www.audioxpress.com

audioxpress

ADVANCING THE EVOLUTION OF AUDIO TECHNOLOGY

Market Update

DSP and Wireless Audio Applications
Technology and Trends
By J. Martins

Fresh From the Bench

Audio Precision APx555 B Series Audio Analyzer
Expanded Capabilities and Software
By Stuart Yaniger

Development Platforms

XMOS xCORE VocalFusion
By David Logvin

You Can DIY!

An 18 W Single-Ended Valve Amplifier with Trans Technology
By Menno van der Veen

Sound Control

Keeping the Sound (or Noise) In or Out
By Richard Honeycutt

Audio Scope

Combining Cascaded Stage Harmonic Distortion
By Michael Steffes

Audio Electronics

Reggie-Phil 2.0: A Current-Regulated Filament Supply
By Christopher Paul

$9.00US $10.00CAN

DECEMBER 2019

Whatever your testing challenge, GRAS has the right acoustic sensor for your audio application.

Only GRAS offers a complete line of high-performance, test microphones and related products ideal for use in consumer audio and electronics applications.

As the leading global provider of microphones, GRAS has a long tradition of working with audio engineers to ensure accurate data is captured, each and every time.

Microphones from GRAS are designed for the high quality, durability and reliability that our customers demand.

Contact GRAS today for a free evaluation of the perfect GRAS microphone for your application.

GRAS

Acoustic Sensors

- Measurement microphone sets
- Microphone cartridges
- Preamplifiers
- Low-noise sensors
- High resolution ear simulators
- Head & torso simulators
- Test fixtures
- Custom designed microphones
- Speech intelligibility
- THD
- Frequency response
- Calibration systems and services

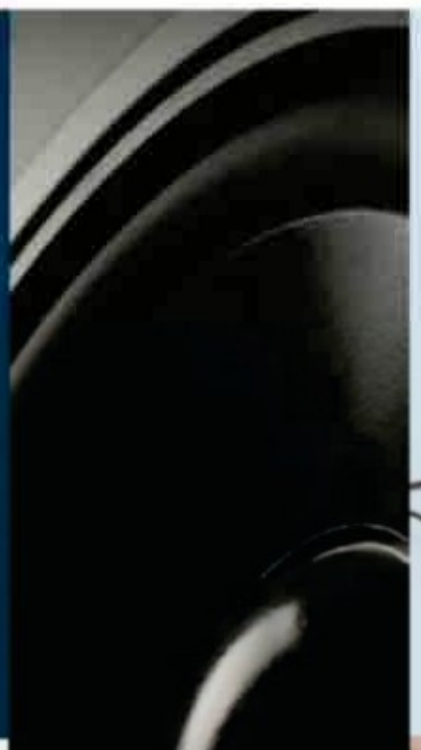

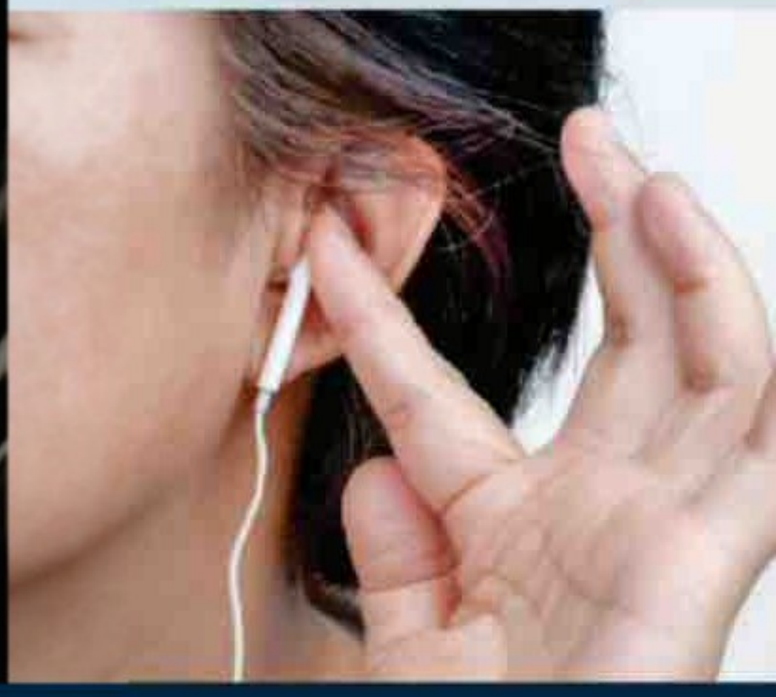

P: 800.579.GRAS
E: sales@gras.us

www.gras.us

HAMMOND MANUFACTURING®

1140 Series

Studio-Grade Impedance Matching Transformers

Rooted In Solid ENGINEERING

- Microphone Input
- Microphone Splitter
- Line Input
- Line Output
- Direct Box

- Designed with Nickel alloy cores for high performance and Wide Bandwidth
- Frequency Response: 20 Hz – 20 kHz ±0.35 dB, (ref. @ 1 kHz)
- Low distortion at levels up to +27 dBu @ 20 Hz
- Excellent Shielding – Mu Metal enclosures
- Line Output Transformers are open style (not Mu Metal)

www.hammondmfg.com

December 2019 **ISSN 1548-0628**

www.audioxpress.com

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

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

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

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

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

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

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.

© KCK Media Corp. 2019
Printed in the US

The Magic of DSP

It's hard to realize how important the use of digital signal processing (DSP) already is in our audio devices. So many good-sounding products we have today have a bit—or a lot!—of DSP magic doing its part, even very high-end products. But of course, the benefit for the vast majority of people is in affordable products manufactured by the millions.

I will always remember a demonstration from my friend John Richards, from Oxford Digital, when he showed me how a pair of cheap plastic computer speakers that sounded absolutely terrible could suddenly be turned into a pretty decent sound system just by applying the company's SmartFx DSP tuning tool, running on a laptop. John Richards is a respected consultant for many consumer electronics giants and he frequently travels to Asia to follow the many OEM projects we probably already have in our homes, bags... or pockets.

Oxford Digital's SmartFx is a relatively simple but extremely effective suite of audio DSP effects designed to overcome limitations in a device's acoustic performance—particularly effective with those devices that have severe limitations. The processing acts at the crossover level to optimize the response correction, adds Low Frequency Boost by applying a compressive style Bass Enhancer to assist the loudspeaker where it has some partial response but is not functioning well and uses Dynamic Gain Control to make the signal louder, finally applying High Frequency Boost as a combination of a Limiter section to make the signal louder, combined with LoudMaster, an effect that can increase the volume without increasing peak signal level. The result is pure magic!

This is how, in large part, we can find products that seem to defy the laws of physics, like the JBL GO portable speaker. I've seen this tiny $25 speaker being used in taxis, rickshaws, and tuk-tuks, from Europe to Africa and Asia, sometimes as the only available entertainment in remote areas of the world. And it sounds fine when there's nothing else, believe me.

No wonder that during the IFA 2019 show in Berlin, Germany, Harman announced that it set the new record with 100 million portable speakers shipped worldwide, cementing its position as leader in the category with a 34.2% volume share.

Speakers like the JBL GO are truly impressive compared with what other products costing many times more can actually do. As Harman explains, "...paying close attention to the habits and lifestyle of its users, and focusing on the real-life scenarios when people needed audio..." But this only a small sample of what it is possible to do with a little DSP magic (yes and good electronics and an excellent driver for its size...).

We are already seeing extreme examples in smartphones, headphones, and miniaturized speakers and drivers. Imagine what we will be able to do in large auditoriums and spaces, controlling hundreds of speakers in the same way, not with rack units of equipment but simply with chips that will be built-in with each driver!

In this edition of *audioXpress*, David Logvin explores the XMOS xCORE VocalFusion Speaker Linear Kit (XK-VF3100-L33), a development platform comprised by a BaseBoard built around the XVF3100-TQ128 dedicated RISC microcontroller, and a separate linear microphone array. This MCU is one of the most basic models from XMOS but is able to reach up to 2400 MIPS, supporting 32×32, 64-bit MAC instructions for DSP. This is a perfect example of an embedded solution that is available to audio developers and OEMs, supporting the features that not long ago would require extremely expensive hardware, and making sophisticated DSP features extremely affordable. This specific development kit was created to enable adding voice control to home appliances, for instance. But as David discusses in his report, and our Market Update focused on DSP and Wireless Audio Applications explores even more, there is a world of exciting possibilities waiting for audio product designers... way beyond voice recognition.

As is frequently the case when an industry or market application grows exponentially, the underlaying technology also evolves at a much faster pace and solutions become much more affordable.

J. Martins
Editor-in-Chief

The Team

President:	KC Prescott	Associate Editor:	Shannon Becker
Controller:	Chuck Fellows	Graphics:	Grace Chen
Editor-in-Chief:	João Martins	Advertising Coordinator:	Nathaniel Black

Technical Editor: Jan Didden

Regular Contributors: Bruce Brown, Bill Christie, Bob Cordell, Joseph D'Appolito, Vance Dickason, Jan Didden, Scott Dorsey, Gary Galo, Gerhard Haas, Chuck Hansen, Richard Honeycutt, Charlie Hughes, Mike Klasco, Ward Maas, Oliver Masciarotte, Nelson Pass, Christopher Paul, Bill Reeve, Fernando Rodrigues, Steve Tatarunis, Stuart Yaniger

OUR NETWORK

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

SUPPORTING COMPANIES

ACO Pacific, Inc.	63	Hypex Electronics BV	51
All Electronics Corp.	29	KAB Electro-Acoustics	43
Audio Precision, Inc.	67	Menlo Scientific, Ltd.	53
Avel Lindberg, Inc.	63	Parts Express International, Inc.	68
Celestion	45	Purifi Audio	11
Comsol, Inc.	37	Redco Audio	29
GRAS Sound & Vibration	2	Sonarworks, SIA	41
Hammond Manufacturing, Ltd.	3	Swans Speakers Kits	43
HEAD acoustics, Inc.	65		

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 edition of the magazine.

COLUMNISTS

Vance Dickason has been working as a professional in the loudspeaker industry since 1974. He is the author of *Loudspeaker Design Cookbook*—now in its seventh edition and published in English, French, German, Dutch, Italian, Spanish, and Portuguese—and *The Loudspeaker Recipes*. Vance is the editor of *Voice Coil: The Periodical for the Loudspeaker Industry*, a monthly publication. Although he has been involved with publishing throughout his career, he still works as an engineering consultant for a number of loudspeaker manufacturers.

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

Mike Klasco is the president of Menlo Scientific, a consulting firm for the loudspeaker industry, located in Richmond, CA. He is the organizer of the Loudspeaker University seminars for speaker engineers. Mike specializes in materials and fabrication techniques to enhance speaker performance.

Michael Steffes has worked for more than 32 years in high-speed signal path IC development, spanning IC design roles with the earliest current feedback amplifiers, roadmap definition, PLM for the BurrBrown high-speed amplifier products, and product launch/application manager roles through Texas Instruments/Intersil. He continues to focus on high-performance ADC interface solutions, DSL/PLC line driver design and optimization, and signal path design optimization tools. In a career spanning five different major IC vendors, he has participated in the introduction of 150 industry leading amplifier products (setting many of the now standard practices in high speed amplifier characterization curves), publishing over 25 application notes and a growing list of contributed articles from 1989 onward.

STATEMENT REQUIRED BY THE ACT OF AUGUST 12, 1970, TITLE 39, UNITED STATES CODE SHOWING THE OWNERSHIP, MANAGEMENT AND COPY CIRCULATION OF audioXpress. Published monthly at 650 Dairy Farm Rd, Red Oak, VA 23964. Annual subscription price is $50.00. Publisher: KC Prescott. The owner is KCK Media Corp., 650 Dairy Farm Rd, Red Oak, VA 23964. The names and addresses of stockholders holding one percent or more of the total amount of stock are: KC Prescott, 650 Dairy Farm Rd., Red Oak VA 23964. EXTENT AND NATURE OF CIRCULATION: Average number of printed copies of each issue published during the preceding twelve months *For each single issue nearest to filing date; total number of copies printed, 1800 *2890; paid/requested mail print subscriptions, 710 *725; sales through dealers and carriers, street vendors and counter sales, 201 *215; paid/requested copies distributed by other mail classes, 24 *22; total paid/requested print circulation, 935 *962; Nonrequested copies distributed outside the mail, 841 *1900; total nonrequested distribution (sum of D.1 & D.4), 841 *1900; total distribution (sum of C & E), 1776 *2862; copies not distributed (office use, leftover, unaccounted, spoiled after printing, returns from news agents), 24 *28; total (sum of F & G), 1800 *28902. Percent Paid Requested: 52.64% *33.6%; total number of paid Electronic Copies, 1233 *11139; Total Paid Print & Paid Electronic Copies 2168 *2102; Total Print Distributions and Paid Electronic Copies 3009 *4001; Percent Paid: 72.05% *52.53% I certify that the statements made by me above are correct and complete. KC Prescott Publisher.

Contents

audioXpress

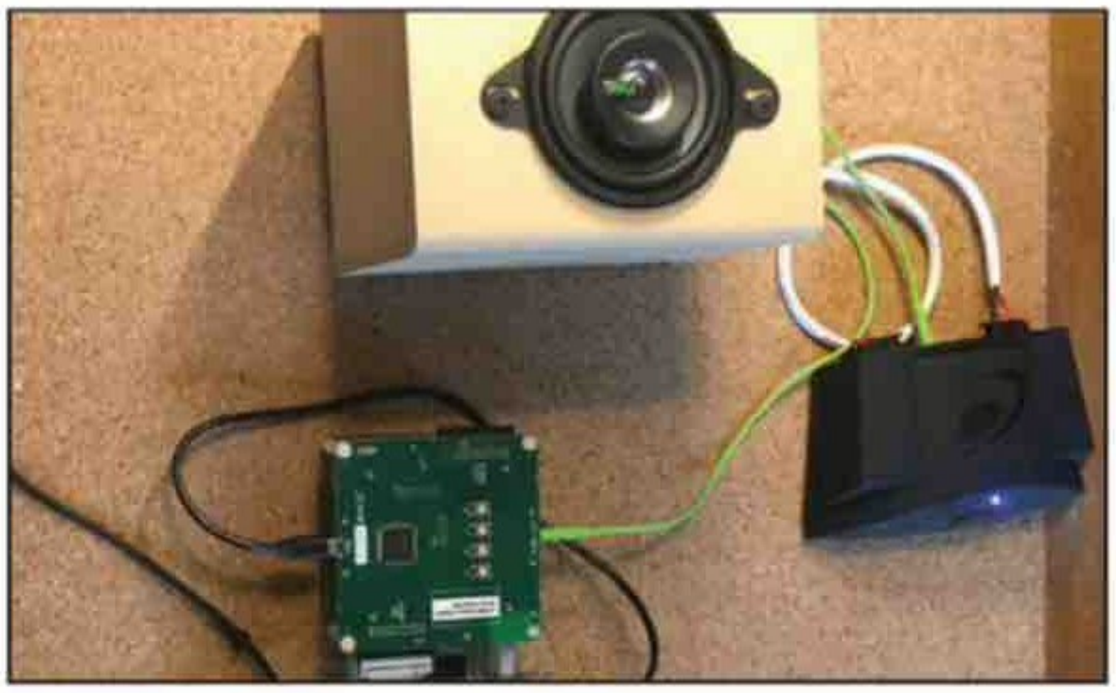

Features

30 Audio Precision APx555 B Series Audio Analyzer

Expanded Capabilities and Software

By Stuart Yaniger

Audio Precision recently loaned us one of its newest analyzers, the APx555 B Series. Stuart Yaniger had one month to test it, and in this article he provides readers with a hint of the new expanded capabilities—both of the APx555 B Series hardware and the new APx500 5.0.1 software.

