upon the cubic volume. These conclusions have been based upon listener preferences for classical-music halls, but may need to be re-examined in regard to rock-pop halls.

Achieving a "flat" RT characteristic is difficult for several reasons. Most lecture and concert halls' acoustical absorption is provided by the

In the design of classical concert halls, consultants usually take bass ratio (the ratio of the reverberation time in the frequencies of 125Hz and 250Hz to the frequencies of 500Hz and 1000Hz) and treble ratio (ratio of the sum of 2000Hz and 4000Hz to mid-frequency late arriving sound levels at 500Hz and 1000Hz) into account. Bass ratio affects the perceived support for bass sounds in the room, and treble ratio affects the perceived brilliance of the room.

I personally design for a 63Hz RT under 1.4× the 500Hz/1000Hz RT, and a 125Hz RT under 1.25× the 500Hz/1000Hz value in multi-purpose auditoria. High-frequency RT is normally designed to decrease above about 1.25kHz, reaching a value of perhaps 70% to 80% of the midband RT in the 8kHz octave. Lecture halls should normally have less variation in RT vs frequency for best intelligibility (**Figure A**).




Figure A: Optimum target design values for RT vs. frequency differ according to the purpose of the venue.

audience, which absorbs more sound as frequency rises (**Figure 4**). Seating—especially upholstered seating—also absorbs more sound at mid and high frequencies. Making matters worse, directional loudspeakers aim most of the mid- and high-frequency sound at the audience, where it is absorbed, while much of the low-frequency sound hits less absorptive surfaces; thus low-frequency RT tends to be higher. Common acoustical treatment provides poor absorption below 250Hz.

Acoustical absorption can be achieved by use of porous materials that absorb sound via internal friction, by thin diaphragms such as 1/4" plywood panels over air spaces, or by a combination of the two (e.g., fiberglass panels with a hardened surface layer). **Figure 5** shows the sound absorption coefficients of several acoustical materials. In general, bass absorption can be expensive.

Conclusion

Understanding the optimal RT-vs-frequency performance of a concert hall used mainly for




Figure 4: Below 2kHz, a seated audience absorbs more sound as frequency rises.

amplified music, especially rock and pop music, is important for acoustical consultants and architects, and for too long it has been ignored, leaving acoustic designers to “fly by the seat of their pants,” with unpredictable and often unintended results. The studies by Adelman-Larsen et al discussed in this article are both valuable and timely. ax




Figure 5: The absorption of common acoustical materials are low at low frequencies.

Resources

[1] V. O. Knudsen, *Architectural Acoustics*, John Wiley and Sons, New York, 1932.

[2] L. Beranek, *Acoustics*, McGraw-Hill, New York, 1954.

[3] L. Beranek, "Concert Hall Acoustics: Recent Findings," *Journal of the Acoustic Society of America* (JASA), 139(4) (April 2016), 1548-1556.

[4] T. Lokki, "Tasting Music Like Wine. Sensory Evaluation of Concert Halls," *Physics Today*, 64(1), 27-32 (2014).

[5] N.W. Adelman-Larsen, E. Thompson, and A. Gade, "Suitable Reverberation Times for Rock and Pop Music," *Journal of the Acoustic Society of America* (JASA), 127(1) January 2010.

[6] N.W. Adelman-Larsen, C-H Cheong, and B. Støfringsdal, "Investigation of Acceptable Reverberation Time at Various Frequency Bands in Halls that Present Amplified Music," *Applied Acoustics*, 129 (2018) 104-107.

[7] L. J. Sivian, H. K. Dunn, and S. D. White, "Absolute Amplitudes and Spectra of Certain Musical Instruments and Orchestras," *Journal of the Acoustic Society of America* (JASA), 2 (1931) (330-371).

[8] H. F. Olson, *Music, Physics, and Engineering*, Dover, New York, 1967, p. 205.

Modern PC Board Design Practices

Part 1 — An Overview of the Process

In this two-part series, Ethan Winer examines current practices and methods for designing four-layer printed circuit boards using surface-mounted components. Part 1 presents an overview of the process, while Part 2 will explain specific strategies. Many related issues will be described along the way.

By **Ethan Winer**




Photo 1: Before printed circuit boards, components and connecting wires were all hand soldered one by one. (Photo credit: Erik von Seggern)

Printed circuit boards (PCBs) first appeared in the 1930s, but they didn't become popular until the 1950s when transistors quickly began to replace tubes. Before PCBs, all the components and wires in electronic devices were manually placed and hand soldered (**Photo 1**). This is not only tedious and labor intensive, it's also prone to wiring errors. Today, many electronic devices are mass produced on miniature circuit boards using parts so tiny they must be placed and soldered by precision machines. But hobbyists and smaller professional companies can still take advantage of modern PCBs.




Photo 2: Shown left to right are a 0603 ceramic capacitor, a 0603 tantalum capacitor with J leads, a SOT363 dual transistor, a TO-236-AB transistor, a 1206 film capacitor, a 14TSSOP quad FET op-amp, and an old school DIP package 741 op-amp.

I learned to design and make PCBs in the 1970s while working for a high-tech firm. In those days we applied a special black paper tape, not unlike masking tape, to clear acetate to draw the circuit patterns. A piece of graph paper taped to the desk under the plastic helped keep the traces straight and orderly. This tape comes in various widths to make traces as wide or narrow as needed, though we taped the board layouts at double size for convenience and better accuracy. Round pads with holes for the electronic parts on the board using the same type of tape was available in various sizes. If you found an error or needed to move a group of traces, you had to lift the tape and pads, then do that section all over again.

Once the board layout was complete, the taped acetate went to a photo studio that created a negative at the correct size. We dipped blank copper-plated hard plastic boards into a bath of photo-resist liquid and hung them on edge to dry. Next, the photo negative was taped to a coated blank board, then exposed to ultraviolet light. Since this is a black negative image with clear slots where the traces will be, light passing through the slots cured those portions of photo-resist that would become the traces. The board was then dipped in acid to dissolve all the copper except where the traces were protected by the cured photo-resist. Finally, all the little pad holes for component leads had to be hand drilled in a miniature drill press.

When I resumed my interest in electronic design a few years ago, I used a pair of free programs from PCB maker ExpressPCB to draw schematics and lay out the boards. This software is very popular and extremely capable, though its internal file format is proprietary, requiring ExpressPCB to make the boards. I recently switched to KiCad, a newer free program that's even more capable and also very popular. KiCad outputs standard Gerber files that can be used by any PCB vendor, such as AllPCB in China that I now use. That said, the tips and methods presented in this article are universal, and apply to any PC board layout software.

SMD Components

Old timers like me grew up using through-hole components and DIP integrated circuits, but many designs today use surface-mounted devices (SMDs). Instead of wire leads that go through holes to traces on the bottom of the board, SMD parts are soldered to flat pads on top of the board. Most SMD parts cost less than the larger older types, and many new parts are available only as SMDs, especially semiconductors. However, some through-hole parts may also be needed, such as large electrolytic power supply filter caps and mechanical pots and switches. But SMD trim pots are available and can be amazingly small, such as the Bournes TC33X series that are less than 4 mm (0.15") in diameter. When possible, it's best to put pots, switches, and connectors on the board to minimize hand wiring.

The first project I did using SMD components was the Null Tester featured in the June 2019 issue of *audioXpress*. I used mainly 0805 size parts because those are small enough to minimize the board size, but large enough to easily solder by hand. You can read about that in my Null Tester article linked under Resources. The size numbers for SMD parts are in 100ths of an inch, so 0805 means the part is 0.08" long by 0.05" wide. There are metric equivalents in 10th millimeters for each size, so 0805 is 2012 metric. For hand soldering, 0805 or larger is preferred, though it's possible to solder 0603 parts by hand. Of course, you need fine tweezers, a soldering iron with a very small tip, and thin 0.025 or 0.020 solder. It also helps to have a lighted magnifier and a PC board vice to hold the board steady as you work.

My current project shown in this article has nearly 200 parts on a board that needs to fit inside a microphone case. So I used mainly 0603 resistors and capacitors, and similarly tiny sizes for the semiconductors. **Photo 2** shows an assortment of the SMD parts I used, with an adjacent ruler to illustrate how incredibly small they really are. An 8-pin DIP IC

is also shown for comparison. The 0603 capacitor, second from left and shown upside down, has "J" leads that wrap around under the part. This takes up even less room on the PC board, though having the pins under the part makes hand soldering much more difficult. Many other SMD parts are available with J leads. Even smaller part sizes are available when the boards will be assembled by machine.



Photo 3: Black belt engineer Julia Truchsess uses a squeegee to apply solder paste to the PC board while the board and template are held in a special jig.




Photo 4: Julia's pick and place machine lets her precisely place teeny SMD components onto the even smaller solder pads of the PC board.



Photo 5: This oven from AliExpress lets you program a sequence of temperature changes to rise, sustain, rise again, then fall at rates matching the solder paste's specs.

Figure 1: This is just a small part of my project originally drawn in KiCad’s schematic editor.



Resistors can be as small as 01005, and an 8-pin DSBGA IC is less than 2mm long!

Rather than try to solder these tiny parts by hand, we used a solder mask template to apply solder paste to the component pads. This template, also called a stencil, is a thin piece of metal with slots that align with all of the pads that require solder. Once the board and template are aligned in a special jig, the paste is applied with a squeegee. This way only the component pads receive the paste. Then after all of the parts have been placed onto their coated pads, the board is baked in a dedicated oven that’s programmed for a temperature sequence matching the solder paste’s requirements. You can see my friend Julia Truchsess apply the paste in **Photo 3**. **Photo 4** shows her Abacom Technologies

pick and place machine (see Sources). This device has a suction arm that grabs each tiny component, then lets you rotate and position the part precisely to drop it onto the board. A foot pedal turns the suction on and off.

Photo 5 shows the oven used to flow the solder paste onto the boards (see Sources). Even if the parts are slightly misaligned on the board, surface tension causes them to slide into place when the paste liquefies in the oven. It’s a good idea to pay a little extra for gold plating—called Electroless Nickel Immersion Gold (ENIG)—on the component pads. When the solder paste is applied to the board through the template slots, you can more easily confirm that the pads are fully covered because the gold color changes to silver.




Figure 2: KiCad’s Choose Symbol dialog contains pre-drawn symbols for thousands of components.

First Things First

Before designing the PC board you need to draw the schematic in the PC board’s companion program. The Audio Input section of my schematic is shown in **Figure 1**. Both KiCad and ExpressPCB link the connections on your schematic to traces on the board. This not only confirms that your board matches the schematic, it can even guide you as you place the traces between parts by highlighting the destination pads. KiCad takes this further by letting you copy the schematic onto the board layout, then it shows a series of lines, called a “rats nest,” of all the board traces needed. As each trace is drawn and connected, that rats nest line disappears. If you remove a trace, the line re-appears so you can’t forget to connect it.

In KiCad, selecting a part in the schematic editor highlights the same part in the PCB program, and

vice versa. This makes it easy to see where each part goes when you place them onto the board. Even better, if you use the optional plug-in to create an HTML Bill of Materials (BOM), it highlights all the parts having the same value. So when you select 100kΩ resistors, the board view shows all of them at once. You may need to flip two-wire parts around on the board layout so their orientation aligns with the schematic. Not only polarized parts like electrolytic capacitors and diodes, but resistors and non-polar capacitors also have a Pin 1 and Pin 2 that must connect correctly in order to match the schematic and avoid warning errors.

In ExpressPCB, you select a component’s physical size and shape—its footprint—when laying out the PC board. But in KiCad footprints are assigned in the schematic editor. This way, when you copy the schematic to the board, the correct footprint is already known and applied.

To insert a component into the schematic you select Add a Symbol, which brings up the Choose Symbol dialog shown in **Figure 2**. There are thousands of symbols from which to choose, so the Filter dialog at the top lets you pare that down by typing res (resistor) or opa (op-amp) or 74xx to show a list of all the known chips in that logic series. You can’t change the pre-defined library symbols, but you can copy, then modify, existing symbols, or create new symbols.

Next you double-click the symbol to bring up the Symbol Properties shown in **Figure 3**. This is where you assign the component’s number, value, and the footprint that defines the pad size and shape on the PC board. You can also include the URL for the part’s datasheet. That is then included when KiCad generates a BOM for your project, which is needed if your boards will be assembled by an outside company.

When you click the Footprint field, the dialog shown in **Figure 4** is displayed, and an image of the footprint is also shown for confirmation. Most of the resistors in this project are 0603 (1608 Metric). After the first resistor or capacitor is added to the circuit and its footprint defined, you might find it easier to copy that part repeatedly rather than bother with Add Symbol each time. When you copy a component, its part number is replaced with a question mark (i.e., R?). Then when you double-click that part KiCad opens the Symbol Properties dialog for you to manually enter the new part number (e.g., R12).

Unless your device is fairly simple, a four-layer board makes sense because it offers ground and power planes in the middle. This eliminates capacitive coupling between the top and bottom




Figure 3: All of the attributes for a component are assigned in KiCad’s Symbol Properties dialog, including the physical footprint on the board.

signal traces, and helps to simplify routing. When you’re unable to route a particular trace on top because it would have to cross other traces, you’ll insert a via to pass that connection through to the bottom layer, then come back up on the other side. It also creates a sort of continuous bypass capacitor between the power and ground layers. The top layer usually holds most of the component pads and connecting traces, with the ground layer directly underneath. Then the power layer, with the fourth layer available for more connecting traces at the bottom and/or a second power source.




Figure 4: KiCad’s Footprint Library Browser contains pre-defined footprints for thousands of component sizes and shapes .

Figure 5: The Null Tester board's power layer is divided into "fingers" that carry both V+ and V- around the board on the same layer.





Figure 6: This board is about eight square inches, shaped to fit inside a microphone case.

About the Author

Ethan Winer has been an audio engineer and professional musician for more than 45 years. His Cello Rondo music video has received nearly 2 million views on YouTube and other websites, and his book *The Audio Expert* published by Focal Press, now in its second edition, is available at amazon.com and Ethan's own website ethanwiner.com. Ethan is also a principal at RealTraps, a US manufacturer of high-quality acoustic treatment.



Resources

AllPCB | www.allpcb.com

ExpressPCB | www.expresspcb.com

KiCad | www.kicad.org

Pragmatic Designs, Inc., Julia Truchsess' website, <https://pragmaticdesigns.com>

E. Winer, "You Can DIY! Building a Null Tester Device," *audioXpress*, June 2019, <https://audioxpress.com/article/you-can-diy-building-a-null-tester-device>

Sources

Manual SMT Pick and Place Machine

Abacom Technologies | <https://abacom-tech.com/product/manual-smt-pick-and-place-machine-ezpick>

Benchtop Reflow Soldering Oven

AliExpress | www.aliexpress.com/item/2255799842971646.html

If many traces are needed on the bottom layer, and the circuit requires both plus and minus power, you can split the power layer to carry both V+ and V- as I did for my Null Tester board shown in **Figure 5**. As with most ICs, the plus and minus power pins for op-amps on the top layer are on opposite sides of the chip. I drew one "phantom" chip here as it would appear on the top, just to show how the chips can straddle gaps. Note that you don't generally draw traces directly on the inner layers. But you can draw gaps, called "splits," as shown here. Otherwise, you can use the power layer for V+ and the bottom layer for V- plus any additional traces needed.

Note that PCBs don't need to be rectangular. Most board manufacturers let you create round boards, or nearly any shape needed for your application. Mounting holes of any size can also be added to the board. **Figure 6** shows the board for my current project, which is designed to fit inside a microphone case. All of the parts except one are on the top board layer and clearly identified as R12, C14, and so forth.