38 The xCORE VocalFusion Speaker Linear Kit by XMOS

How Many Microphones Does It Take to Run Your Life?

By David Logvin

Learn more about the VocalFusion Speaker Linear Kit, designed for those who want to integrate voice user interfaces into consumer electronics typically found at the edge of the room, such as smart home control panels and washing machines.

46 Reggie-Phil 2.0: A Current-Regulated Filament Supply

By Christopher Paul

Inspired by previous approaches published by *audioXpress*, this article explores alternative options to the design of a linear current regulator for audio tube amplifiers, including employing feedback to sense load and other design strategies.

54 Build an 18 W Single-Ended Valve Amplifier with Trans Technology

By Menno van der Veen

This article shows an easy-to-build 18 W single-ended valve amplifier using Trans technology with pentodes, current sources, and toroidal transformers. The amplifier was created by the author to support his TubeSociety students in the Netherlands.

Volume 50 – No. 12

December 2019

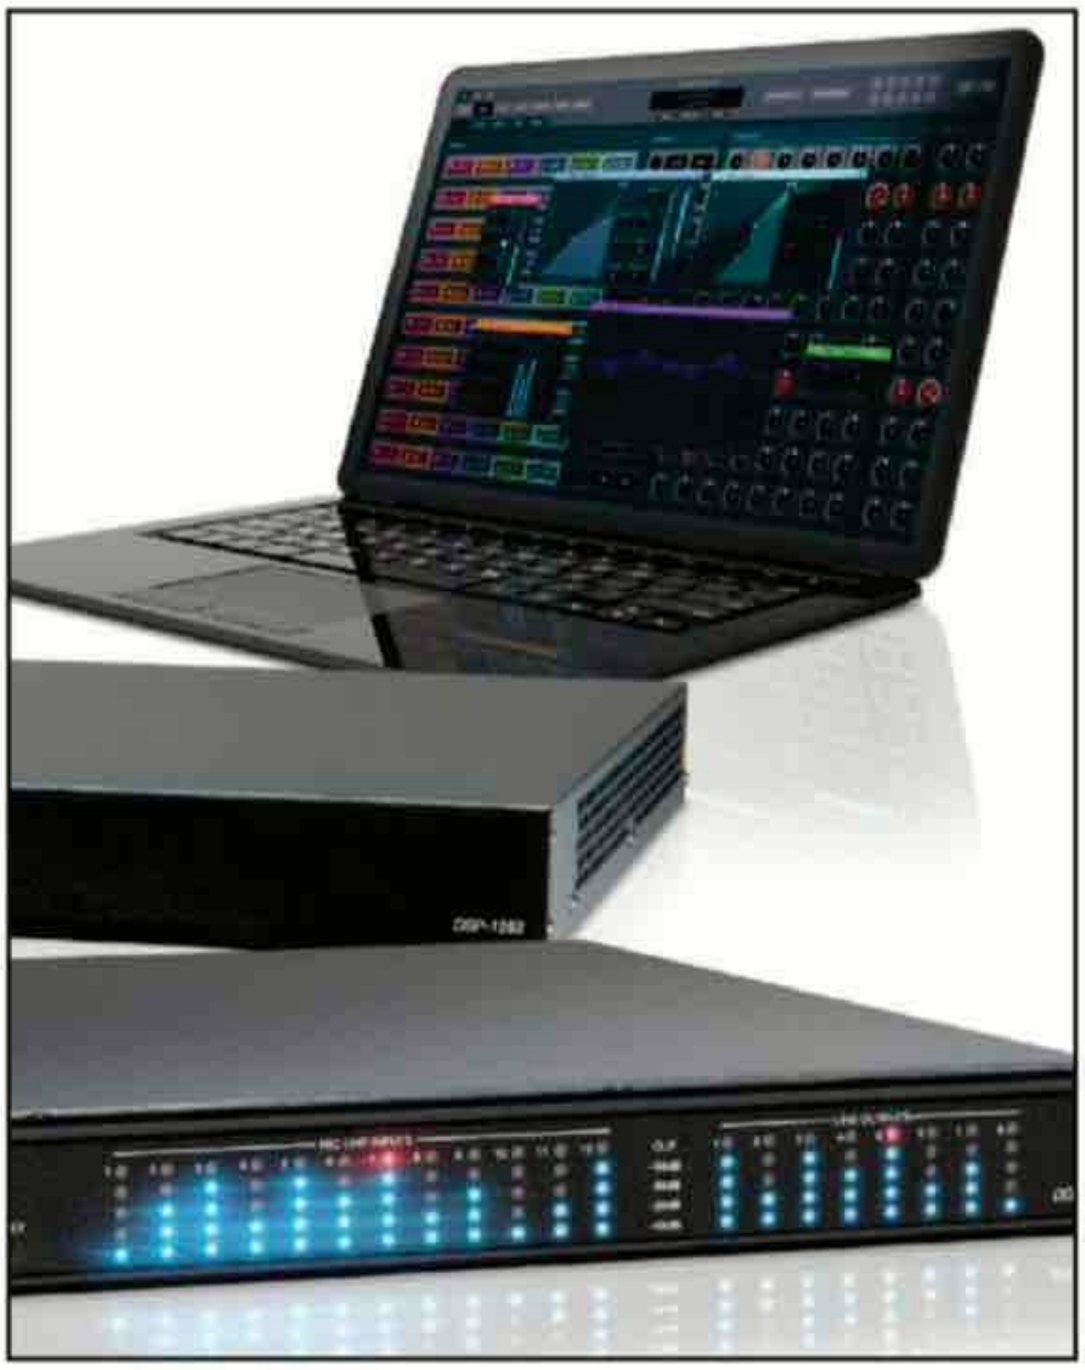

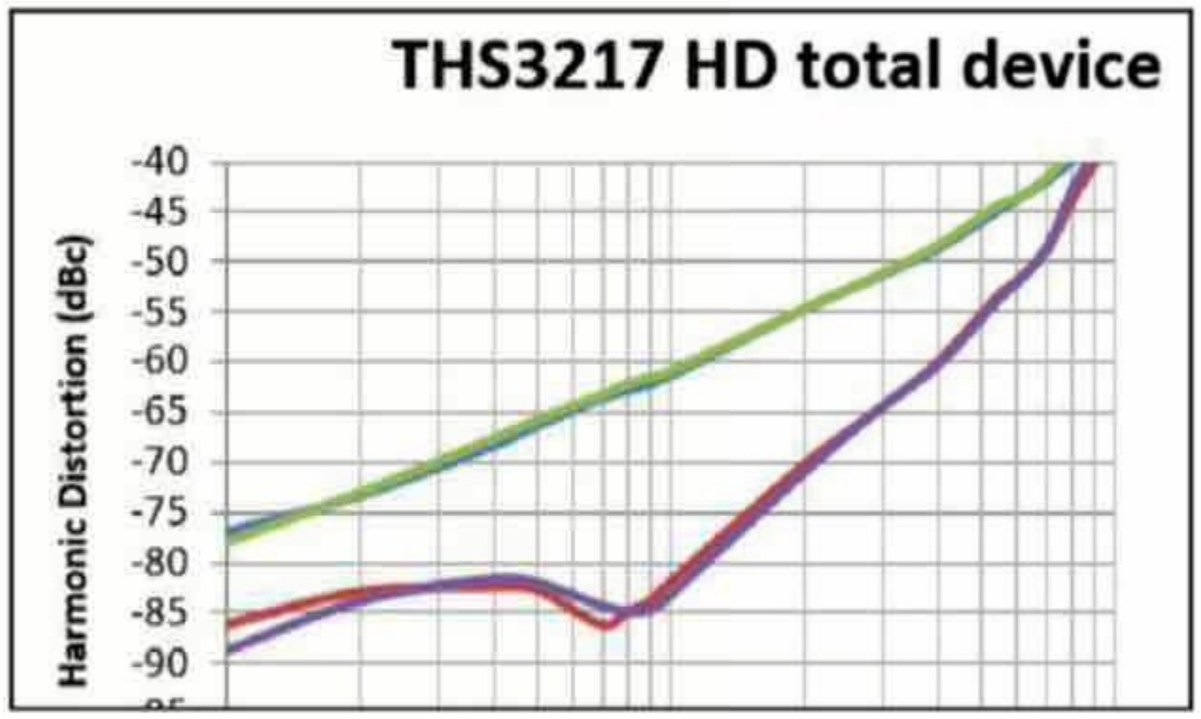

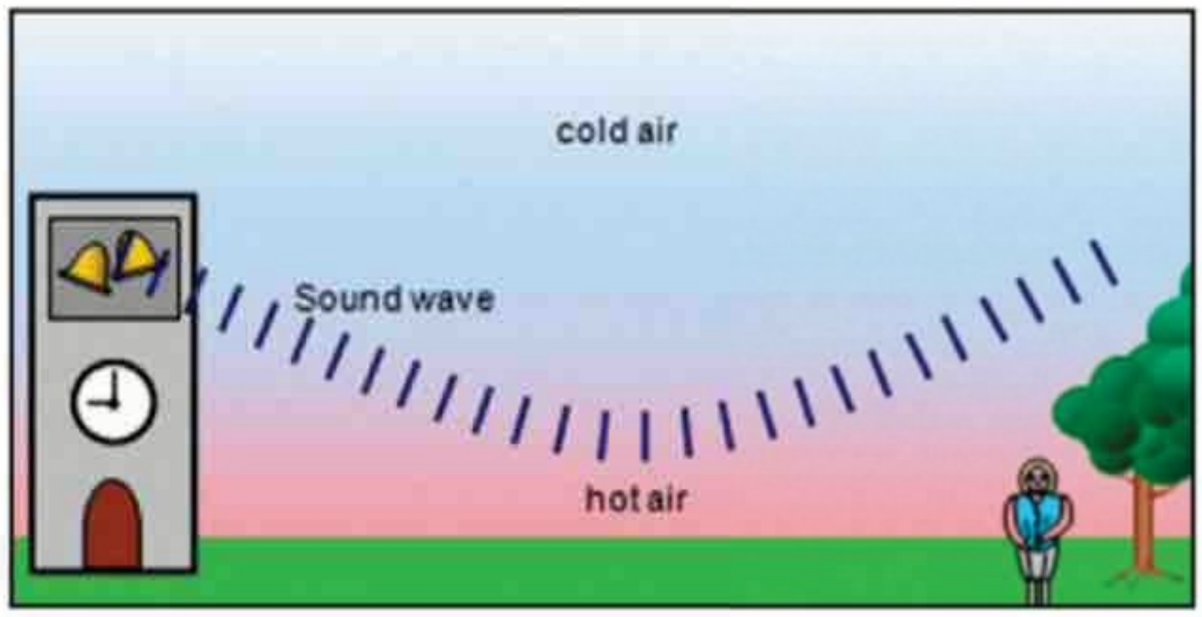

Market Update—DSP and Wireless Audio Applications

8 Cutting-Edge Audio

No Wires, Room Optimized

By J. Martins

In this Market Update, we examine wireless audio and digital processing technologies, with an overview of all the latest solutions and trends in these two critical but extremely exciting domains for audio product designs.

Columns

SOUND CONTROL

24 Keeping the Sound (or Noise) In or Out

By Richard Honeycutt

AUDIO SCOPE

28 Combining Cascaded Stage Harmonic Distortion

By Michael Steffes

HOLLOW-STATE ELECTRONICS

60 Plug-In Modifications for Hollow-State Guitar Amplifiers?

By Richard Honeycutt

Departments

4 From the Editor's Desk

5 Client Index

66 Industry Calendar

Websites

www.audioxpress.com
www.voicecoilmagazine.com
www.cc-webshop.com
www.loudspeakerindustrysourcebook.com
www.circuitcellar.com
www.linuxgizmos.com

@audioxp_editor audioxpresscommunity
linkedin.com/company/audioxpress

Cutting-Edge Audio

No Wires, Room Optimized

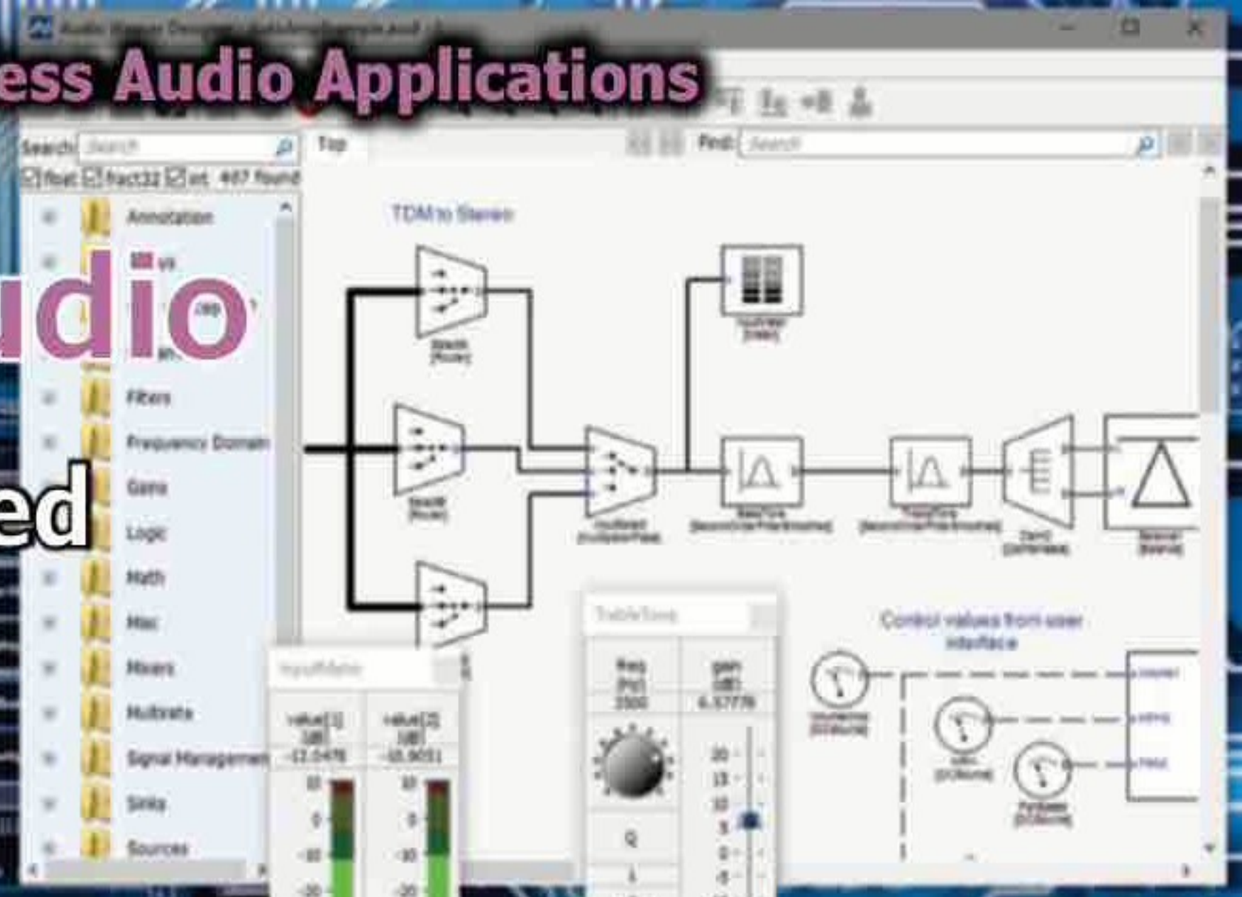

For this edition of *audioXpress,* we are examining wireless audio and digital processing technologies, with an overview of all the latest solutions and trends in these two critical but extremely exciting domains for audio product designs.