KiCad includes the 3D viewer shown in **Figure 7** that lets you see a virtual photograph of the finished board. When using a separate CAD program to design an enclosure, this helps verify that the finished board will fit as intended. The pins sticking up from the bottom of the board near the middle left are for a 1" diameter dynamic loudspeaker used here as a microphone, mounted on the bottom. Notice the U-shaped test points scattered around the board—four are at the far right near two large capacitors and the XLR connector. These let you connect an oscilloscope or voltmeter probe to various parts of the circuit for troubleshooting. A final board design probably won't need as many as used here.

Avoiding Errors

Both KiCad and ExpressPCB let you confirm that your board layout matches the schematic. Besides guiding you when connecting components, ExpressPCB lets you highlight related connections. With that tool selected, when you click on a component pad, all of the other pads that need to be connected are highlighted in blue. However, it doesn't flag existing connections that are incorrect, or prevent you from connecting a via on top to more than one other layer. KiCad is much better in this regard. The schematic editor's Electrical Rules Check provides warnings, such as parts that are not connected to anything, and also fatal errors, such as IC outputs that are connected together or to ground. KiCad's PCB editor has its own error reporting module that displays warnings (e.g.,

silkscreen text items that overlap) and hard errors (e.g., short circuits between component pads or traces).

After a circuit is designed, I always buy all the parts before I even start the board layout. Then I can print the board layout in progress on paper to confirm that the part outlines and pads are in fact the correct size for what I received. Many SMD components are available in a bewildering number of minor variations, such as the time I received ICs much smaller than I thought I had ordered. Another time I accidentally bought a “wide” 1206 half-watt resistor that probably should have been listed as an 0612.

And Finally

Even if you need only one board, order at least two so you’ll have a blank board to examine and test without resistors and other parts blocking any labels or covering any traces. Some PCB manufacturers have minimums, such as five boards from AllPCB, which is the company that I use. But the cost for five boards is barely more than it is for one.

This article has presented an overview of the methods and procedures used to create modern

PC boards. It also compared the features of two popular freeware programs for drawing schematics and laying out boards. Next month, Part 2 will go into greater detail, and present specific techniques to help you design PC boards that work correctly and are efficient. 




Figure 7: KiCad’s 3D Viewer lets you rotate the board to see the top and bottom from all angles.

Ribbon Loudspeakers
what you
need to
know.

Ribbon Loudspeakers
Theory and Construction
By Arthur V. D'Angelo, Ph.D.

Order your copy
cc-webshop.com

SpiralAirtm
Interconnects

*Coax
Is Cool
Again!*

Cardas Golden Strand
Air Filled Cotton Dielectric
Unique Spiral Litz Shield
ETI Tellurium Copper Connectors

Affordable Litz Interconnects and Bulk Wire
Tonearm - Phono - Line - Loudspeaker

KAB
Preserving The Sounds
Of A Lifetime
www.kabusa.com

Power Transformer Parameters, Selection, and Testing

Part 4 — Parameters and Their Losses

The initial group of articles focused on the history of transformer cores, their construction methods, their materials, and various testing methods. Part 4 of this series discusses the parameters of transformers and their losses.

By
Chuck Hansen

In its simplest form, a transformer is an electrical device which, by mutual electromagnetic induction, transfers electrical energy from one isolated circuit to another by means of coils called windings. The two windings are wound on a common iron core. **Figure 7** shows the core and coils of a conventional real transformer, along with its equivalent schematic. If there are multiple secondary and/or primary windings, they are also wound on the same iron core.

As a result, additional primary voltage would be required to establish the desired full-load secondary voltage, E_S . Since you can't change the line voltage from the power company, this is normally accomplished by reducing the primary turns during the design phase, so the secondary turns ratio appears larger. This higher secondary voltage appears during no load operation where the losses are lower.

AC current in the primary winding provides the time-varying magnetomotive force (mmf), which creates an AC magnetic flux in the core. The iron core provides a high-permeability magnetic path that ensures that a high percentage of the flux will be linked to the secondary winding. The voltage induced in the secondary winding and delivered to the load by the AC core flux variations is given by Faraday's law:

$$e = k \times N \times \frac{d\Phi}{dt}$$

In other words, the voltage induced in the secondary winding is proportional to the number of turns (N) and the time rate of change of flux ($d\Phi/dt$). The voltage can easily be changed from one level to another, lower or higher, by proper selection of the number of primary and secondary turns. At a lower frequency, more turns and/or a larger core area are required to develop a given flux density than at a higher frequency.

In terms of real-world transformer design, we still employ Michael Faraday's 1830 law:

$$e = 2.865 \times N \times B \times f \times A \times 10^{-4}$$

where the voltage developed in a winding is proportional to the number of turns (N), the flux density in kilogauss (B), the AC frequency (f), and the core cross-sectional area in square inches (A). This is the formula we generally use in aerospace because many core catalogs still give core dimensions in inches, and core loss in watts/pound. The numeric constant is 4.44 for metric dimensions.

When voltage is applied to the left, or primary winding (V_p), the largest part of the current (I_p) goes toward creating the mutual core flux. But not all of this current produces useful output volt-amperes (VA) at the secondary. The remaining portion of the primary current is required to provide the core exciting current, I_E .

Exciting current (I_E) is the penalty for the significant advantages of using high permeability

transformer cores. The exciting current is divided into magnetizing current (I_M) and core loss current (I_C).

These loss components cause a voltage division to occur due to the voltage drop in the primary winding impedance, Z_p , composed of resistance R_p and inductance X_p . This reduces the voltage (E_p) available at the input to the ideal (no loss) transformer portion of the equivalent circuit and reduces the useful primary current (I_1). This is why transformers can never be 100% efficient.

Magnetizing current is the primary current that flows when the secondary is open-circuit and has no load, and is determined by:

$$I_M = 0.794 \times H \times \frac{Ml}{N}$$

where the magnetizing current is proportional to the mmf (H) times the mean magnetic path length of the core, Ml , divided by the number of turns N . Transformer designers attempt to reduce the magnetizing current by minimizing the magnetic path length. This allows the fewest turns for a given primary voltage and has the added benefit of reducing the leakage flux.

The polarity dot marks in the coils are the instantaneous polarities of the primary and secondary waveforms of the AC transformer. There are three equivalent ways of defining the dot notations:

- Voltages induced in the two windings due to changes in the mutual flux will have the same relative polarity at the dot-marked terminals.
- If positive currents flow into the dot-marked terminals, the flux mmf produced in the two windings will have additive polarity.
- If one winding is open, and if the current flowing into the dot-marked terminal of the

other winding has a positive rate of change, the voltage induced in the open winding will also be positive at the dot-marked terminal.

The leads connected to transformer windings are usually color-coded and a polarity schematic is either printed on the case or furnished separately in the package. Some smaller transformers have PC board pins and they are arranged such that the transformer can only be inserted in one orientation.

Transformer Energy Losses

Energy loss in a transformer is almost totally converted to heat, and consists of eight components:

- Hysteresis losses (I_{HYST}): Each time the magnetic field is reversed, a small amount of energy is lost due to hysteresis within the core (think of it as “magnetic friction”). For a given core material, the loss is proportional to the frequency, and is a function of the peak flux density to which it is subjected.
- Eddy currents (I_{EDDY}): Ferromagnetic materials are also good electrical conductors, and a solid core made from such a material would constitute a single turn low impedance electrical path throughout its entire length. Eddy currents circulate within the core in a plane perpendicular to the flux, and are responsible for resistive heating of the core material. The eddy current loss is a complex function of the square of the ac supply frequency and inverse square of the material thickness.
- Excess eddy currents (I_{EX}): Non-sinusoidal, non-uniform and non-repetitive magnetic domain motion in the core material contributes to excess eddy currents. (ieeexplore.ieee.org/document/5324814)




Figure 7: The equivalent circuit for a real transformer (Image Source: Static Electromagnetic Devices, Hunt and Stein; Allyn and Bacon Inc., Boston 1970)

- Stray losses: Leakage flux, especially across core air gaps in laminated transformers, can intercept nearby conductive materials such as the transformer's support structure, and give rise to eddy currents that are converted to heat.
- Winding resistance: Current flowing through the windings causes resistive heating in the copper wires, and the winding resistance increases with higher temperatures. At higher frequencies, skin effect and proximity effect further increase the effective winding resistance and losses. The wire diameter should not exceed twice the skin depth in mils at the operating frequency.
- Dielectric loss: Dielectric loss occurs in the materials used to insulate the core and windings, and the capacitive coupling in the insulation system to the steel core and to other conductors used in the transformer support structure.
- Magnetostriction: Magnetic flux in a ferromagnetic core causes it to physically expand and contract slightly with each cycle of the magnetic field. This produces the buzzing sound at twice line frequency commonly associated with transformers, and in turn causes losses due to frictional heating.
- Mechanical losses: In addition to magnetostriction, the alternating magnetic field causes fluctuating electromagnetic forces between the primary and secondary windings. These create vibrations within the structure,

adding to the buzzing noise, and consuming a small amount of power.