By
J. Martins
(Editor-in-Chief)

In this overview, we will mainly look at the latest technologies covering configurable audio processing, in particular dedicated to room compensation approaches and speaker frequency response correction, among other key application areas.

audioXpress readers are probably getting tired of repeatedly reading references to "smart speakers," "smart amps," "smart home," smart this, and smart that. Why such an insistence on "smart," sometimes to characterize completely different solutions? The most obvious reference—and the most mainstream—sees "smart" referring to connected devices that respond to voice commands—hence the reference to smart speakers being commonly associated with speakers from Google or Amazon, which have popularized the use of personal voice assistants using the English language as the interface. The same happened in China, where "speakers" from Baidu, Alibaba, Xiaomi, Huawei, and Tencent do the same with Chinese (standard Chinese or Mandarin).

If we look beyond voice applications in English and Chinese, those "smart" devices have to depend largely on other types of "intelligence" since support for truly interactive voice assistants in other languages is still primitive—and most often limited to voice command recognition, not natural language.

Because current audio electronics projects are mostly microprocessor, microcontroller, and increasingly System-on-Chip (SoC) based, audio designers are able to explore the platforms' abilities to combine cutting-edge features, among which the most "intelligent" are effectively wireless audio and DSP.

Among the most critical areas of additional "intelligence," are the ability for connected speakers to support additional wireless audio protocols and content streaming services, while at the same time using its built-in DSP abilities to optimize the speaker's drivers and amps own responses. The speaker optimization approach is consensually what audio designers mostly associate with truly "smart" audio devices, not only able to use DSP to optimize the system itself, but increasingly able to compensate and adjust to the surrounding acoustical conditions or room response.

While wireless audio features typically support Bluetooth and Wi-Fi protocols, there are many more options to consider when designing a connected audio product, from supporting different wireless protocols, to direct support of streaming services such as Spotify Connect, where a device—a wireless speaker, soundbar, AV receiver, or smart speaker—is able to play content from Spotify's streaming service directly over Wi-Fi, without being dependent of a computer,

smartphone of tablet, which people typically use as sources to access their music services. And Spotify Connect is just one of the examples of a "must have" option in any audio product launched in the market in 2019, as well as Apple's AirPlay 2 support or support for the Wireless Speaker and Audio Association (WiSA) technology specification, which allows wireless audio support for home-theater systems. And those are just some of the basic wireless audio options to consider...

Processing Hardware

When we address the topic of DSP in audio products, the first thing that comes to mind are obviously the signal processing functions of filters, crossovers, noise and echo removers, equalization and dynamic processing, limiting, and so forth. But DSP capabilities in modern audio platforms are responsible for doing a lot more at the basic level—from managing network and serial interfaces, memory, and displays to audio sample rate conversion, media formats decoding, and conversion, timers and audio peripherals. Those are all essential and structural features on the design of any current audio solution, consumer or professional.

In the professional audio industry, there are dedicated audio processing platforms that are able to power most professional audio and networked audio applications. Companies such as Xilica Networked Audio DSP & Control Solutions and Symetrix, to name just a few, are typical examples in this domain, offering modern platforms that are able to manage and distribute audio processing power wherever it is needed. Many of the dedicated DSP solutions for professional and commercial applications were typically ASIC platforms that gradually evolved to field-programmable gate array (FPGA) and increasingly to fully software-based solutions now running on Intel x86 and ARM computing hardware. As an example, a typical music-oriented hardware effects processor, such as the Eventide H9000 Harmonizer (the latest generation), relies on four quad core ARM processors for all its DSP.

On another more "creative" level, in the audio industry—and in particular in the music industry and recording world—we tend to associate DSP capabilities with audio effects processing, increasingly powering extremely sophisticated algorithms, such as convolution reverb or pitch-independent time-stretch manipulation, to name just two of the most "colorful."

But DSP is also being increasingly used to power a new level of functions that we can—and should—associate with truly "smart" devices. Processing and analyzing the information that is generated by sensors—and microphones are just some of the "basic" sensors that we normally have in audio designs.

Crestron Avia multichannel digital signal processors (DSP) and software tools for commercial applications

On the commercial, residential, and high-end side of the industry, sophisticated audio processing solutions (e.g., those from Trinnov Audio) used for loudspeaker/room optimization and immersive audio installations in home theaters, professional studios, and commercial movie theaters use Intel CPU-based platforms, with all the audio circuit boards designed and manufactured by the company. Like many of the current dedicated professional or high-end audio processing solutions, these are basically powerful PCs that are custom designed for audio applications and environments.

That can be an ideal solution for very powerful applications, mainly fixed installations, where separated hardware components are viable. But increasingly for cutting-edge audio development, the industry is moving toward fully integrated solutions, such as speakers with fully integrated amplification, electronics, connectivity, and... intelligence.

Self-Aware Audio

When we discuss "smart platforms" and we mention the "intelligence" that should be inherent in new-generation audio systems, we need to consider that everything depends on the basic connectivity

The Eventide H9000 Harmonizer effects processor relies on four quad core ARM processors.

CEVA's high-performance audio and voice DSPs are optimized for intensive sound processing applications. They have been specifically designed to tackle multi-microphone speech processing use-cases, high-quality audio playback and post-processing, and on-device sound neural network implementations.

and sensing capabilities in first place, immediately followed by the ability to process the information generated by those interfaces. DSP functions are typically used to process the information generated by one or multiple microphones in order to recognize speech and command words in voice systems, but also to analyze the response of the built-in speaker system in the room, analyzing echoes, reverberation, delays, and interferences.

Using aggregated data from more sensors (e.g., a motion sensor), a smart speaker "knows" that it has been moved and should generate a new calibration process to optimize the speaker's response for the new position. And all the data generated by multiple sensors—accelerometer, gyroscope, magnetometer, thermometer, barometer, air humidity sensor, or even GPS, proximity sensors, light sensors, etc.—also needs to be filtered, calibrated, and so forth. And in portable and mobile devices, which are battery-dependent, all those functions need to run with as low power as possible.

That is why companies such as CEVA have created integrated processing solutions that are able to combine and process multiple sensors, what the company calls "sensor fusion" or a DSP-based sensing solution that is also able to perform multiple sensing tasks and is even able to "learn"—powering locally what the industry likes to call artificial intelligence (AI) at the edge. Taking it further, the same DSP core is able to run sensor fusion, wireless protocols such as Bluetooth Low Energy (BLE), and even audio codecs—all with extremely low power levels.

For most generic consumer electronics, connected home or mains powered applications where such low power requirements are not as demanding, it becomes possible to build hardware modules that combine audio processing in embedded platforms to create the core for most of the features needed in current audio designs. As an example, many "smart" audio products are currently designed using solutions from StreamUnlimited, a supplier of software solutions (StreamSDK) and hardware modules for connected audio products.

Partnering with all major semiconductor companies and technology providers in the consumer electronics channel, StreamUnlimited is able to support original device manufacturers (ODMs) in launching high-end audio, residential integration (CEDIA-channel), premium, and mainstream consumer electronic products. The company's latest Stream1955 hardware module powers soundbars with a powerful Quad Core SoC with ample processing power and a multitude of signal interfaces, helping brands to support features from Dolby Atmos decoding directly in the SoC (no extra DSP required), to extended wireless connectivity, and support of all major voice assistants.

Another more generic hardware module from StreamUnlimited is the Stream1832, meeting all Google Assistant specification requirements. This module used typically in smart speakers or network audio streaming players, provides a powerful quad core platform and 1.5 GHz CPU designed for Google Assistant, with Chromecast built-in, AirPlay 2, Spotify Connect, and Amazon Voice Services. In combination with the company's StreamSDK software, the module provides a development solution optimized for cost and low power and is only one of the examples of ready-available solutions for connected audio designs,

StreamUnlimited Stream1832 module is a quad core platform supporting Chromecast built-in, AirPlay 2, Spotify Connect, Amazon Voice Services and Google Assistant for connected audio designs.

PURIFI

www.purifi-audio.com

PURE SOUND

Building a Straight Wire
to the Soul of Music

FIRST PRODUCTS
AVAILABLE NOW

PTT6.5W04

6.5" Ultra Low Distortion Long Stroke Driver

- Powerful, precise bass
- Lifelike and uncluttered mids
- Effortless musicality and extreme detail retrieval
- Get "beyond 3-way" performance from a simple 2-way

1ET400A

Ultra High Performance Class D Amplifier Module

- Unrivaled measured and sonic performance
- Outperforms amplifiers of any operating class
- Very low idle losses
- Straightforward interfacing and system integration

When Sound Comes First
sales@purifi-audio.com

allowing brands to address the ever-increasing technical complexity of dealing with multiple voice and streaming services ecosystems.

Being a flexible and modular software solution, StreamUnlimited's StreamSDK also helps ODMs to add support for a vast range of music and Internet radio services, and support for the latest technologies, formats, and protocols, for high-resolution audio, UPnP/DLNA, USB, Bluetooth, wireless and Ethernet connections, GUI, Android and iOS applications and even adding a proprietary Multiroom streaming solution.

Qualcomm Smart Audio Platform

Recently, Qualcomm Technologies also made two important announcements for audio developers, offering new alternatives for development of smart speakers and new-generation whole-home audio systems. The San Diego, CA company introduced the new QCS400 family of audio SoCs and the Smart Audio Platform 400 development kit, combining high-performance processing, superior connectivity, and multiple advanced audio and visual display capabilities in a single, power-optimized chip architecture.

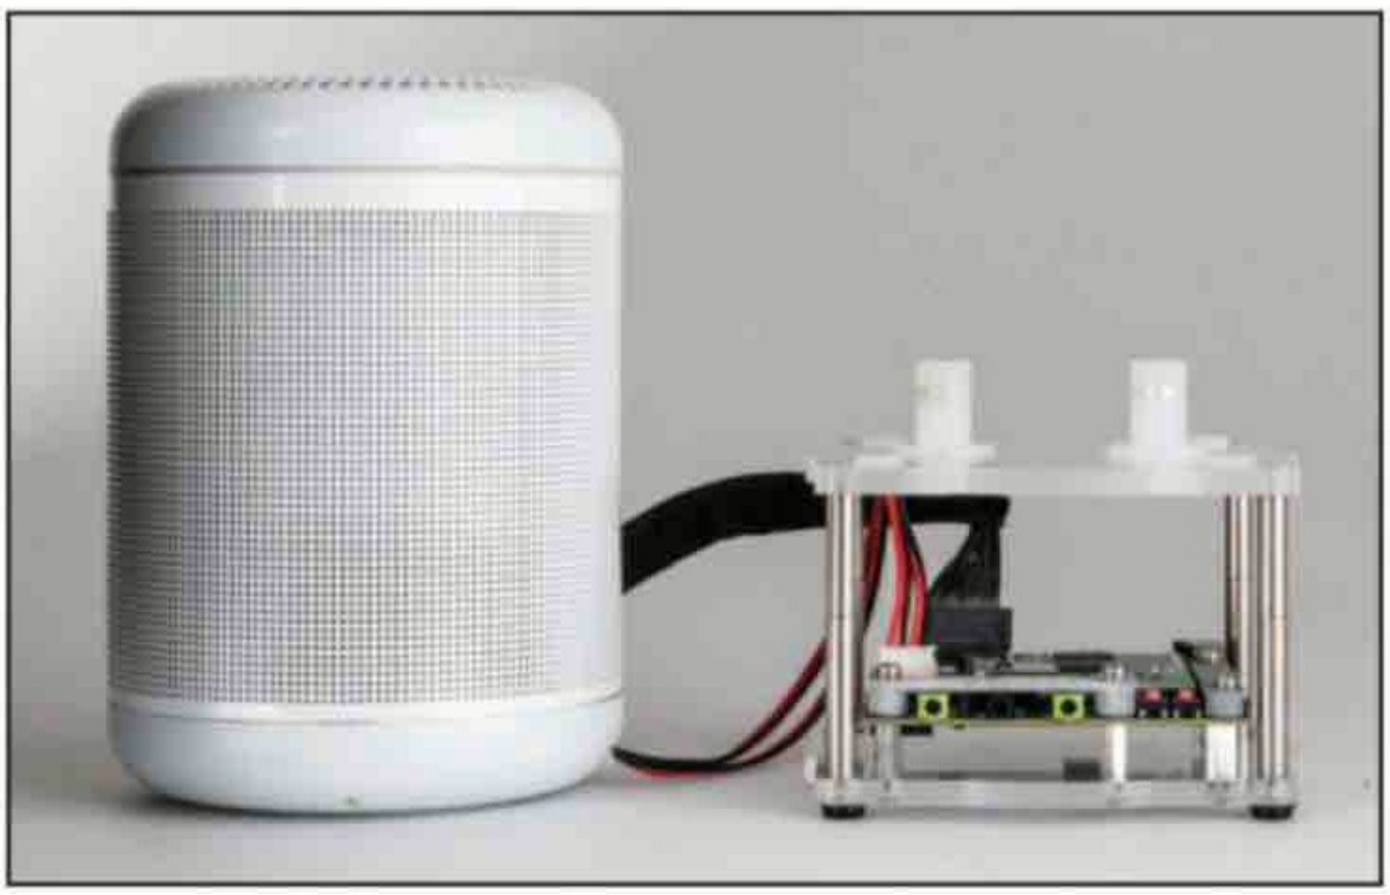

Qualcomm Smart Audio Platform 400 development kit

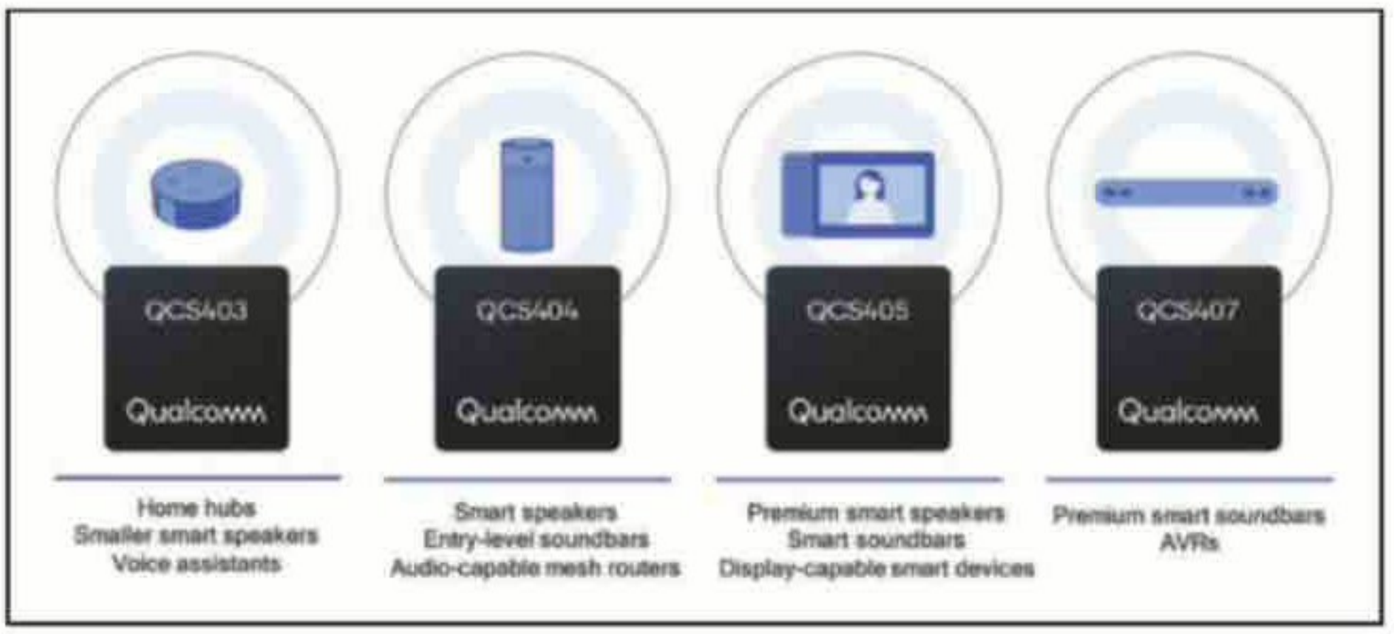

The Qualcomm QCS400 SoC series features the Qualcomm AI Engine, dual DSPs, and up to four processor cores, engineered specifically for home audio applications.