Core Losses

The core loss current, I_C , is divided into hysteresis (I_{HYST}), eddy current (I_{EDDY}), and excess (I_{EX}) losses. The total core losses are represented by:

$$P_T = P_{HYST} + P_{EDDY} + P_{EX}$$

Hysteresis Losses

When an AC current is applied to a coil surrounding a magnetic core, it generates a magnetic flux within the core. The elementary dipole magnetic domains within the core try to align themselves with the continuously reversing magnetic field. Current is the forcing function for magnetizing force (H). There is a lag between the magnetizing force and the induced flux density (B), see **Figure 8**. This results in an energy loss within the core called hysteresis loss, P_{HYST} , which is proportional to frequency, the peak flux density, and the area under the B-H curve for the particular magnetic material.

$$P_{HYST} = \frac{(2 \times S \times Bp^2 \times f)}{\mu}$$

where S = hysteresis curve shape factor, Bp = peak sinusoidal magnetic flux density, f = frequency (Hz), and μ = permeability of the core material.

The hysteresis loop shows the nonlinear relationship between B and H, and key points along the loop vectors. Segment 1 is the normal initial magnetizing curve where B and H start at zero then progress toward core saturation at B_{max} and H_{max} . This initial curve is called the First Quadrant and is used to define the B-H characteristics of the magnetic material (see Figure 5 in Part 2, *audioXpress*, December 2022).

As the AC magnetizing current begins to decrease horizontally along curve segment 2, the flux density lags behind. When H reaches the initial starting point at $H = 0$, the flux density B still has a non-zero value, at point B_r along the vertical axis. Not all of the magnetic domains return to their original random orientations. B_r is the residual magnetic flux (retentivity or remanence).

In order for B to return to zero, H has to change polarity along the horizontal axis until it reaches value $-H_c$ where the flux density reaches zero. H_c is defined as the coercive force. $-H_c$ will continue to increase until $-B$ approaches negative saturation at $-H_{max}$ and $-B_{max}$.




Figure 8: Hysteresis loop for a soft magnetic material (Image Source: *Alternating Current Machines*, 5th Edition, MG Say, Section 2, John Wiley & Son, 1989).

Now B-H curve segment 3 will trace a mirror image of segment 2 with respect to the intersection of the B and H axes. When H reaches zero, the flux density will be non-zero at point -Br. Then when the magnetizing force reaches -Hc, the flux density will be zero again.

Hysteresis loss is proportional to the area under the hysteresis loop for each core material (**Figure 9**), and it is the power required to magnetize the core and produce the AC flux.

The core consists of tiny magnetic dipoles, or domains. These domains align themselves and rotate with the AC magnetic field. The hysteresis losses can be thought of as the friction caused by these rotating magnetic domains. Soft magnetic materials, such as silicon steel, have a narrow hysteresis loop with a small area, and are easily magnetized and demagnetized. Hard magnetic materials, such as permanent magnets, have a wide hysteresis loop with a large area, and take a large coercive force to magnetize, and are more difficult to demagnetize. Hysteresis loss is proportional to the flux density, the magnetic core volume and the frequency.

Eddy Current Losses

Eddy current loss is caused by the current that wants to flow across the cross-section of the laminations, perpendicular to and opposing the primary flux in the core, and results in increased temperature rise due to power wasted by current flow in the conductive iron core (**Figure 10**).

Current is induced in any conductor that is perpendicular to the flux path. This includes the windings and the area of the iron core. If the core was solid iron, it would act like a shorted turn. High current would circulate at the core surface. The amount of current depends on the resistivity of the core.

For this reason, transformer cores—rather than being a single mass of electrical steel—are laminated into thin, electrically insulated steel strips whose long dimension is parallel to the core flux. The thinner the lamination, the shorter the path available for the eddy currents, and the lower the eddy current losses.

Eddy current loss per unit volume is proportional to the squares of the core material thickness t , the peak flux density B_{max} , and the AC frequency f , and inversely proportional to the resistivity ρ , of the core material. It causes resistive heating in the core material.

$$P_{EDDY} = \frac{(\pi \times t \times B_{max} \times f)^2}{6 \times \rho}$$

where t = core tape thickness, B_{max} = peak sinusoidal magnetic flux density, f = frequency (Hz), and ρ = resistivity of the core material.

Using thin oxide-insulated lamination thickness results in a squared reduction in the eddy current losses. Lamination thickness is given in mm or mils (0.001"). The thinner the lamination, the shorter the path available for the eddy currents, and the lower the eddy current losses. In addition, silicon is added to the iron to reduce its conductivity and




Figure 9: Hysteresis loops for soft and hard magnetic materials (Image Source: Electronics Tutorials)




Figure 10: Eddy currents in solid and laminated ferromagnetic materials

About the Author

Chuck Hansen is an Electrical Engineer and holds five patents in his field of electric power engineering. He has worked in both the electric utility and aerospace electric power industries. He was assigned as the Instrument and Controls Startup Group Lead Engineer for two nuclear power generating units. He was the Electrical Systems and Controls Supervisor for a number of military programs as well as commercial and business jet electric power programs. He has written two books for Audio Amateur publications, and more than 260 magazine articles. He began building vacuum-tube audio equipment in college, and enjoys restoring and modifying audio equipment and test equipment. In his most recent roles as Senior Design Engineer and Lab Calibration Manager, he designed specialized test equipment particular to the aerospace electric power sector.



References

[1] "Anomalous Loss Hysteresis Loop," ResearchGate, March 2014, www.researchgate.net/publication/273943168_

Resources

Electronics Tutorials, www.electronics-tutorials.ws/electromagnetism/magnetic-hysteresis.html

Table 2: Recommended core thicknesses for various operating frequencies, and the core stacking factors (Image Source: Arnold, Magnetics Inc., and Magnetic Metals core catalogs)

Tape		ISIA GOSS	ISIA NGO	Stacking Factor		Recommended Fmax
14 mils	0.36mm	M6	M15, M19	0.95 SiFe	0.95 NiFe	50Hz
12 mils	0.30mm	M5		0.945 SiFe	0.95 NiFe	50Hz, 60Hz
6 mils	0.15mm			0.915 SiFe		400Hz
4 mils	0.10mm			0.90 SiFe	0.90 NiFe	400Hz to 1,000Hz
2 mils	0.051mm			0.85 SiFe	0.85 NiFe	2,000Hz
1 mils	0.025mm			0.75 SiFe	0.75 NiFe	5,000Hz
0.5 mils	0.013mm				0.50 NiFe	10,000Hz

further reduce the eddy currents. 3% silicon steel has a resistivity of $48 \mu\Omega\cdot\text{cm}$, vs. $9.7 \mu\Omega\cdot\text{cm}$ for iron.

There is an eddy current increase at the ends of the transformer windings, especially for E-I and R-type laminated cores and tape-wound C- and -E cores, where the core is exposed beyond the ends of the windings. This eddy current increase is due to the crowding effect of the magnetic field leakage at the exposed ends of the windings. Toroidal and O core transformers can be designed with continuous overlapping windings such that the core is not exposed.