The new Qualcomm QCS400 SoCs are designed to power "smart" home hubs, speakers, soundbars, and AV receivers, with support for advanced Bluetooth and Wi-Fi connectivity, and enough processing power to decode Dolby Atmos or DTS:X immersive formats and still allow for advanced voice assistant implementations. A single-chip amplifier helps manufacturers create portable and power-efficient speakers that can also offer differentiated audio quality.

This type of solutions enable many ODMs to bring to market products that they would not have the ability to fully develop in house, and make the most advanced technologies available on a large scale. Qualcomm's SoCs and full development kits have enough processing power to enable separating voice commands from other surrounding voices and sounds in noisy environments and support AI-based local automatic speech recognition for faster voice-recognition, offline commands and extended battery life.

And the QCS400 SoC series combines high-performance audio technologies to deliver evolved premium audio experiences, including speaker optimizations, beamforming, echo cancellation, and other advanced signal processing functions. Yes, those are available because Qualcomm wanted to combine low-power computational power with AI acceleration and low-latency audio distribution in a single chipset, optimized for voice recognition.

But audio developers can also leverage this power to build the intelligence they may find better suited to differentiate their "smart" designs. This is similar to soundbars with multiple drivers that are able to generate complex immersive experiences in Dolby Atmos or portable speakers that can generate completely different sound fields when placed at the center of the room for an omnidirectional experience, or combined in a pair or multiple connected speakers to cover larger spaces. All with a single SoC that is able to support up-to 32 channels of integrated audio processing.

And with advanced connectivity where Wi-Fi, Bluetooth, and Zigbee coexist, the QCS400 series is even able to support Qualcomm aptX Adaptive (see "Qualcomm aptX Adaptive Audio Codec" side text) designed to deliver low-latency audio streaming in the room or throughout the home. To meet the growing demands of audio processing, QCS400 SoCs include a quad-core CPU, Qualcomm Adreno GPU and Qualcomm Hexagon DSP with Hexagon Vector eXtensions. The Qualcomm AI Engine and

advanced AI hardware and software suite offers support for myriad neural network frameworks including Tensorflow, PyTorch, and ONNX, as well as the Qualcomm Neural Processing SDK for easily executing AI algorithms on device. The QSC400 SoC's Audio Development Kit (ADK) and the Qualcomm AI Engine can be combined by developers to fully develop the most sophisticated product vision.

The Qualcomm Smart Audio Platform 400, built around the Qualcomm QCS400 SoCs, is a complete solution designed specifically to help reduce the materials cost and development time associated with creating next-generation smart speakers. The Qualcomm Smart Audio Platform 400 is also a comprehensive, customizable development kit and example design.

Exploring Voice DSP Solutions

Due to the fast market evolution in smart speakers, the audio development platforms available are now more powerful than ever, and can be leveraged for audio development in general. As an example, CEVA, a leading licensor of signal processing platforms and artificial intelligence processors for connected devices, announced that DSP Concepts' industry-leading Audio Weaver tools and its TalkTo voice front-end now support the CEVA-X2 and CEVA-BX family of DSPs. Audio Weaver is used by many companies to develop advanced audio and voice applications for CEVA DSP-based designs. This combination creates a powerful solution to explore very sophisticated audio processing algorithms on areas such as speaker calibration and room acoustic compensation—including the ability to easily integrate the latest multi-microphone arrays.

Audio Weaver is a graphical configuration tool that enables rapid development of audio/voice DSP software, eliminating the need for native coding and kernel optimization. Audio Weaver is easy to use and enables the design, debugging, and tuning of audio processing flows in real-time, dramatically reducing product development time. Audio Weaver also enables product designers to build complex digital audio applications by simply choosing the required audio modules from a library.

TalkTo is a high-performance far-field voice front-end that leverages Audio Weaver for ease of integration and tuning. When used in conjunction with CEVA's audio and voice DSPs, product makers can fully exploit the processing performance of the DSP engine and create unique audio and voice features to differentiate their products. The CEVA-X2 and CEVA-BX family of DSPs are ideal for any power-constrained audio/voice use case and are widely adopted in SoCs for smartphones, smart speakers, voice assistants, action cameras, and more.

The combination of these systems enables the intensive sound processing applications that are required for real-time acoustic analysis and calibration, and to enhance quality audio playback and post-processing, including the possibility to explore on-device sound neural network implementations—an area that promises to differentiate new developments in this area, creating truly "intelligent speakers."

Another advantage for ODMs is the fact that CEVA and DSP Concepts both support a large third-party ecosystem of audio/voice software, hardware, and development tools.

Cadence HiFi 5

Another similar combination is offered by Cadence in cooperation with DSP Concepts. The latest Cadence Tensilica HiFi 5 DSP for audio and voice is the company's first IP core optimized for high-performance far-field processing and AI-based speech recognition processing. This fifth-generation HiFi DSP offers 2× audio processing and 4× neural network (NN) processing improvements versus the HiFi 4 DSP. This solution was specifically developed to answer the needs of advanced DSP algorithms that are needed to eliminate and reduce noise and echoes, or isolate sound signals such as voice in new applications that not require not only higher processing capabilities but also improved energy efficiency. Additionally, NN-based speech recognition algorithms are performing more tasks locally, rather than in the cloud, due to concerns of latency, privacy, and network availability.

Ambiq Micro was the first silicon licensee of Cadence's HiFi 5 DSP porting it to the company's

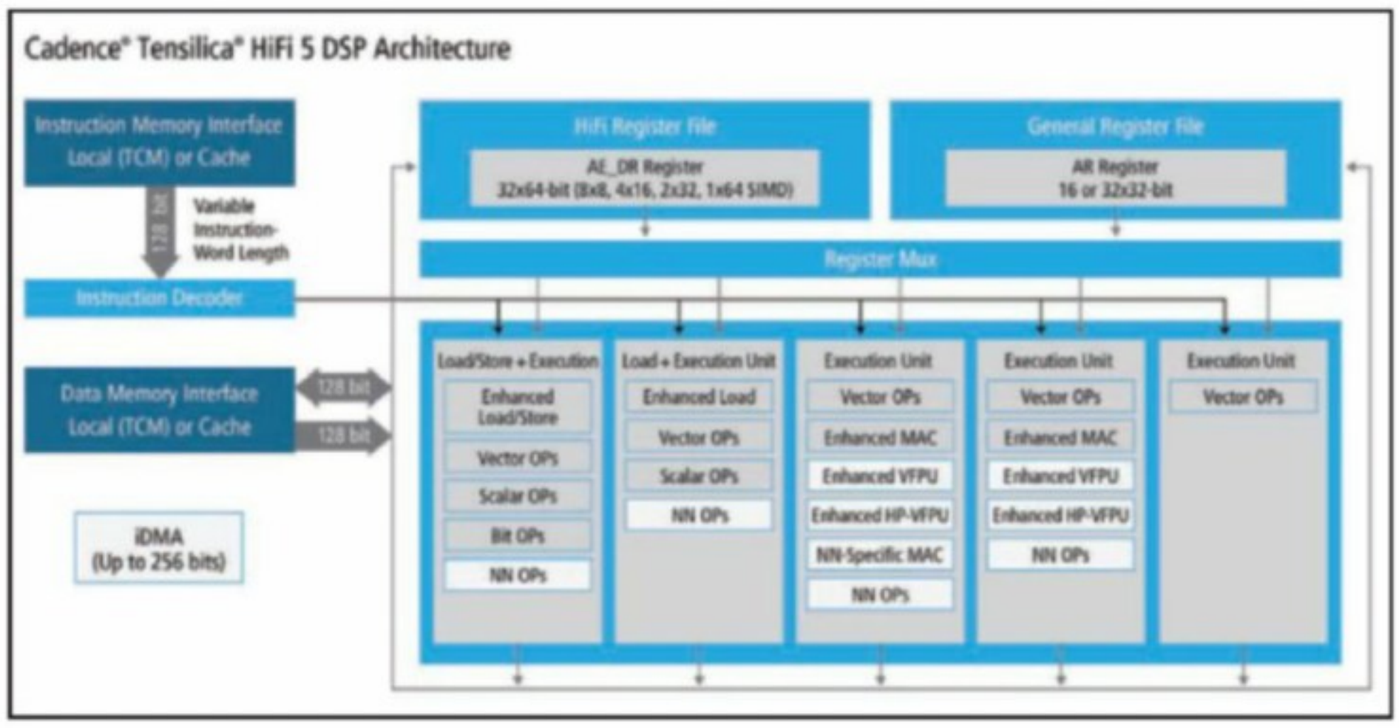

Processing cores such as the Cadence Tensilica HiFi 5 DSP and ready-to-use graphical configuration tools (e.g., Audio Weaver Designer) combine to create a powerful solution for new DSP-intensive audio applications, while at the same time supporting shorter and shorter development schedules.

The Microsemi (Microchip) Timberwolf Audio Processor family supports automatic speech recognition and audio development with advanced algorithms.

Subthreshold Power Optimized Technology (SPOT) platform that enables product designers, ODMs, and OEMs to leverage audio development tools from DSP Concepts and Sensory.

The Cadence Tensilica HiFi 5 DSP solution is also software compatible with the complete HiFi product line totaling more than 300 HiFi-optimized audio and voice codecs and audio enhancement software packages, meaning a large ecosystem of audio software is already available.

The Dedicated DSP Route

But of course, many developers still prefer to root their own creation at a more fundamental level, working at the chip (microprocessor, microcontroller, or SoC) level with many of the currently available software tools. In that aspect, the processing power available currently to audio developers is both an enabler and a challenge since the amount of DSP programming expertise and level of hardware integration is exponentially more complex as are the ambitions of such projects.

That is why we mentioned the StreamUnlimited and Qualcomm audio development kits and hardware modules. They are completely integrated solutions with ready-to-use functions and algorithms that can serve basic designs truly well.

Otherwise, options to create audio processing features can resort to using open hardware (common embedded platforms are increasingly powerful and can already rival many other alternative platforms for audio processing) or using FPGAs (Altera, Xilinx, Microsemi/Microchip) and the many available development tools.

Those who want to start a development based on FPGAs will also find many System-on-Module (SOM)-type of solutions, combining an FPGA SoC-type solution with an ARM-based processor, peripherals, and interfaces. The logic and programming tools will remain the same for FPGA development, but it becomes possible to easily implement readily available libraries for wireless audio and interfaces, while leveraging popular ARM-based development tools for embedded audio applications.

As an example, the ST-AudioWeaver free development tool from STMicroelectronics simplifies development using 32-bit ARM Cortex-M based STM32 microcontrollers and was developed in cooperation with DSP Concepts. The tool contains a library of hundreds of audio algorithms, with GUI for configuring and fine tuning, including royalty-free digital audio-processing modules for powerful filtering, array-processing, and noise-cancelling modules. The Audio Weaver Designer PC GUI is also included for free, which lets developers configure the modules and do fine tuning in the GUI or on prototype hardware. This creates a customized audio-processing engine that runs directly on the production embedded system with no further tailoring required. Users are also free to take advantage of DSP Concepts' commercial offers to scale and upgrade ST-Audio Weaver (e.g., by adding further audio modules or MATLAB API support).

STM32 microcontrollers are rich and efficient solutions for many audio devices, featuring networking and sensor fusion. There also are many low-power options to explore the energy efficiency of the ARM Cortex-M core architecture for battery-operated or self-powered devices. Plentiful audio features also add to the STM32 credentials, including I^2S and digital audio interfaces. The DSP core extensions currently supported in ST-AudioWeaver enable a single-chip solution capable of high-performance audio processing, which would otherwise require a separate DSP co-processor.

Other development options include to fully explore the FPGA ecosystem and all the extensive development tools that are available or go directly to dedicated DSP processor. There is a reason why many professional audio processors are based on Analog Devices' SHARC processors, because it can be easier to explore complex mathematical algorithms than on FPGA.

And, Analog Devices also offers its SigmaDSP audio digital signal processor range, complemented with the SigmaStudio graphical tool that doesn't require coding. Many SigmaDSP processors also integrate high-performance audio ADCs and DACs to build a full audio SoC.

Another company that we mentioned for FPGAs is also responsible for an interesting range of dedicated DSP processors for audio and voice. Microsemi, now part of Microchip, offers the Timberwolf Audio Processor family for automatic speech recognition and audio development. The Timberwolf platform is a third-generation product from Microsemi, supporting dedicated hardware accelerators. Its architecture also supports low-power applications in mobile devices and wearables. It is also a field-upgradable platform, which when combined with algorithmic firmware, enables beamforming, multichannel Acoustic Echo Cancellation (AEC), direction of arrival, and far-field microphone capabilities.

Microchip/Microsemi also offers its range of AcuEdge acoustic technology solutions for audio

processing and license-free, royalty-free intelligent audio IP algorithms, which can be combined with the company's range of voice and audio development platforms for smart speakers, soundbars, and any voice application—all featuring the Timberwolf family of processors. The Microsemi AcuEdge integrated algorithms support trigger word detection and voice commands for connected devices.

Developers working in acoustic compensation will also benefit from the interesting MiTuner Automatic Tuning Kit, a complete development package including hardware, software, and support for the automatic tuning of AcuEdge and Timberwolf processors. Microchip also offers development support to individual projects using the platform, including test and measurement, electrical analysis, mechanical analysis, and recommendations on transducers and plastics.

The SHARC Bytes Back

Among the recent solutions announced in 2019 for audio processing, Analog Devices launched its new SHARC Audio Module (ADZS-SC589), a hardware/software platform that facilitates efficient product prototyping, development, and production of a variety of digital audio products. Delivering an innovative combination of high-performance audio signal processing components and a comprehensive software development environment, the SHARC Audio Module is ideal for effects processors, multichannel audio systems, MIDI synthesizers, and many other DSP-based audio projects.

Traditional product development efforts usually require a significant investment in time and effort to select components, prototype a board, and develop a basic software structure before developers can even begin to address their own product needs. The new Analog Devices SHARC Audio Module significantly accelerates this process by providing an audio platform integrating a wide selection of analog and digital I/O with a fundamental software package and development environment. The SHARC Audio Module can be used as a self-contained system but can also be expanded, allowing for the creation of custom I/O and control surfaces.

The centerpiece of the SHARC Audio Module is Analog Devices' ADSP-SC589 processor, a state-of-the-art engine for audio processing, combining two 500 MHz SHARC+ DSP cores and an ARM Cortex-A5 core. Connected to the processor is a carefully designed set of system peripherals, including two 2 Gbit DDR3 memories and a 512 Mbit SPI flash memory, a SigmaDSP ADAU1761 audio codec chip for 24-bit up to 96 kHz audio and a AD2425W A2B multichannel audio interface, which can be used to connect the Class-D Amplifier Module, an A2B connected board that contains two Class-D SSM3582 amplifier ICs, providing four channels of audio output.

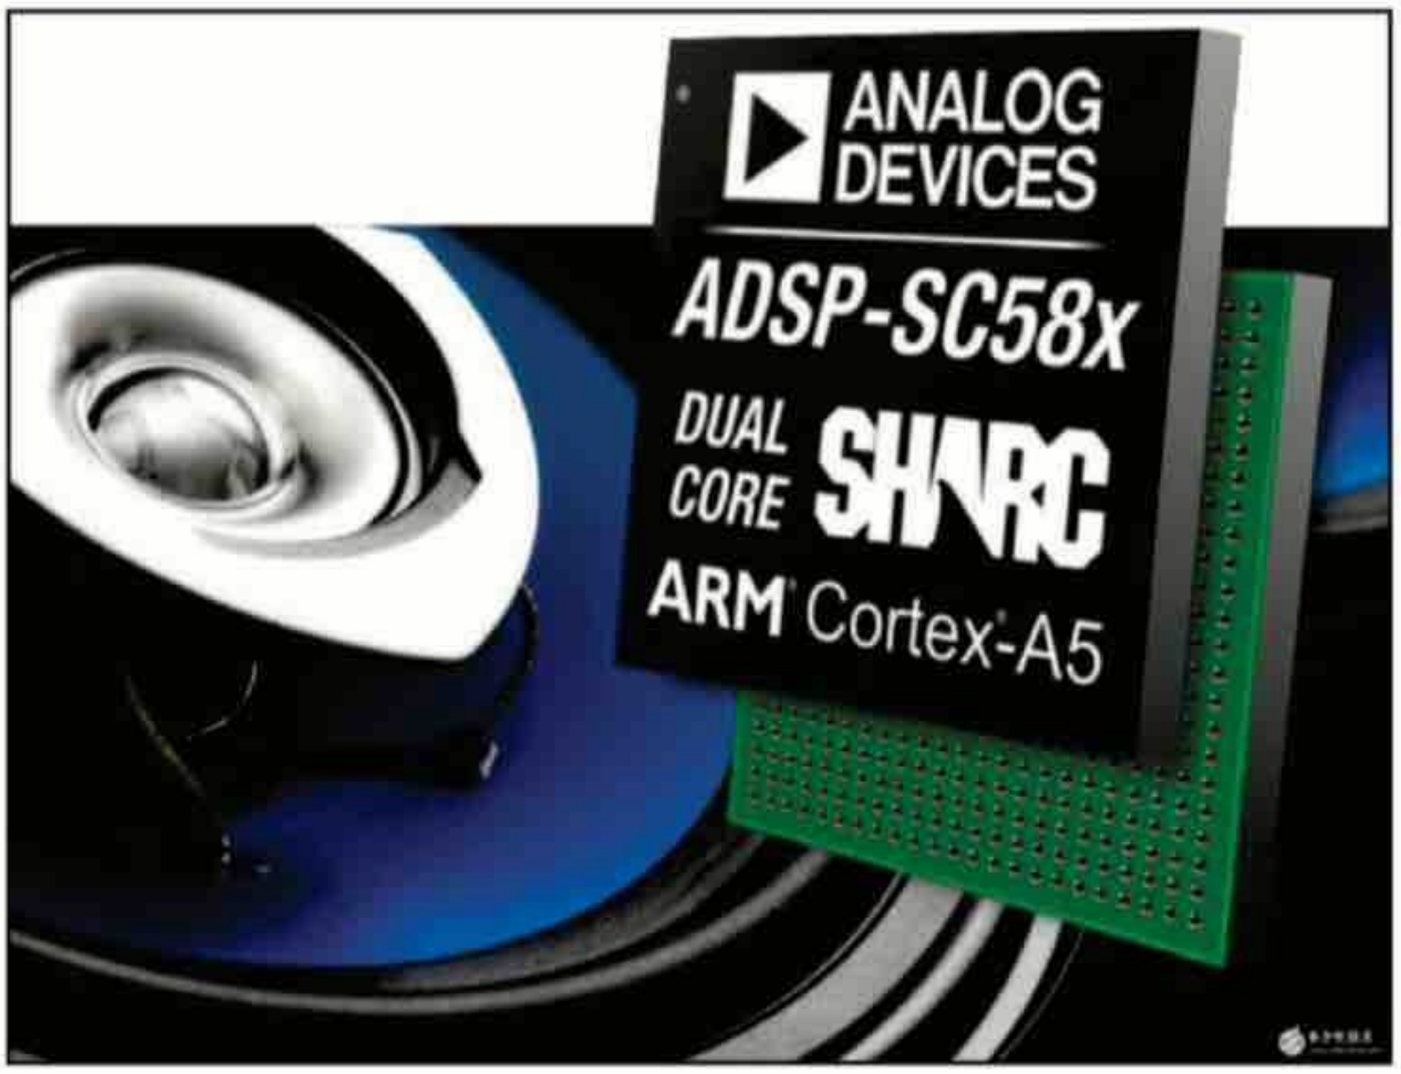

The new ADSP-SC589 is a state-of-the-art engine for audio processing, combining two 500 MHz SHARC+ DSP cores and an ARM Cortex-A5 core.

The complete system-in-a-module solution also includes Gigabit Ethernet and USB (OTG & HS) interfaces, S/PDIF and 1/8" stereo jacks, UART (for MIDI and more) and two 64-pin expansion connectors. This board is designed to be used in conjunction with the CrossCore Embedded Studio (CCES) development tools for advanced application code development and debugging.

In addition to the SHARC Audio Module, Analog Devices also offers expansion boards called "Fins," which provides additional functionality. The first Audio

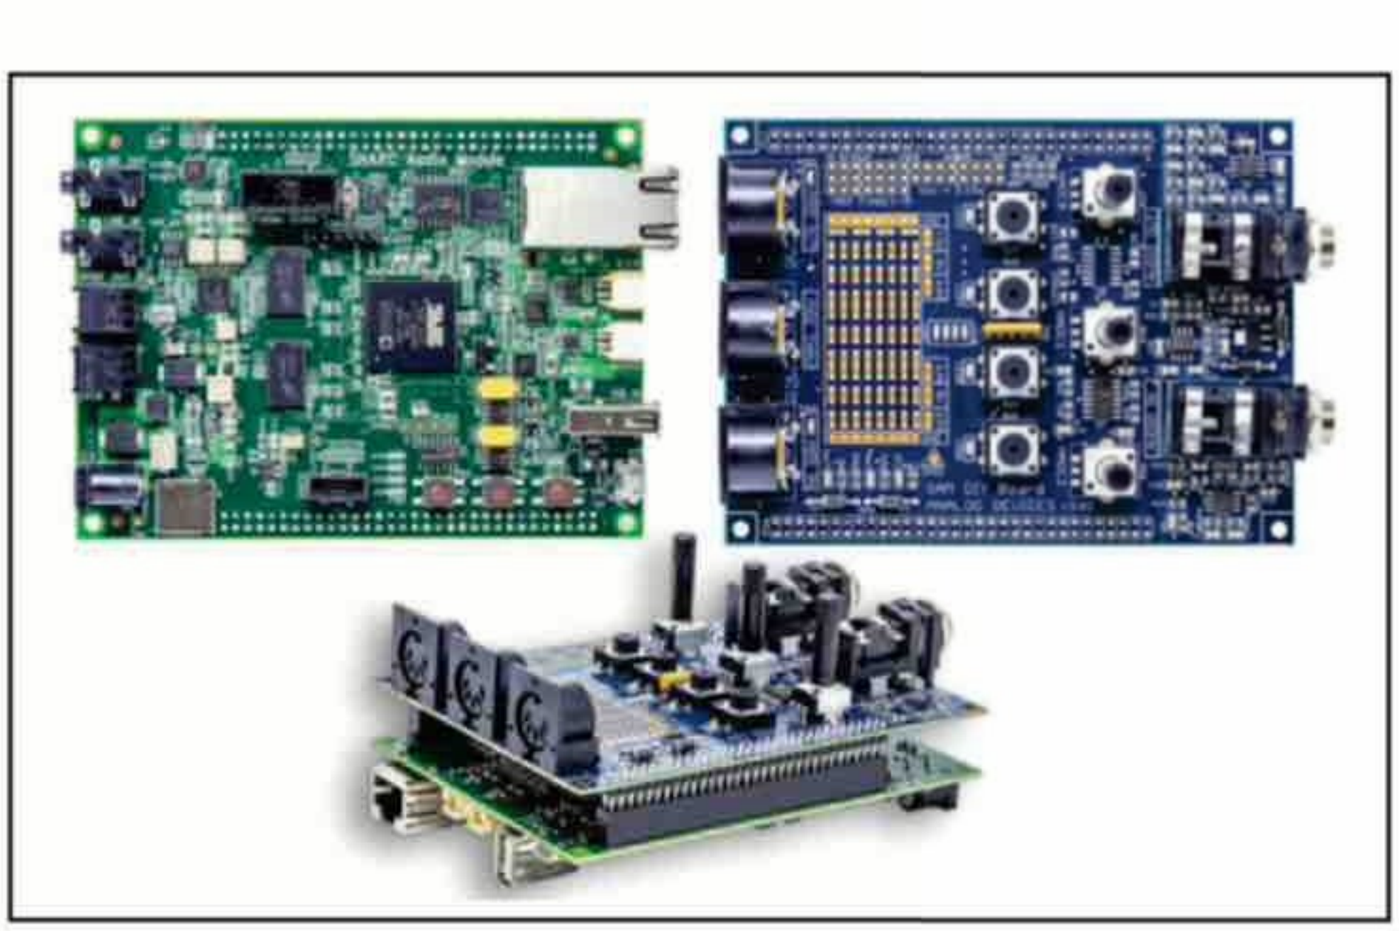

Development boards for the SHARC Audio Module allow developers to focus on their own algorithm and user interface development using highly optimized software and tools.

Project Fin (ADZS-AUDIOPROJECT) was a MIDI and musical instrument/FX daughter card that provides a control surface with 1/4" stereo I/O, additional connectivity, and expanded signal accessibility for the SHARC Audio Module. Since all of the major functions and GPIO of the board can be accessed via the two multi-pin connectors, developers can also easily design and build their own expanders using the available documentation.

In addition to traditional C and C++ development, the SHARC Audio Module supports Faust (Functional Audio Stream), an open-source, functional programming language, specifically designed for real-time audio signal processing and synthesis. Faust generates C++, as well as other target languages, for signal processing applications. The Faust library provides a rich set of audio DSP objects that can be used in creating DSP algorithms for SHARC-based platforms.

Reacting to this announcement from Analog Devices, audio development specialist Danville Signal Processing immediately announced its own dspSoM 589 platform—its most powerful DSP module to date—based on the ADSP-SC589. The SOM adds external memory, connectivity, and peripheral set to deliver an optimized development platform for audio designs and a production-ready target for manufacturing, including dual USB2 ports.

From Danville Signal Processing, there's also the dspblok 21489 lower-cost solution, based on the ADSP-21489 SHARC DSP Module, a highly integrated DSP platform supported by the company's family of dspblok power supplies, data converter modules, and digital I/O modules.

Development on these Danville platforms includes Analog Devices' Visual DSP++, CrossCore Embedded Studio and In-Circuit Emulators (ICE) for software, while peripheral integration can be initiated with the dspblok Development Board, which includes a power supply and computer connections.

Danville Signal is also a valuable partner for ODM and OEM audio companies that want to combine high-performance wireless audio features, as well as bass management and room correction.

And room correction is precisely one of the processing applications we will discuss in a future article, addressing also the evolution towards speaker self-calibration according to room response.

A Semi-Customizable Audio Processing Platform from Danville Signal Processing

The dspMusik 2/8 is a stereo in, eight out multichannel DAC platform suitable for high-performance DSP loudspeaker crossovers, music decoders, and many other high-quality audio algorithm applications, with an audio-optimized 32-bit/40-bit floating point DSP, based on Analog Devices' SHARC. An LCD/encoder provides source selection, volume control and access to special features and users can also use an Apple Remote for control. The dspMusik 2/8 platform supports digital audio up to 384 kHz sample rate and features balanced and unbalanced inputs and outputs with and an optional analog input module, XMOS USB Audio Class 2, and AKM DACs. Each part has been carefully selected to build a precisely integrated system.

The dspMusik starts by using low phase jitter clocks driving all the data converters. The optional ADC incorporates very high common mode rejection/low noise balanced inputs and DACs with minimum phase anti-imaging filters (desirable to avoid pre-echo effects).

The dspMusik 2/8 is supported by a custom version of DSP Concept's Audio Weaver programming software, enabling users to easily create their own signal processing chain by simply dragging and dropping highly optimized audio processing modules with EQs, crossovers, mixers, etc. Once the design is done on a PC, the file can be directly saved and run stand-alone on the dspMusik platform.

The dspMusik 2/8 platform is available in a variety of configurations including PCB assemblies where users provide their own housing, or it can be provided in an extruded aluminum enclosure. Panels can also be rebranded and Danville also provides customized display software that includes a company and product name as well as access to special features that may be part of a specific system design. OEM/ODM options include individual board assemblies, power supply adapters for Pascal and Hypex power amplifiers, a 16 channel expansion and alternative dspblok modules.

www.danvillesignal.com

Wireless Audio Technology and Trends

For audio brands looking at developing new-generation products, nothing is currently more relevant than taking into account the ability to wirelessly transmit audio signals. Yes, using Bluetooth, Wi-Fi, or dedicated radio-transmission approaches, there are many alternatives to consider when designing new audio solutions for the home, personal audio or even professional applications. Let's look at some of the products and platforms in wireless audio connectivity that have recently been announced and how audio designers can leverage the possibilities.

As the global audio market continues its growth trajectory, manufacturers can feel confident that their investments in R&D will pay off if they target the most popular categories. According to the Futuresource Consulting, the worldwide home audio hardware market grew 20% in 2018 to reach 154 million units, with a further 19% growth projected in 2019. The report, which considers wireless speakers, soundbars, A/V receivers, Hi-Fi systems, and dedicated speaker docks, also signals 5% revenue growth, with trade values reaching $14.5 billion in 2018. Wining product categories all involve designs that support wireless audio technologies.