Transformer core lamination thicknesses are available from 20 mils down to 1 mil for silicon-steel, from 14 mils down to 0.125 mils (and less on special order) for nickel alloys, and even thinner laminations result in less effective core area for the same physical cross-sectional area. This is because the thinner laminations have more total oxide film thickness for a given core area, since each lamination is insulated on both sides. This

results in more individual oxide films across the total core thickness. This reduction in net metallic cross-sectional area is called the stacking factor (SF), as shown in **Table 2**.

The 50Hz transformers (and other electromagnetic devices) use 14-mil thick laminations. The 60Hz devices use 14-mil or 12-mil laminations. However, the 400Hz transformers use 4-mil laminations, which produce less eddy current losses (**Figure 11**). Even higher frequencies can be used with thinner laminations, we use 2-mil up to 2kHz and 1-mil up to 5kHz. At even higher frequencies we would use powder metal cores such as molybdenum permalloy (MPP, 81Ni-17Fe-2Mo), softer magnetic materials or ferrite cores, all as dictated by a SWaP-c analysis (low size, weight, and power, with as low a cost and still meets the requirements).

Excess (Anomalous) Losses

Another type of core loss is excess (anomalous) eddy current loss, P_{EX} , which has to do with a non-sinusoidal, non-uniform and non-repetitive domain wall movement in the core material associated with magnetic domain structure [1]. It is not usually specified for the small types of C, E, and toroidal cores used in commercial and aerospace applications, but it is a very big deal in terms of material cost vs. power loss over the installed life for large utility power transformers.

A utility transformer dissipates a constant no-load magnetizing current loss as long as it is energized, 24 hours a day, for all conditions of loading. This power loss represents a cost to the user during the lifetime of the transformer. Initial no-load loss measurements are made by the manufacturer to verify that the no-load loss does not exceed the specified or guaranteed value.

Next Month

Next month, in Part 5 we will continue to discuss electrical and electromagnetic losses in transformers. 




Figure 11: Typical core loss curves vs. frequency for 4-mil and 12-mil GOSS (Image Source: Microsil, Magnetic Metals Tape Wound Core Design Manual C1-171-20M, pg. 27)

Cine Gear EXPO 2023

New York
March 10-11

Los Angeles
June 1-4

Atlanta
October 6-7



cinegearexpo.com

The Pre-5

A High-Performance Tube Hybrid Preamp



Photo 1: The Pre-5 is a tube-based preamp with a 10Hz to 100kHz frequency bandwidth.

In a never-ending search to improve the performance of tube-based audio circuits, this project, which the author calls the Pre-5, proposes pushing the envelope to achieve even lower distortion from an already great-sounding concept.

By
James Skov

The stimulus for this project started, maybe 10 years ago when I bought and built a TU-8500 preamp from Elekit, a kit that I still use today. After doing a SPICE simulation of the main 12AU7 circuit and measuring its performance with my APx 515 I felt that the tube section could be improved. Don't misunderstand me, the TU-8500 does an amazing job given everything it does and its price point. But I was challenged to do better. Before I get too far into the details, I wish to thank John Broskie for his *Tube CAD Journal* website of which I reference constantly. You can find the ideas for this circuit on his website, although this complete circuit is not there.

I remember reading this statement years ago, I don't know where, "The way to remove distortion is to keep the current constant through the device and the voltage constant across the device, then you will achieve low distortion." Yes, if you have no current change and no voltage change then you also have no gain, you have a wire, so of course distortion is low. But still that concept plays a key part in this design. Let's get started.

I had several basic requirements for my Pre-5 tube hybrid preamp (**Photo 1**). It had to be tube-based with a 10Hz to 100kHz frequency bandwidth. The output needed to be 2Vrms. The total harmonic distortion (THD) needed to be: <0.001%, 1kHz, 1Vrms, full bandwidth, and no A weighting. I wanted the THD+N to be: < 0.005%, 1kHz, 1Vrms, full bandwidth, no A weighting with a 10dB gain. For mute, I wanted a 30-second power up, auto mute on power loss.

Figure 1 shows the schematic of a hybrid preamp that easily meets these requirements. (You are free to use this for personal DIY projects, but for production projects please contact me.)

The Circuit

Now let's breakdown the circuit, what it does, and some of the choices I made. The basic topology is a difference amp with the tube grid being one input and the base of the PNP transistor being the other input. You can see that the tube is surrounded on the top and the bottom by solid-state devices. This helps to implement the "no voltage change" part of the low distortion statement. Because of the cascode, the anode can't move, "no voltage change," Miller effects are also minimized.

Above the cascode is a constant current source (CCS), set for approximately 4mA. I chose 4mA based on several observations that for a 12AU7, the 3mA to 4mA range can actually have lower distortion than a higher current value. I have measured results that confirm that 3mA to 4mA is better than say 10mA, as it completely depends on the circuit topology and use case. The 4mA current source manages the "no current change" part of the low distortion statement. The voltage across the tube changes in the range of the grid input voltage * 0.96, which is a small change. The tube is in complete control of the circuit but is placed in an environment that minimizes distortion, and maximizes gain and frequency bandwidth.

The PNP at the cathode also helps to minimize any noise contribution. This is the same idea that

is often seen in Recording Industry Association of America (RIAA) preamps, where the cathode is biased by an LED instead of a resistor with a bypass capacitor, lower noise, and higher gain. The choice of the IXCP10M45S versus the DM2540 came down to consistent distortion performance. I have found that both devices are extremely useful, but you need to check which device performs better in your specific application.

With all the device shortages lately, I have also used the IXCP10M90S and the DM2540 as the cascode device, but I prefer the IXCP10M45/90S. Why not an NPN device? Again distortion, I tried the MJE340 since I am using the MJE350, I also tried the MJE15034 but didn't like the results in this specific application. The choice of the KSA992 comes down to low noise, high voltage, and stable gain. It was the best I could find using a device search and my criteria.

The next thing to note is that the bias for the anode CCS is itself coming from a CCS, which is generating the reference voltage for the cascode device. This CCS is set to approximately 1.1mA ($0.0677/61.9 = 1.09\text{mA}$). There is a 50kΩ trim pot in this path, which allows for the adjustment of the anode voltage. The reason for this is that tube halves are never equal and the goal is to generate 4.85V at the cathode of the tube to set the current in the output leg to approximately 11mA.

From measuring different tubes, I have found the anode voltage value will range from 110V to 136V. The 50kΩ pot along with the fixed 80.6kΩ

resistor allows for a range of 88V to 143V. The anode voltage directly controls the amount of current in the output branch and should be set such that the voltage measured at the 1kΩ output resistor is approximately 15.2V (10.9mA).

As a side note, the initial voltage should be set to 14.5V and then wait 15 minutes. The LM334 is temperature sensitive and will drift as it is warmed by the tube, and the tube will also drift as it warms. After the 15-minute period make one more adjustment to 15.2V. This setting is on the high side of what is required to meet performance so it will allow for tube aging. The output current will slowly drop as the tube ages and yet still provide excellent performance. It is not something you need to constantly check or worry about.

Why use a CCS and not a resistor in this branch? The answer is noise. I found that, with an early prototype, the noise floor was lowered by almost 10dB when I switched from a fixed resistor to the current source in this branch. I think the reason is power supply noise transfer.

If you study the circuit, you will find that the top section is referenced to 180V and the bottom section is referenced to ground. Between these two sections are three current sources, input constant current, bias constant current, and output current. The three current sources allow the top and bottom sections to interact without transferring very much of the noise between them.

The power supply rejection ratio (PSRR) of this preamp is about -75dB, not amazing, but not bad.