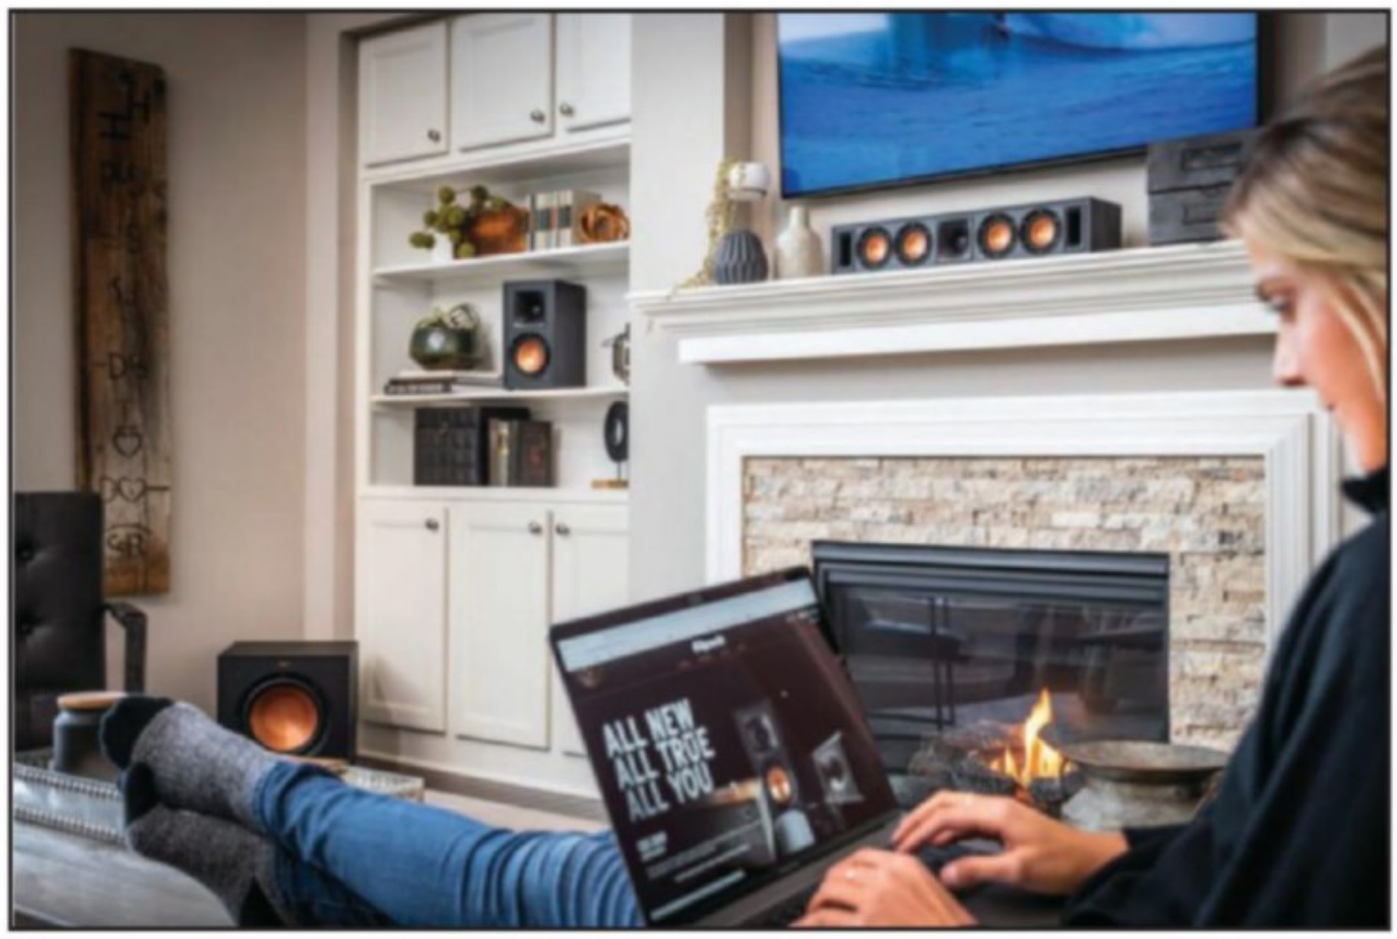

The much-anticipated WiSA Certified Klipsch Reference Wireless speaker series received multiple industry awards since its introduction at CES 2019 and is now available.

And according to the latest market report from GfK (www.gfk.com) released at IFA 2019 in Berlin, Germany, between January and June 2019, the global audio devices market, excluding North America, recorded total sales of €7.9 billion, or approximately $9 billion USD. Compared to the first half of 2018, this is an increase of 15%. Headphones and headsets are the key driver of growth, together with portable Bluetooth speakers, audio home systems, and loudspeakers for the home.

Most important, almost 50% of all personal audio devices sold are wireless, and home and multiroom devices are also becoming increasingly wireless and affordable for a mass market.

Sales of multiroom devices, which play different songs in different rooms, increased by 4% in the first half of 2019. Turnover with these Wi-Fi-enabled speaker systems was stable. Manufacturers are increasingly equipping their devices with one or even more smart assistants. This not only enables users to control the device via their voice, but also to integrate it into their in-home smart device ecosystem. This capability has increased consumers' willingness to invest in these devices. As a result, turnover for multiroom devices with a built-in smart assistant grew by more than 100% and the average sales price increased by 67%.

Although Gfk's consumer data suggests that people are in general very happy with the image quality on their newly purchased TVs, sound quality scores below-average on the sentiment score. This creates an opportunity for external speaker systems. So, while on the one hand consumers are spending less on their TV sets and average sales prices are decreasing, on the other hand, they are increasingly conscious about sound quality and are spending that saved budget on soundbars. As a result, the turnover for soundbars has been increasing year-on-year, and in 2019 it rose by 7%. Single-sold soundbars contributed with an increase of 13% and wireless subwoofers grew by 5%.

Note that soundbars are typically connected to the TV via HDMI, given the proximity and the convenience of the Audio Return Channel (ARC) and the ability to control the TV and soundbar with the same remote. That connection will keep improving with HDMI 2.1 and the high bitrate audio from Enhanced Audio Return Channel (eARC) technology. But otherwise, soundbars are essentially systems that wirelessly connect with multiroom systems and subwoofers. Even the most powerful soundbar models with Dolby Atmos and 4K/HDR compatible HDMI ports, usually have Chromecast built in, and Spotify Connect and AirPlay 2 support.

And the fact that soundbars sit at the center of home entertainment systems is the reason why all manufacturers have launched their own soundbars to support their own ecosystems. That was the reason why Sonos spent a good amount of time making sure

its Sonos Beam compact smart soundbar would tick all the right boxes, such as support for 80+ music services, support for Amazon Alexa at launch (with multiple voice assistants to come over time), Apple AirPlay 2, Ethernet, HDMI-ARC, and all the features of a smart speaker: becoming both a powerful soundbar, and a music and streaming center for the living room. More importantly, seamlessly connecting to Sonos multi-room set-ups and the Sonos SUB... wirelessly.

And speaking of Sonos, the company is using Skyworks Wi-Fi solutions to transmit high-quality sound throughout the home and power its latest products. Since Skyworks acquired Avnera Corp. in August 2018, the company is now able to complement its own range of wireless connectivity modules with Avnera's own ultra-low power analog circuits and offer a complete range of custom Analog System-on-Chip (ASoC) solutions for audio, voice, speech, sensor, and AI applications. These solutions are critical for manufacturers that are looking to a completely integrated approach to the smart home, combining hi-fi audio systems with intelligent audio technologies.

According to Juniper Research, the smart audio devices market is expected to exceed $10 billion by 2022, up from an estimated $2.5 billion in 2017, as consumers opt for homes and entertainment systems with increasing levels of connectivity. The proliferation of connected devices, coupled with artificial-intelligence voice assistant integration, easy smartphone interfacing and a growing number of mobile apps are driving this exponential growth.

A portable speaker from Sonos needs to support multiple wireless technologies in order to respond to current user expectations, supporting Sonos own multi-room platform and Apple AirPlay 2 over Wi-Fi.

Sonos Beam soundbar. An example of a system that needs to offer all connectivity options, wired and wireless, frequently being at the center of new home entertainment experiences.

Wireless Everywhere

As another example of the way wireless audio technologies are affecting the global audio industry in multiple applications segments, we can increasingly see the adoption of Bluetooth and Wi-Fi radio modules in musical instruments amplification, sound reinforcement systems, and even recording studio equipment.

Portable speakers for musicians or corporate events now require Bluetooth as a mandatory feature, simply because most of the audio sources also support Bluetooth streaming and many of the smartphones no longer have an analog output jack. That's why the new studio monitor company, California-based Kali Audio, has decided to create a Bluetooth Input Module as its third product. The Kali MV-BT allows users to connect Bluetooth devices (smartphones being the dominant ones) with professional speakers, mixers, controllers, and interfaces. The device is easy to pair, and features a large and precise volume control, balanced outputs and Qualcomm aptX support.

Another example comes from Nashville, TN-based Angry Audio, which offers the versatile Bluetooth Gadget, a Bluetooth solution to connect phones, tablets, and computers to any studio or professional audio production setup, including radio stations. Equipped with AAC and the virtually-lossless aptX codec algorithm, the Bluetooth Gadget was originally designed to help radio stations put Skype and smartphone calls on the air, but in the end, it offers a complete integration box for any Bluetooth device with high-quality balanced circuits.

Audio manufacturers are also starting to explore updated Bluetooth 5 specifications, in particular when combined with Class 1 antennas and radio modules, in order to provide longer and more robust transmission with longer distances—essential in fixed installation sound reinforcement systems

and portable PAs. Most audio manufacturers, particularly companies in the music instruments (MI) and professional audio segments, increasingly rely on cheap, ready-to-use Bluetooth modules and solutions, mostly still based on 4.1 specifications. That will be an area where we could see more companies improving its products, offering Bluetooth options that will be easier to pair with new devices and are more robust and reliable.

Bluetooth Codecs

The other thing that's changing in the Bluetooth domain is the increasing use of greater quality audio codecs, which enable 24-bit audio or close to CD quality, while being advertised for High-Resolution Audio (HRA) options.

Qualcomm's aptX technology is the oldest and most widely adopted solution in this domain, even if Apple only supports it in its mains-powered products, such as Mac computers, and doesn't support it in the iPhone or iPad. On the other hand, aptX and the higher quality tier of aptX HD became increasingly popular because most Android smartphones and headphones do support it.

Unfortunately, in the world of giant companies, adopting a rival technology is not something that happens often, and even though the mainstream Bluetooth music transmission codecs continue to be Fraunhofer's AAC, Qualcomm's aptX, and more recently Sony's LDAC, most consumers still listen daily to the basic SBC codec mostly without noticing, because not only both sides of the wireless chain need to support the same codec, but very frequently they need the user to configure a different codec option from the default selection.

Following the decision by Sony to introduce its LDAC audio coding option as a way to offer the transmission of audio content at 990 kbps over a Bluetooth connection, other Asian smartphone manufacturers thought it would a good idea to add their own Hi-Res "codec" options, even if most also support Qualcomm's aptX. To that effect, Huawei announced its own HWA HD Bluetooth audio transmission technology "with better sound quality than LDAC" at 900 kbps. Independently developed by Huawei, HWA uses what the company called the low latency and high-definition (LDHC) audio codec from Savitech, a Chinese technology company.

Samsung apparently also decided it would be a good idea to implement its own proprietary codec and offered the ultra-high-quality audio (UHQA), a scalable codec that was not so much designed to rival Sony's LDAC, but more with Qualcomm's AptX Adaptive (see "Qualcomm aptX Adaptive Audio Codec" side text).

The Kali MV-BT enables connecting Bluetooth devices in the studio with balanced outputs.

Cable Replacement

For more robust connections, Swiss-based Neutrik—the leading global supplier of professional entertainment connector products—is gradually expanding its Xirium Pro range of dedicated wireless solutions. Knowing how even professionals are increasingly ditching cabled connections for convenience, Neutrik designed the Xirium Pro range of transmitters and receivers for fixed installation, touring, and event applications, providing a dedicate "studio quality" digital wireless audio link.

Xirium Pro is a professional cable replacement system that is able to transmit high-resolution, full-bandwidth digital audio signals over long distances with low latency.

As a rapid, easy means of audio connectivity for indoor and outdoor events and venues, of the most demanding scale and complexity, the neutral system eradicates the need for time-consuming or

The search to extract the maximum possible quality for high-resolution audio with the restrictions of Bluetooth transmission has generated a new series of codecs from Qualcomm, Sony, and Huawei.

operationally impractical cable runs; including over very long distances. The system is in use, with increasing frequency, by hire companies and sound designers, across a wide range of events and installations. The signal transmission is based on Neutrik's proprietary Digital Wireless Audio (DiWA) technology, which sends and receives totally uncompressed 24-bit/48 kHz audio, using the license-free 5 GHz band. Up to six channels of audio can be transmitted to an unlimited number of receivers, up to a usable distance of 500 meters in line-of-sight and point-to-multiple point. Greater distances can be achieved, as required, with the use of repeater units. More recently, Neutrik even launched Xirium Pro transmitters and receivers with integrated Dante network connectivity modules, so the signals can be input directly into a network infrastructure without conversion to analog or signal degradation.

Microchip launched a fully-certified, Bluetooth 5-compliant System-on-Chip (SoC) supporting Sony's LDAC audio codec. The IS2064GM-0L3 SoC allows manufacturers to develop mass market higher-quality Bluetooth wireless applications.

The Neutrik Xirium Pro is a wireless cable replacement system that is ideal for long distances with low latency.

Wi-Fi Landscape

In the original Standards Review articles dedicated to Wireless Audio published by *audioXpress* in 2014 and 2015 (see Resources) we have fundamentally covered the most important wireless audio protocols and technologies that have since then been available to audio developers, including Qualcomm's AllPlay and DTS Play-Fi, or SONOS proprietary AES-encrypted peer-to-peer mesh network SonosNet technology.

Since then, a few things have happened. First and foremost, Google launched the Chromecast audio ecosystem, offering an inexpensive way to stream music across multiple products from various manufacturers, with multiple operating system support and support for high-resolution audio. The fact that Google offered its own hardware dongles for a really low price created considerable initial support for the technology, but didn't generate much sympathy in the audio industry in general. As a result, just four years past, Google has officially discontinued all Chromecast Audio hardware in January 2019, and decided to focus instead on selling its own Google Assistant speakers and support Google Home products or Cast-enabled speakers from partners.

At the other extreme, we have Apple with its AirPlay ecosystem, which has recently been upgraded to AirPlay 2, adapting the streaming protocol to better support whole-home and multiroom audio over Wi-Fi. Announced in 2017, the update received extensive industry support including initial commitments from brands such as Sonos, Naim, Libratone, Bowers & Wilkins, Bang & Olufsen, Dynaudio, and even Xperi/DTS, which made its own Play-Fi products compatible with AirPlay 2 and allows users to choose between a wide variety of compatible products and solutions. Airplay is one of the key technologies that Apple licenses to audio manufacturers to allow compatibility with Apple's devices and is certainly one of the most important components for home entertainment systems and speaker manufacturers.

So, essentially, when streaming audio wireless over Wi-Fi, it is not as if users are particularly worried about the available options. Apart from the multiroom protocol compatibilities that still oppose Sonos to other companies and continues to be aggressively protected, users don't really care too much what wireless audio technology they are using, and all

basically work flawlessly. In general, the industry is much more focused on the ecosystem integration for media streaming and interoperability with streaming services and voice interfaces.

If something will change in the future, it will be the consolidation of these technologies around Wi-Fi 6, basically allowing existing protocols for wireless audio streaming to work even better and seamlessly support streaming lossless and uncompressed high-resolution audio. The next-generation Wi-Fi connectivity technology, or 802.11ax, enables faster, more reliable connectivity for more devices but it will not motivate any companies to change existing streaming technologies.

Wi-Fi 6 networks and devices are simply able to handle more traffic, more efficiently, thanks to multiuser, multiple-input multiple-output communication (MU-MIMO) on 2.4 GHz and 5 GHz bands, Orthogonal Frequency Division Multiple Access (OFDMA) allowing the efficient management of available spectrum, and 1024 quadrature amplitude modulation (QAM), allowing for a 35% increase in data rate over Wi-Fi 5. This will be critical in homes where users are simultaneously streaming 4K/HDR video content from different sources to different devices—all wirelessly—and for supporting new immersive audio streaming services now being introduced with Dolby Atmos and other object-based immersive formats.