Figure 1: This is a schematic of a hybrid preamp that easily meets these requirements

About the Author

James B. V. Skov recalls as a 2-year-old sitting in front of his parents' phonograph and listening to classical music on 78s for hours. As a high school and university student (Electrical Engineering), he went to swap meets to buy cheap tube suitcase record players, tube audio electronics, and TVs to fix them for resale or personal use. For employment, James has worked at Rockwell Semiconductor, Conexant, Broadcom, Cypress, and Infineon. Always with some involvement in audio chip and/or software design. James has always loved tubes and has a small but sophisticated audio lab at home. Currently retired, he is back to doing what he loves full time—tube audio projects.



I've included a graph (Figure 21) of the PSRR at the end of this article.

The Output

Next is the output section. It is basically a current output converted to a voltage through the $1k\Omega$ and 392Ω resistors. With the resistor values defining the cathode voltage and the final gain. $1K$ divided by $392 = 2.55$, plus 1 gives a rough gain of 3.55. The actual gain is around 3.33 or 10dB. With an output current of 10.9mA the 392Ω resistor will have 4.27V across it. Then adding the base emitter junction voltage gets us to the cathode voltage of 4.85V. I picked the target idle current of 10mA in order to be able to drive output cable capacitance and then adjusted up from there to get better high-frequency distortion results.

Photo 2 shows an inside look at the PCB. The reason for the PMOS/PNP pair is simple. The PMOS (P-channel metal-oxide-semiconductor) provides almost no load on the input section (input capacitance only). It also gives the input CCS some voltage with which to work. The PMOS/ PNP combination has a turn on threshold of 4.5V.

I couldn't find any high turn on threshold PMOS device that could meet my distortion goal by itself. The solution was to keep the low load with the PMOS device and drive a PNP device to get the higher current and lower distortion. I'm sure that there are other device pairs that would work just as well. The key is the high turn on threshold and the low input capacitance. **Figure 2** shows a top view (a) and bottom view (b) of the preamp's PCB.

For the PNP, it should be an audio/video driver with a stable gain. The MJE15035 and the MJE15033 are possible TO220 device examples but have a 2MHz




Photo 2: Here is a look at the PCB, which measures 5.2" x 4.0".




Figure 2: This is the top view (a) and bottom view (b) of the preamp's PCB.

MENLO SCIENTIFIC

Engineering
Success
since 1983

Engineering Consulting for the Audio Industry

We Know

- Product Development
- We Provide
- Design and Technical Support
- We Deliver
- Liaison, Sourcing, and Guidance

We've worked with the world's largest brands, supported large companies outsourcing production, advised manufacturing facilities on streamlining their process, and helped companies large and small to develop new products, and launch revolutionary technologies. We've helped startups find investors, and introduced entirely new concepts. Many companies learn their lesson by failing. Others have hired us first.

MENLO
SCIENTIFIC

Richmond, California, USA

- tel 510.758.9014
- mike@menloscientific.com
- www.menloscientific.com



bandwidth instead of the 3MHz bandwidth of the MJE350. The 100Ω output resistor is a capacitive cable decoupling agent and then the 10μf is to block the DC offset of the preamp. The 15pf and 56pf capacitors keeps things stable. The entire preamp is DC coupled except for the final output capacitor. This is a feature and a warning, any DC

on the input will shift the bias of the entire circuit. **Figure 3** shows a peak hold FFT capture of a 1kHz to 950kHz sweep. What matters is that the top edge has a smooth drop-off and doesn't rise. You can see roughly a 0.5dB drop from 1kHz to 950kHz. Plenty of bandwidth available. I've added measure points and zoomed in a little to make the frequency response clearer. The frequency response is down 20dB at 8Mhz. Since I am going to include several graphs, I will also include the 20Hz, 1kHz, and 20kHz square waves while I'm at it (**Figure 4**). The output load is 22kΩ, which explains the 20Hz tilt.

Mute Control

Last is the mute control section. Obviously if you have a processor in your preamp driving a display, input selection, and receiving IR remote control commands, then the CPU could handle mute and power loss as well. But what if you don't like or don't want the noise a CPU might add? I have devised a circuit, which is on the preamp PCB and will handle the mute task.

Figure 5 shows a retriggerable one-shot and a time delay made from a dual 555 timer chip (NE556). The one-shot, on the left, has a timeout of roughly three 50Hz/60Hz power cycles. It is retriggered by an opto-isolator driven from the transformer secondary in the power supply. Pin 3 and Pin 4 are shorted together during the cycle peaks. As long




Figure 3: This is a peak hold FFT capture of a 1kHz to 950kHz sweep.




Figure 4: Here are the 20Hz (a), 1kHz (b), and 20kHz (c) square waves. The output load is 22kΩ, which explains the 20Hz tilt.

as power is present, the one-shot is retriggered on every power cycle and won't time out so mute won't be set. If three power cycles are missed, the one-shot will timeout and the preamp output will mute. In the non-triggered state, the output pin is 5V, which then supplies the voltage for the 30-second timer. When power is present and stable the timer will start timing for 30-seconds and then release mute by activating the relay.

The discharge pin of the one-shot drains the capacitor of the timer forcing the 30-timer to do a full restart if there is a power glitch. The 43Ω

resistor is a current limit for the discharge pin. The discharge pin is rated for 200mA but a fully charged capacitor could provide a much higher current for a very short instant. Between the one-shot and the timer is a connector exposing an interface for a mute switch on the front panel. It could also be driven by a processor through an opto-isolator. The resistor and capacitor values were chosen to create a switch debounce time of 100ms for both set and reset to avoid any relay chatter.

The output of the timer drives the mute relay. In the normally closed (mute) state the preamp has




Figure 5: I devised a circuit that is on the preamp PCB and will handle the mute task.




Figure 6: The power supply I designed for this project provides both 12.6V DC and 180V DC. This is the 12.6V supply.




Figure 7: The 180V DC power supply




Figure 8: Here are the top view (a) and the bottom view (b) of the power supply's PCB.

a 47k Ω load to ground, and the preamp output is directly shorted to ground. In the normally open state (un-mute), the output to ground is removed and the 47k Ω resistor is shorted, connecting the preamp output directly to the output jack. What this means is that the relay is powered as long as the preamp is active. Several relay manufacturers warn against this long-term use case because the heat from the relay coil can degrade the contact surface. That is the problem the RC filter between the relay coil and the timer output is meant to address. Most relays require 80% of their rated coil voltage to latch but only 20% of their rated coil voltage to hold in the latched state.

The 165 Ω resistor was chosen to match the 165 Ω coil resistance. Initially the 470 μ f capacitor is discharged and the coil sees the full voltage transition when the timer drives low. The capacitor quickly charges leaving the relay coil with 50% of the coil voltage. This has the effect of latching the relay at full voltage and then lowering the voltage

to 50% for holding. The 50% level is still well above the required holding voltage and will generate less heat, which helps the contacts last longer.

The specific relay used in this design does not have a warning about long-term latched contact degradation, but I still included the holding voltage feature for added longevity. There are connectors for a front panel mute LED and a front panel power LED. A single dual-color LED could also be used to give an indication of both power and mute states. I don't like bright LEDs so I keep the current low, 5mA, with 620 Ω resistors, but these could be changed to whatever value suits your needs. The mute circuit is powered from the 12.6V heater voltage isolated by a 5V regulator.

The Power Supply

The power supply I designed for this project provides both 12.6V DC and 180V DC. I'll start with the 12.6V supply (**Figure 6**). This is a fairly straightforward 18V AC to 12.6V DC power supply. The two RC filters clean up the DC and lower the thermal stress on the regulator. My goal was to have no more than 2.3W on the regulator so that the heatsinks can be smaller. The R26-90 transformer provides 22V AC unloaded or 24.2V DC loaded. The AS-05T200M provides 16.6V AC unloaded or 19.2V DC loaded. My goal is to have 2.4V to 5V across the regulator to limit the regulator heat to 2.25W, assuming a 0.45A load. The regulator requires at least 1.2V to regulate. The two RC filter resistors provide the needed voltage drop and a slower turn on. The only thing to note is the opto-isolator and 7.5k Ω resistor providing power present switching for the relay mute circuit on the preamp PCB.

The 180V DC power supply may need a little more explanation (**Figure 7**). In the bridge the 10 Ω resistors lessen charging current spikes and give the 1nF noise capacitors some extra help. The 100 Ω resistor is part of an RC filter to lower the ripple voltage. Next is the IXTP3N100D2. It is a 3A, 1000V




Photo 3: This is the populated power supply circuit board. This design has proven beneficial as the power related harmonics in the output of the preamp are reduced to almost nothing as you can see in the 1kHz FFT.

The *Loudspeaker Industry Sourcebook 2023* Is Taking Shape...

Act Now!

Don't miss an opportunity to be a part of this comprehensive collection of listings on loudspeaker materials! With 16 sections, 200 separate categories, and more than 600 companies from around the globe, the *Loudspeaker Industry Sourcebook* is an indispensable resource for the audio industry OEMs, specifiers, and buyers.

Boost Your Business!!

The Loudspeaker Industry Sourcebook is designed to help engineers, designers, purchasing managers, and key executives locate the resources that are essential to loudspeaker manufacturing operations.

Print edition • Online PDF edition • Online database

Secure your premium advertising and **expanded vendor listing** for 2023.

Also, be sure to fill out a **FREE text-only listing** at www.loudspeakerindustrysourcebook.com

For Advertising Opportunities contact:
• Strategic Media Marketing, Inc.
• P: 978.281.7708 | Email: LIS@smmarketing.us






Photo 4: This is my test setup, which includes RCA jacks, a selector switch, a volume pot, and cabling.




Figure 9: This graph shows the results from an earlier prototype, but it provides you an idea of what to expect with different tubes.

depletion mode, always on, device. Since its gate is referenced to the regulator output, it provides a fairly noise free voltage of 2.4V or greater across the regulator. The regulator requires at least 1.2V to regulate so there is some margin. The Zener keeps the gate from going more than 10V below the source in the event of a short. The 22 μ F input capacitor is there to keep the regulator from oscillating.

The MJE1535 provides a low noise virtual ground for the regulator. The base of the virtual ground comes from a temperature compensated current source and a couple of resistors to ground. The 1.25V from the regulator reference plus the 0.6V from the PNP transistor provide at least 1.85V for the current source, which requires at least 1V. The base of the PNP is connected to the current source reference, which means that its base current is factored into the 1.044mA current from the regulator. This keeps the output voltage from drifting as the base current changes. The 133 Ω and 1.33k Ω resistors along with the 1N457 provide the temperature compensation for the LM334.

Although the regulator can get by with a minimum load of 5mA, it is standard practice to draw 10mA at no-load current. The 137 Ω resistor provides a 9mA load, 1.25/137, with current source drawing a little over 1mA so there is a total of 10mA for the no-load case. The slow turn on feature comes from the 1mA current source charging the 22 μ F ripple reduction capacitor. Beyond that, it looks like the normal programmable regulator with the ripple reducing capacitor and a UF4004 to help with power down discharge. Last is the final 330 μ F output capacitor bypassed with a 2.2 μ F film capacitor. The result is a 180V regulator with a 130dB noise floor. **Figure 8** shows a top view (a) and a bottom view (b) of the power supply's PCB.

Is all that extra attention to the power supply really necessary or could you get by with something simpler? The extra attention is needed. The power-related harmonics in the output of the preamp are reduced to almost nothing as you can see in the 1kHz

Channel	Gain	THD 20kHz	THD+N 20kHz	SNR 2Vrms	Crosstalk 10kHz	CCIF	Phase Degree 10kHz	Noise Level
Left	10.53dB	0.000443%	0.002329%	107.72dB	-90.11dB	-112.56dB	Ref	8.53 μ Vrms
Right	10.49dB	0.000417%	0.002412%	108.03dB	-89.89dB	-112.74dB	-0.017	7.82 μ Vrms

Table 1: Performance measurements of the Pre-5 taken with APx515 and ECC99 tube

Channel	Noise Level	THD 20kHz	THD+Noise 20kHz	SNR 600mV	Crosstalk 10kHz	CCIF 600mV, 19kHz/20kHz
Left	6.450 μ Vrms	0.000347%	0.002176%	114.845dB	-129.242dB	-112.674dB
Right	6.743 μ Vrms	0.000314%	0.002290%	114.808dB	-128.639dB	-111.873dB

Table 2: Performance measurements of the APx515 analyzer in loopback for comparison




Figure 10: The Pre-5 THD sweep 20Hz to 20kHz 1Vrms output. The sharp rise below 50Hz is due to 60Hz noise, not distortion.




Figure 11: Here is the THD+N sweep 20Hz to 20kHz 1Vrms output.

fast Fourier transform (FFT) shown in **Photo 3**. All of the components needed to build the Pre-5 can be found in the Parts List that is posted in the Supplementary Materials section of the audioXpress website.

The Test Setup

Photo 4 shows my test setup. RCA jacks, a selector switch, a volume pot, and cabling were all included in the test results. The boards are sitting on my wooden work bench, not in a shielded case. The transformer shown is a custom AS-05T200M - 50VA 200V Antek transformer. I used the 180V AC

tap because my house voltage is high, 122V AC. I recommend that the 200V AC tap is used in general. On request, John at Antek, changed the standard 6.3V AC secondaries to 9V AC secondaries for me. (Thank you, John!)

Performance Graphs

I'll start with the THD+N of the three tubes I have tested in this preamp. **Figure 9** was derived from an earlier prototype, but this will give you an idea of what to expect with different tubes. The difference is primarily more second harmonic and a




Figure 12: This is the THD sweep of a 1kHz signal from 100mV to 10Vrms output.




Figure 13: The graph here shows the FFT of a 1kHz signal, notice the 180Hz and 300Hz power supply harmonics are barely visible.

higher noise floor. The only downside of the ECC99 is that I find it to be microphonic. If I snap my fingers, clap, or sneeze it will show up in the output. The 12BH7A and the ECC82 (12AU7) were not nearly as bad. Except for the 5963, I used JJ Electronic tubes. At some point, I would love to test the ECC802, just to see where it falls in this comparison. The 5687 is another possibility but that requires pin translation. Note: The ECC99 tubes I am testing with are at least 5 years old so the microphonic behavior observed may have changed with current production. The measurements shown in **Table 1** were taken with my APx 515 using the 5.1 PCB and an ECC99 tube. For reference, **Table 2** shows my Audio Precision 515 in loopback with the 300mV output level used to test the preamp. **Figure 10** shows a THD sweep 20Hz to 20kHz at 1Vrms output. The sharp rise below 50Hz is due to 60Hz noise, not distortion. **Figure 11** shows the THD+N sweep 20Hz to 20kHz 1Vrms output. **Figure 12** shows the THD sweep of a 1kHz signal from 100mV to 10Vrms output. **Figure 13** shows the FFT of a 1kHz signal, notice the 180Hz and 300Hz power supply harmonics are barely visible.

Figures 14–17 shows some 20kHz harmonic charts using different tubes. Pay attention to the fourth harmonic—it is an indicator of the tube noise floor. If you like cascading harmonics then the 5963 is a good choice. **Figure 14** shows the ECC99 at




Figure 14: This graph shows the ECC99 at 20kHz, with the relative harmonic levels.




Figure 16: The 5963 (Beckman, pulled from a 1950s computer) is shown at 20kHz, noise floor is similar to the ECC99.

20kHz showing the relative harmonic levels. **Figure 15** shows the ECC82 (12AU7) at 20kHz showing the relative harmonic levels, this doesn't show the high noise floor. **Figure 16** shows the 5963 (Beckman, pulled from a 1950s computer) at 20kHz, noise floor is similar to the ECC99. **Figure 17** shows the 12BH7A at 20kHz showing the relative harmonic levels.