That is precisely why the audio industry does have a unique opportunity to get together around the Wireless Speaker and Audio Association (WiSA) technology as a standard solution for wireless audio in the home.

The WiSA Momentum

Already encompassing more than eight years of technical research and industry-wide promotion, the Wireless Speaker and Audio Association (WiSA) effort was built around a simple concept to create a wireless audio standard with improved audio transmission for multichannel (home theater) applications. WiSA transmits 24-bit audio at 48 kHz/96 kHz using no compression, with all the speakers in the system certified to be synchronized within 1/1,000,000 of a second even when using all eight channels, and the total system adding less than 1/10th the latency of a Bluetooth audio device—5.2 or 2.6 ms depending of the source sampling rate. The technology was designed to be completely plug-and-play, being able to automatically recognize audio configurations from 2.0 to 7.1 and even 5.1.2. More important, the WiSA transmission doesn't depend upon Wi-Fi and uses Dynamic Frequency Selection (DFS) channels between 5.2 and 5.8 GHz. These channels—previously reserved for weather and military applications—are unlicensed, but not unregulated. Consumer devices can use these DFS channels, as long as conflicts with these important services are actively avoided and the WiSA implementation makes sure a new channel is selected in the event of any interference.

The WiSA technical work for the association is backed by Summit Wireless Technologies, the founding member, now supported by more than 60 leading

You know the market trends are clear when even companies such as McIntosh announces a product like the RS200 ($3,000 USD), featuring adjustable EQ and support for Apple AirPlay 2; Bluetooth 5.0 with support for AAC and Qualcomm aptX HD; and DTS Play-Fi whole-home wireless audio.

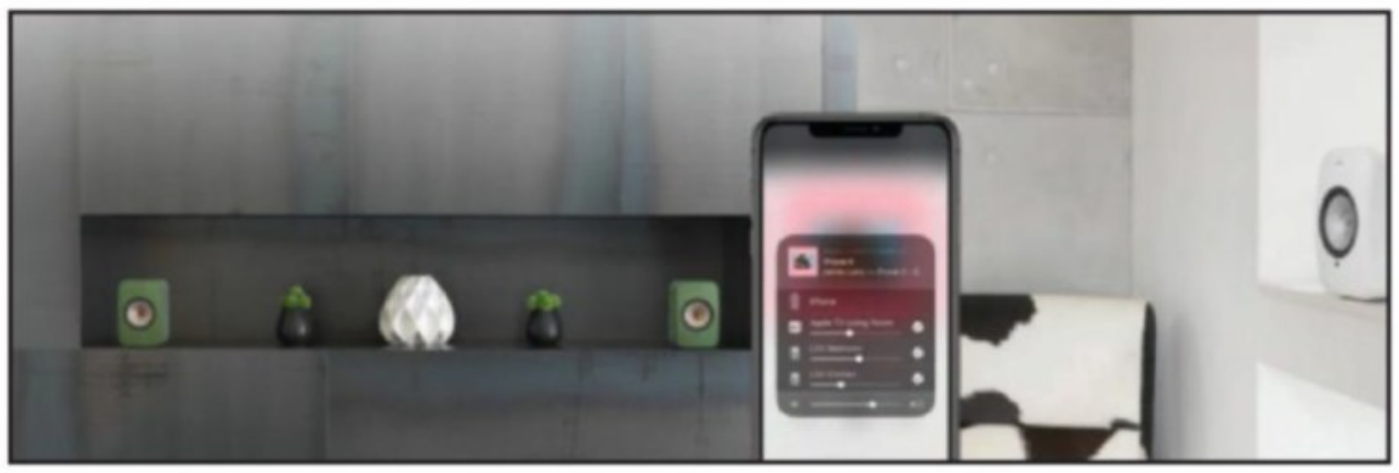

KEF LSX wireless speakers featuring AirPlay 2

Tony Ostrom is the president of the Wireless Speaker and Audio Association (WiSA).

Market Update—DSP and Wireless Audio Applications

WiSA Ready TVs are now available in the market and are a critical incentive to generate consumer awareness and demonstrate there is an easy and convenient way to set up 5.1 or 7.1 home theater systems without wires.

Axiim Link USB wireless audio transmitter and USB dongle allow WiSA Ready products to transmitted audio to WiSA-certified speakers.

consumer electronics brands. All products that adopt WiSA technology—WiSA Certified and WiSA Ready components—work together seamlessly. The scope of members now includes all main global TV brands, shipping WiSA Ready TVs, which are "ready" to transmit audio to WiSA Certified speakers when a WiSA USB Transmitter is plugged in and a user interface is activated through an app or product like LG TVs.

"WiSA's rapid adoption as the de facto standard in wireless, high-resolution audio is apparent with 20 million WiSA Ready TVs anticipated in market in 2020," sates Tony Ostrom, president of WiSA. "Speakerless WiSA Ready TVs are anticipated to enter the market in 2020, and when connected with a WiSA USB Transmitter, will seamlessly deliver immersive audio to WiSA Certified soundbars and speaker systems. WiSA is poised to play a critical role as the industry moves in this direction."

WiSA's recent growth was heavily fueled by the Association's WiSA Ready platform, which clearly identifies entertainment sources—such as TVs, Xbox consoles, and Windows 10 Gaming PCs—that are equipped to deliver wireless audio to WiSA Certified speakers connected with a low-cost WiSA USB transmitter.

With this level of industry support WiSA's momentum has never been greater. The number of companies joining the association is increasing every month. And WiSA is already defining an ambitious technology roadmap, maintaining interoperability with technologies such as AirPlay 2 and existing audio multi-room systems at its core. All reasons to keep a close eye on WiSA developments as the most important technology for wireless audio for the foreseeable future.

Resources

"HDMI 2.1 Specification Sees Strong Adoption Numbers and New Announcements at CES 2019," *audioXpress*, January 2019, https://www.audioxpress.com/news/hdmi-2-1-specification-sees-strong-adoption-numbers-and-new-announcements-at-ces-2019

J. Martins, "Wired-Quality Sound Without the Wires," *audioXpress*, October 2014.

———, "Wireless Audio (Part 2): Wireless Personal Audio Networks," *audioXpress*, November 2014.

———, "Wireless Audio (Part 3): Seeking Better Quality Wireless Audio Solutions," *audioXpress*, December 2014.

———, "Wireless Audio (Part 4): Are We There Yet?," *audioXpress*, January 2015.

———, "Wireless Audio (Part 5): Approaches to Wireless Audio Implementation," *audioXpress*, February 2015.

Sources

Analog Devices | www.analog.com
Cadence Design Systems | www.cadence.com
CEVA DSP | https://www.ceva-dsp.com
Danville Signal Processing | https://www.danvillesignal.com
DSP Concepts | www.dspconcepts.com
Futuresource Consulting | www.futuresource-consulting.com
GfK SE | www.gfk.com
Qualcomm | www.qualcomm.com
Meridian Audio | https://www.meridian-audio.com
miniDSP | https://www.minidsp.com
The Neutrik Group | www.neutrik.com
StreamUnlimited | www.StreamUnlimited.com
Skyworks | www.skyworksinc.com
The Wireless Speaker and Audio Association (WiSA) | www.wisaassociation.org
Xilica | Networked Audio DSP & Control Solutions | https://xilica.com

Qualcomm aptX Adaptive Audio Codec

In 2018, Qualcomm promoted the very first demonstration of its Qualcomm aptX Adaptive technology to audio manufacturers. This addition to Qualcomm's audio portfolio combines enhancements in audio transmission to support an improved wireless listening user experience for today's most popular and demanding applications including mobile gaming, watching video and music listening.

The next-generation audio codec takes the existing and established aptX technology and makes it dynamically adjustable. In that respect, it is a breakthrough audio codec technology that combines robustness, premium audio quality, scalability and low-latency, designed to provide a single flexible solution for virtually all key wireless audio applications.

The new aptX Adaptive technology was engineered to automatically adjust to provide optimum audio quality or latency depending on the type of content being played on the device while also taking into account the external RF environment to support a robust connection.

The codec recognizes the type of content that is being played back and, depending on the content, delivers audio using low-latency aptX, aptX, or aptX HD, combining the features of all of these solutions. The technology is backwards compatible with aptX and aptX HD enabled devices, supports compression ratios from 5:1 to 10:1 and supports 24-bit audio at 48 kHz with 276 kbit/s and 420 kbit/s with bit-error resilience. According to Qualcomm the new audio codec also supports legacy Qualcomm TrueWireless Stereo and dual-mono Qualcomm TrueWireless Stereo modes.

"Most audio codecs available today are largely static in nature with fixed bit-rates, which can lead to wireless connectivity issues in challenging RF environments," explains Jonny McClintock, director, product marketing, Qualcomm Technologies International, Ltd. "Many of these codecs were designed for music listening and are simply not suitable for low-latency gaming and audio/video applications. aptX was one of the first technologies to really help make a difference to the Bluetooth music listening experience and we've continued working hard to create a technology that will help transform the listening experience for the next generation of consumers who want to replace the wire completely. With this new product, listeners don't have to do anything while aptX Adaptive dynamically adjusts performance to best deliver exceptional audio whether the user is playing games or listening to a song."

Since media content can vary between high resolution (24-bit/96 kHz) and CD (16-bit/44.1 kHz), at a low bit rate (280 kbps) aptX Adaptive can achieve approximately 40ms latency—the actual codec latency is under 2 ms and it's the Bluetooth packet structure and buffer settings that attribute much of the system latency. At a high bit rate (420 kbps) aptX Adaptive supports high-definition Bluetooth audio at a significantly lower bit rate than previous HD codecs like aptx HD (576 kbps) or LDAC.

In Application

The aptX Adaptive decoder, for devices such as headsets, headphones and speakers, is available on the Qualcomm CSRA68100 and the Qualcomm QCC5100 series Bluetooth Audio SoCs and a version of the aptX Adaptive encoder for devices such as smartphones and tablets runs on Android Pie.

The flagship Qualcomm Snapdragon 855 Mobile Platform also supports aptX Adaptive and TWS+ wireless audio experiences. As audio manufacturers are announcing new or updated versions, the first examples of products implementing aptX Adaptive have been announced.

Bowers & Wilkins was the first company to confirm the implementation with its latest headphones—two new in-ear, one over-ear, and one on-ear model—all adopting the latest available platform and technologies from Qualcomm, including the latest aptX Adaptive Bluetooth transmission. The range includes the top-of-the line PX7 Wireless over-ear model, based on the ultra-low power Qualcomm QCC5100 Series Bluetooth Audio SoCs, and which includes ANC that runs at 16 times the sampling rate to ensure the best performance with no degradation to the audio quality.

Qualcomm | www.qualcomm.com

Heading the new 2019 headphone range from Bowers & Wilkins is the PX7 Wireless over-ear design ($399.99 retail) with Active Noise Cancellation and aptX Adaptive Bluetooth transmission.

Keeping the Sound (or Noise) In or Out

In this article, our columnist tackles sound isolation and ways in which noise can be controlled, taking into account the differences between outdoor and indoor conditions.

By
Richard Honeycutt
(United States)

One of the important and challenging aspects of acoustical design is sound isolation. Whether you're trying to keep the sound of trains, motor traffic, airplanes, and neighbors out of your auditorium, sanctuary, or bedroom; or trying to keep the sound of your stereo or lounge band from annoying neighbors, there are important principles to keep in mind. As with so many acoustical issues, "what everybody knows" may not be based in fact.

Outdoor Sound Transmission

The first principle to remember is that in noise isolation, distance from source to listener is your friend: In a reflection-free environment (outdoors, away from large walls, etc.), sound pressure level (SPL) decreases by 6 dB for every doubling of distance. So if a band is playing in an open field in a park, and the average SPL 10' from the band is 96 dB, at 20', it will be 90 dB. To get down to a level that is comfortable for conversation (~65 dB), you have to move 320' away.

Several factors affect the transmission of sound outdoors. Perhaps the least commonly known is refraction. If the air through which a sound wave passes is uniform, the wavefront will move in a straight line outward from the source. If the wave gradually encounters a medium whose density and/or stiffness is different from those of the medium from which the sound originated, the sound will proceed with a different direction from that of the original wave.

An example of this effect occurs in very large outdoor concerts, where changes in air temperature with elevation above the ground cause the speed of sound to change. The result is that the sound from the stage can curve up, missing the rear of the audience; or it may curve down, making the sound in the rear unusually loud (see **Figure 1** and **Figure 2**). This curving of the sound's path is called refraction, and is analogous to the refraction of light waves by eyeglasses, or to the apparent bending of a fork that is partially submerged in a glass of water.

In both of these cases, the refraction is caused by the wave passing from one medium to another in which the wave travels at a different speed. Refraction of sound in air occurs because of a change in the properties (temperature, humidity, etc.) of the air, resulting in a change in the speed of sound in the air as the wave progresses.

Another factor in outdoor sound transmission that is more commonly misunderstood than

understood is sound absorption by trees and foliage. Although the visual barrier trees provide may lead us to expect that trees will block sound, in fact, they provide only limited benefits. Some real estate zoning regulations require a certain width of trees or bushes, ostensibly as both a visual and an acoustical barrier. In fact, trees provide almost no reduction in the sound level at frequencies below 500 Hz. And from 500 Hz to 2,000 Hz, the main sound reduction is due to ground effect, resulting from the roots making the ground more porous. Typical values of ground-effect attenuation in this frequency range vary from about 5 dB at 15′ distance to perhaps 10 dB at distances of 30′ or more. At frequencies above 2 kHz, foliage does provide some attenuation, but this is limited to about 1 dB per 30′ of distance and a maximum of 10 dB for distances greater than 300′.

An example of the difficulties attending sound isolation outdoors is a concert venue in which the sound level was kept at safe values for listeners' hearing protection in the venue, and a 30′ high barrier wall flanked by a small forest was expected to prevent the sound from offending neighbors.

In fact, these methods did prevent near neighbors from being annoyed by the sound, but on numerous occasions, neighbors located about 3 miles away experienced noise levels near 65 dB in their front yards. The culprits were refraction of the sound above the barrier wall and the ineffectiveness of the forest in attenuating sound. The sound refracted upward over the warmer audience, then over the wall, and downward beyond the barriers.

Under certain weather conditions, the sound, after traveling through the limited absorption provided by the air well above the ground, refracted downward into the residential neighborhood. (A very interesting aspect of this case is that when the consultant who investigated it tallied the number of noise-annoyance complaints, he found significant differences depending upon the type of music being performed. Perhaps surprisingly, Southern Rock acts produced higher measured sound levels but fewer noise complaints than several other types of music!)

Indoor Sound Transmission

When the listener is indoors, other factors come into play. The first is the ability of the building to block the transmission of sound. This ability is quantified in two ways: for walls, it is specified as sound transmission class (STC), and for floor-ceiling systems, it is specified as impact isolation class (IIC). Both are specified as numbers, with a higher number indicating greater acoustical insulation. Typical single-family residential construction provides walls with an STC of 30 to 40, and typical floor-ceiling systems provide an IIC in the 25 to 35 range if the floor is hard. Higher values of STC and IIC are designed into partitions and floor-ceilings between units of multifamily structures: a minimum of 50, and up to 60 in wood frame structures, or higher in concrete construction.

Strictly STC only applies when both the sound source and the listener are in the same building. If the sound source is outdoors, the Outdoor-Indoor Transmission Class (OITC) is used. The OITC better characterizes the performance of a sound barrier at low frequencies, and outdoor noises are more heavily balanced toward low frequencies than are indoor noises. For impulsive noise such as heel-strike noise on wooden floors in the room above, Impact Isolation Class (IIC) is the applicable metric.