Figure 18 shows the frequency response into a 100kΩ load. The slight bump at 20Hz is caused by the AP DC blocking filter set to 10Hz. It is not caused by the preamp. **Figure 19** shows the frequency response into a 600Ω load showing the effect of the 10μf output capacitor. **Figure 20** shows the THD+N 20Hz to 20kHz sweep 864mVrms into a 600Ω load. The voltage drop is due to the 100Ω resistor in series. **Figure 21** shows the SPICE-generated PSRR 1Vrms, solid line. Phase, dashed line. The -75dB reference is because I consider 300Hz to be the end point of the power supply related noise. Below 300Hz it is actually better.

Overall Impressions

I have a considerable vinyl and CD collection to which I enjoy listening. My listening system consists of this Pre-5 preamp, a FET-tube RIAA preamp I designed, a Technics 1300 turntable, a Marantz CD5004 CD player, Benchmark AHB2 power amp and




Figure 15: The ECC82 (12AU7) at 20kHz shows the relative harmonic levels, this doesn't show the high noise floor.




Figure 17: The 12BH7A at 20kHz shows the relative harmonic levels.




Figure 18: Here is the frequency response into a 100kΩ load.




Figure 19: This graph shows the frequency response into a 600Ω load showing the effect of the 10μf output capacitor.

GETTING THE MOST OUT OF
VACUUM TUBES

By Robert B. Tomer

Vacuum Tubes—The Important 4/5ths of Tube Circuits

Over 80% of all tube equipment defects result, directly or indirectly, from tube failures. Why do tubes fail? What can be done to prevent them from failing before their time? How can you determine whether a tube is good or bad, or how well and how long it will work in a given circuit? Should tubes be replaced periodically, whether they've failed or not... or should they be tested every so often, and replaced if indications show them to be below par? This book supplies the answers to these profound questions...plus many, many more.

Like Watching a Play From Backstage

Author "Bud" Tomer, an international authority on tubes, is well-known for his candid "Tech Tip" bulletins and his slide-film presentations to service, military and engineering groups. Within this book you'll find a virtual encyclopedia on vacuum tubes—yet the author has not delved into mathematical details that would be of interest to only a few tube engineers. Everything is focused toward the interest of those who use tubes—either for replacement purposes or in original equipment designs.

In his refreshing, down-to-earth style, Mr. Tomer sheds light on the "Why so many tube types?" "What About Tube Testers?", etc.

Other Interesting and Informative Topics:

- 10 reasons why tubes fail
- Predicting tube performance
- Selected and proper tube replacement
- Lengthening tube life

Anyone who has anything at all to do with tube equipment design, or proper use—should read this book. It is a must for anyone who designs, builds, repairs, and maintains consumer goods such as radio, TV, and audio equipment, or in production and manufacturing processes—or even computers and radar gear.

You'll also learn about:

- Catastrophic failures—what they are, when and how they occur
- Degenerative failures—what happens to tubes over time
- Subjective failures—why a seemingly normal tube may be failing
- Characteristic variables—differences that make tubes unique
- Special-purpose tubes—considerations for filament and thyratron types.

Order your copy
cc-webshop.com

DO IT YOUR WAY

Panels/Boxes/Cables

WITH REDCO AUDIO'S

"Build Your Own" web pages

Full control for design, editing and more.

WE'LL BUILD IT JUST THE WAY YOU WANT!

REDCO AUDIO

www.redco.com 203-502-7600




Figure 20: The THD+N is shown in a 20Hz to 20kHz sweep 864mVrms into a 600Ω load.




Figure 21: This graph shows the SPICE-generated PSRR 1Vrms, solid line. Phase, dashed line.




Photo 5: My listening system consists of this Pre-5 preamp, a FET-tube RIAA preamp I designed, a Technics 1300 turntable, a Marantz CD5004 CD player, Benchmark AHB2 power amp and GoldenEar Aon-2 book shelf speakers.

Project Files

To download the Parts List for this project, visit:

<http://audioxpress.com/page/audioXpress-Supplementary-Material.html>

Resource

J. Broskie, Tube CAD Journal,

<https://tubecad.com/PublishingInfo.htm>

GoldenEar Aon-2 book shelf speakers (**Photo 5**). The sound of this simple system is amazing. I love the highs and mids, it brings out the space in the music. Since I am an electric bass player, I pay special attention to the sound of the bass. The sound reflects the actual recording without adding any tube exaggerations.

There is one weakness that gets exposed when using a subwoofer and playing music with a lot of very low-frequency energy. Think rap and low peddle notes in organ music. These notes can cause significant air pressure changes in the room. The pressure changes can interact with the microphonics of the ECC99 tube and create some messy, muddy low frequency resonances. Other than that, this preamp is a real pleasure. If lots of loud low subwoofer driven bass is a requirement, one of the other tubes mentioned might be a better choice or simply provide better isolation between the tube and the environment.

As a side note, the first implementation of this circuit was built in April of 2019 and I called it Pre-3. Next, I built Pre-4, which made some noise floor improvements but went crazy with PIC micro controllers for IR remote control, relay volume control, relay input selection, and mute management. This version, the Pre-5, changed the power supply from 250V to 180V and removed the CPU control, going back to simple analog controls.

Author's Note: Readers may contact me with questions or comments at: jbvskov@gmail.com.

J SERIES

THE LEGENDARY J SERIES IS BACK!



TRIAD'S CLASSIC J SERIES TRANSFORMER DESIGN COMES OUT OF THE VAULT WITH ITS LEGENDARY PERFORMANCE PEDIGREE AND MORE. BORN IN THE 1950S, THE J SERIES TRANSFORMERS FEATURES HIGH-NICKEL LAMINATIONS, PAPER SECTION WINDINGS AND A SPECIALLY TOOLED MU-METAL SHIELDING ENCLOSURE. LONG RECOGNIZED BY ENGINEERS FOR ITS QUALITY SOUND, TODAY'S J SERIES IS IDEAL FOR NEW OR VINTAGE AUDIO GEAR, SOUND PROCESSING AND ANALOG STUDIO RECORDING EQUIPMENT.

SHIELDED AUDIO IMPEDANCE MATCHING
DESIGN FOR A WARM SIGNAL.

SINGLE-HOLE MOUNTING ALLOWS ROTATION
FOR MAXIMUM HUM REDUCTION.

VINTAGE CIRCUITRY, PREMIUM
MATERIALS, HAND-CRAFTED QUALITY.

VISIT ONLINECOMPONENTS.COM TO
ORDER OR LEARN MORE.

MADE IN THE U.S.A.



onlinecomponents.com

THE ONLINE DISTRIBUTOR OF ELECTRONIC COMPONENTS

Passive and Aggressive

EPIQUE
by Dayton Audio



Patented MMAG Motor System

Capable of unheard-of excursion for their size, Epique drivers are engineered to produce exceedingly low distortion with controlled response. Allowing more versatility than ever, you now have the opportunity to create systems that will deliver game changing performance.

Carbon Fiber Is All About Performance

Dayton Audio's new ground-breaking Epique E150HE-PR 5-1/2" and E180HE-PR 7" passive radiators with carbon fiber cones are the perfect companion for the Epique Series MMAG by retaining the same family look. Engineered to replace the excessive port lengths required for low-tuned enclosures, these passives give you the opportunity to create systems that provide extraordinary accuracy and fidelity without the extraordinary size.

Epique low-distortion drivers and matching passive radiators

Epique DVC MMAG Extended Range Subwoofers

- Patented dual gap motor design allows exceptional linearity
- Over 14 mm of excursion (at 70% BL), 28 mm peak to peak
- Damped woven carbon fiber cone delivers remarkably smooth response
- Inductance control for extremely low distortion

Epique Carbon Fiber Cone Passive Radiators

- Carbon fiber cone is strong enough to handle the pressure of high output drivers
- Incredibly linear 19mm xmech enhances output at low tunings
- Extremely versatile, perfect compliment for any compact subwoofer

Shop the Entire Epique Series



DAYTONAUDIO
daytonaudio.com

Distributed by:
PARTS EXPRESS
AUDIO COMPONENTS & SOLUTIONS
parts-express.com 1-800-338-0531
Custom OEM inquiries welcome

parts-express.com/brand/
EPIQUE-by-Dayton-Audio