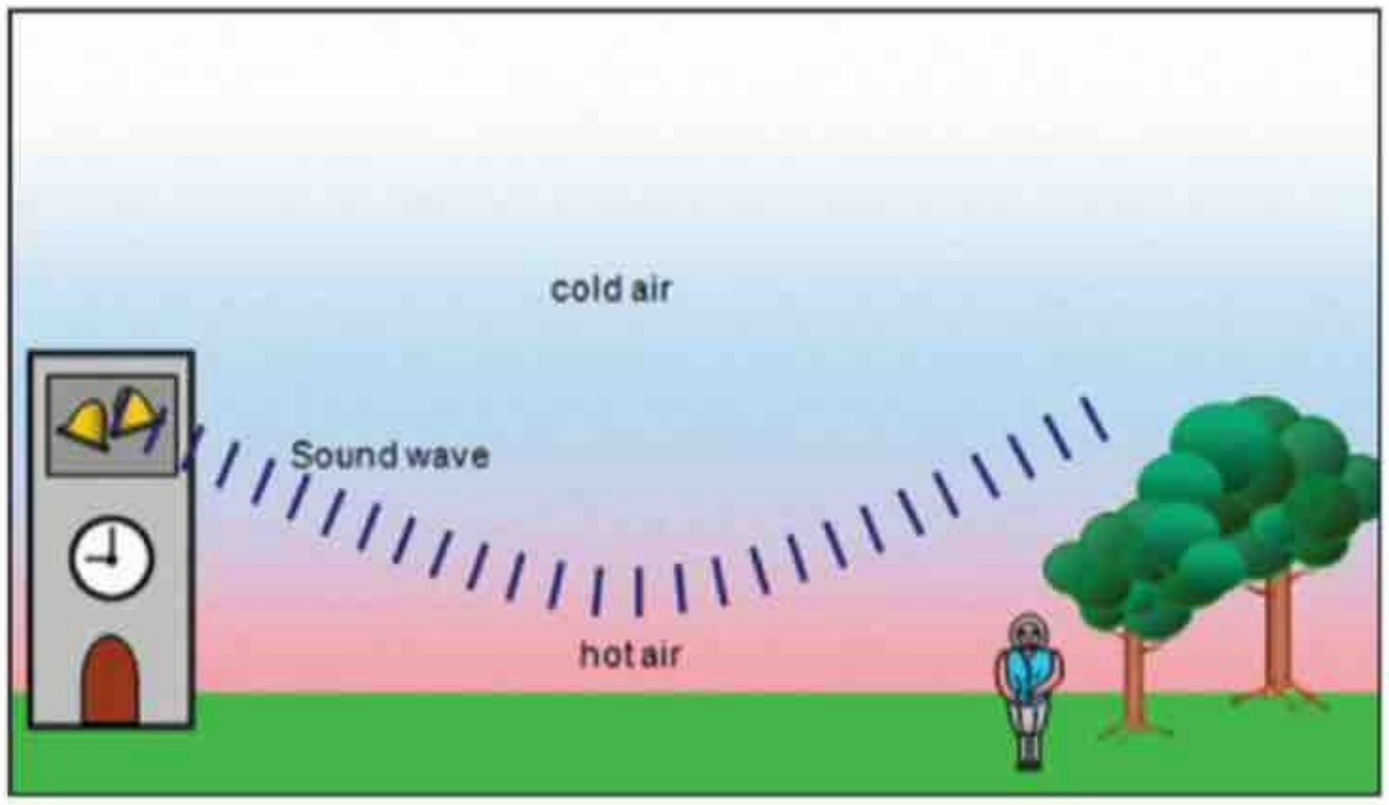

Figure 1: Refraction of sound when air is cooler at greater heights (Image courtesy of schoolphysics)

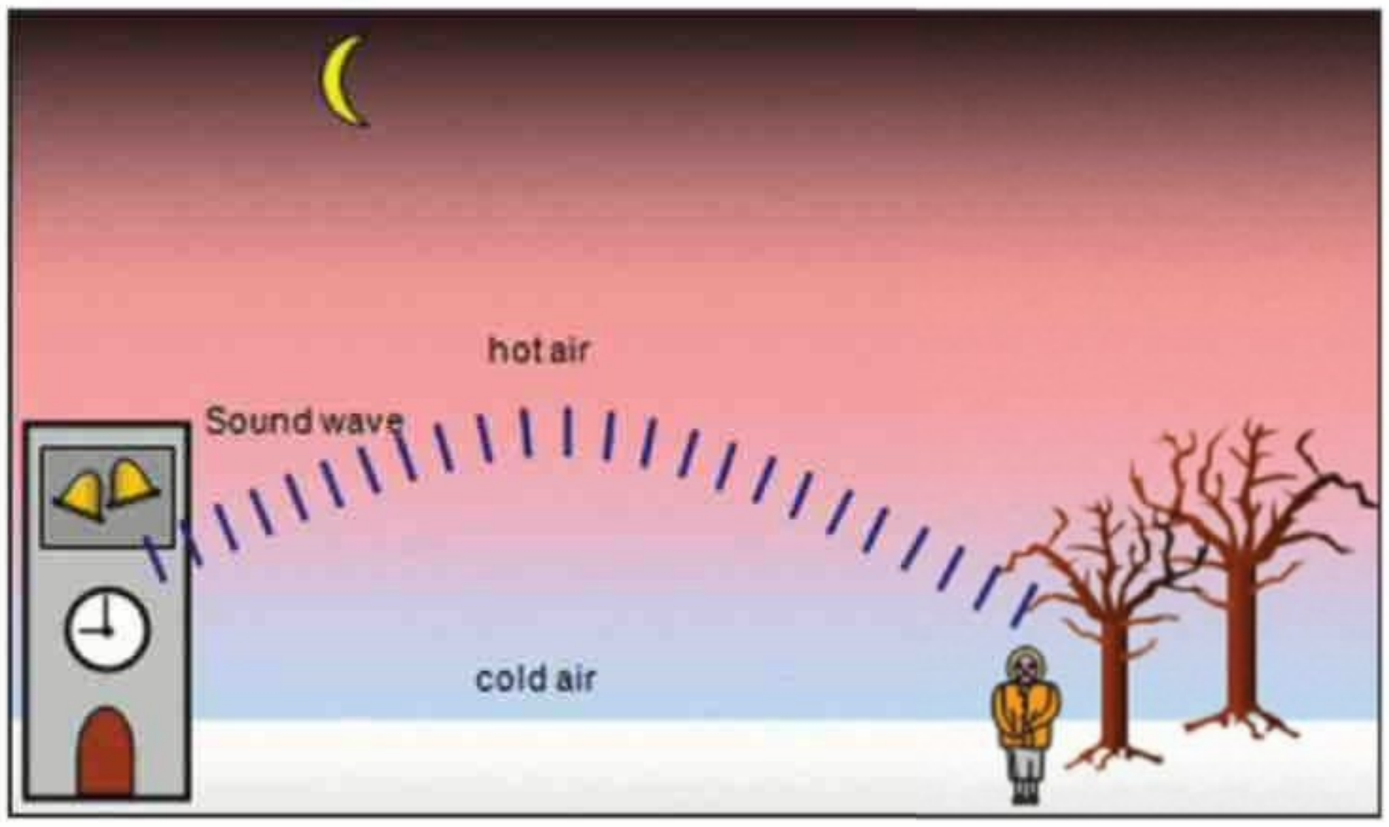

Figure 2: Refraction of sound when air is warmer at greater heights (Image courtesy of schoolphysics)

Photo 1: QuietRock 510 is a 1/2" cost-effective sound damping gypsum panel. (Image courtesy of QuietRock)

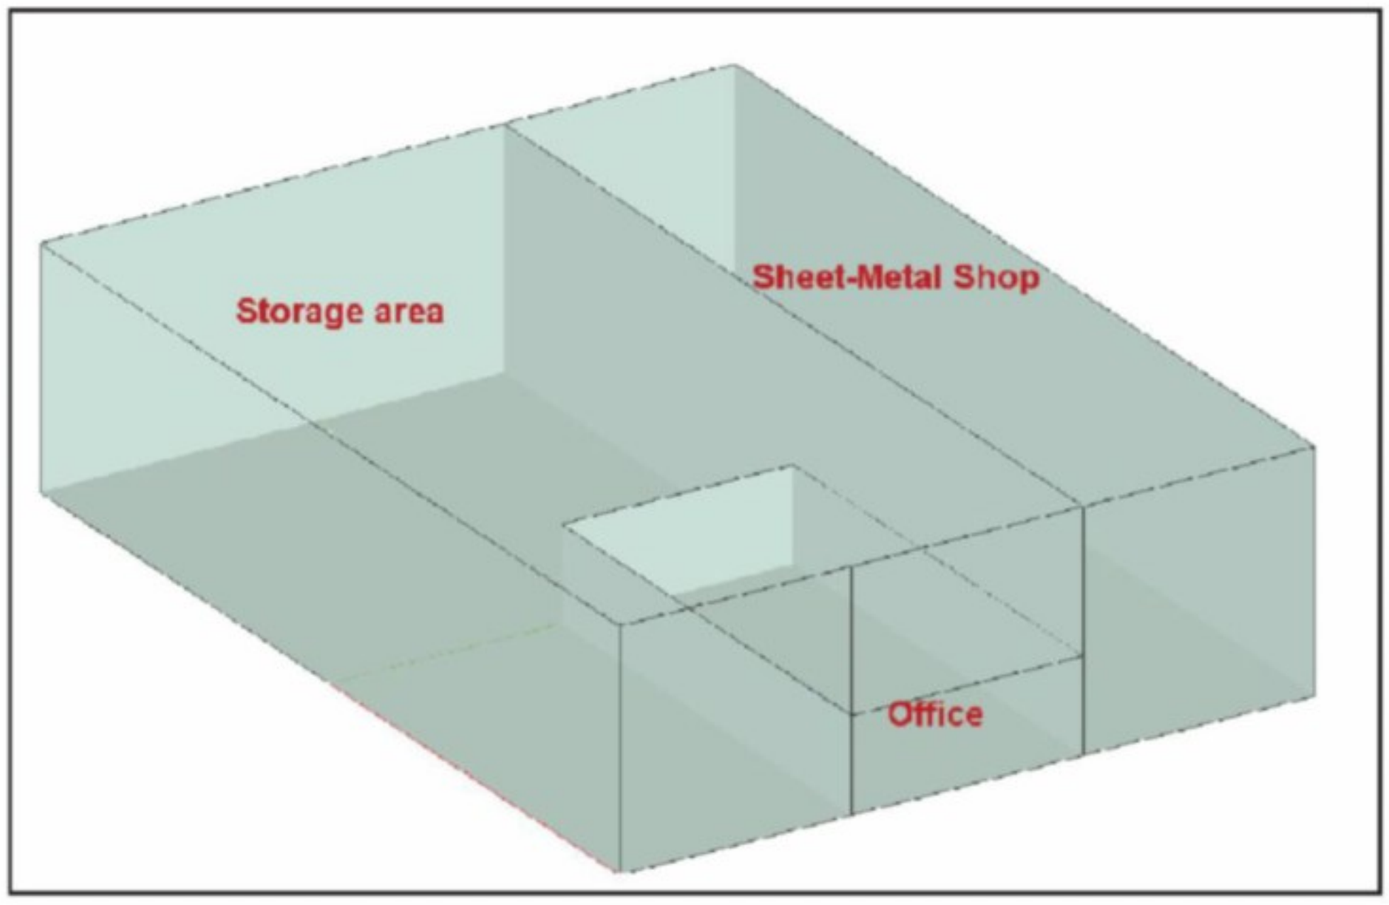

Figure 3: This is the building layout of a client who complained of office noise.

Photo 2: Mass-loaded plastic sheeting can be used to improve noise isolation on all room boundaries. (Photo courtesy of Acoustiblok, Inc.)

Isolating Sound

Contrary to popular opinion, soft fuzzy materials such as fiberglass and rock wool provide little attenuation of transmitted sound. The primary factor affecting STC or OITC is the mass per unit area of the barrier material. This conforms to the mass law, which states that the isolation of a partition is proportional to the distributed mass (or "surface weight"). A 3"-thick concrete wall having a surface weight of 39 pounds/square foot (PSF) has an STC of about 47; a 6"-thick concrete wall having a surface weight of 78 PSF has an STC of about 53—higher by 6 units. (Note that despite numerous Internet pages stipulating that STC is measured in decibels, the current standard definition refers to "integer units," not decibels.)

Obviously, there is a limit to the STC improvement that can be obtained by increasing surface weight or thickness. If you need to isolate the sound generated in room A from room B, and the intervening wall is made of a layer of 1/2" gypsum board on each side, framed wooden 2×4s on 16" centers, your STC will be about 32. If you add fiberglass batts in the cavity between the gypsum panels to reduce the reflection of sound back and forth inside the wall, the STC rises to about 35. Doubling the gypsum to a 1" thickness on each side brings us to an STC of about 40. Getting an additional 5 to 6 units of STC would require that we double the gypsum thickness again—to 2" on each side. Clearly, we are encountering diminishing returns.

One way of increasing STC that does not depend upon the mass law is to use constrained-layer damping. This involves inserting a layer of viscoelastic material between two massive rigid layers of wall material. An example is the use of QuietRock (see **Photo 1**). These panels are composed of multiple layers of gypsum laminated with viscoelastic adhesive. If the STC 32 wall described above were built with one side using QuietRock 510, the STC would go up to 50.

A similar approach for improving sound isolation in existing buildings involves attaching normal gypsum to existing walls using viscoelastic adhesive such as Green Glue (https://www.greengluecompany.com/data-resources/test-data). Adding a layer of 1/2" gypsum laminated with Green Glue to an existing 1/2" gypsum panel on an STC 32 wall with a common 1/2" gypsum panel on the other side will raise the STC to 42.

One of my clients had his office just across a wall from a sheet-metal shop, and the noise of the metal-cutting saws made teleconferencing quite difficult.

The wall was constructed of metal studs on 24″ centers, with a 5/8″ layer of gypsum on both sides. Adding two layers of 5/8″ gypsum laminated to the existing wall with Green Glue significantly reduced the noise in the office.

In this case, however, there were other problems. The gypsum was not sealed to the corrugated steel roof deck, so we added acoustical caulk to prevent sound leakage through the gap. This improved STC of the wall assembly by about 5 units. **Figure 3** shows the layout of the rooms involved. Notice that much more partition-wall area in the storage area is exposed to the noisy sheet-metal shop between the sheet-metal shop and the office. Therefore, the noise level in the storage area was quite high, and quite a bit of noise entered the office from the storage area through the office ceiling, which was made of mineral-fiber acoustical tile having a surface weight of only about 1 PSF. (This route for noise entry is called a flanking path.)

We had two options: we could have extended the walls between the storage area and the office to the ceiling and caulked the joint at the top, or we could apply mass-loaded plastic sheeting (https://www.acoustiblok.com/acoustiblok-product-lines/acoustiblok-soundproofing-material) over the office ceiling. The latter was less expensive, so that's what we did. **Photo 2** is an example showing mass-loaded plastic sheeting. Other significant factors in sound isolation include: frequency dependence, low-STC and leakage-prone features, and type of sound exposure.

High-frequency sound is easier to isolate than low-frequency sound, because it is more easily absorbed by mechanisms other than the mechanical energy expended in vibrating massive walls. However, a "double-leaf" gypsum wall (gypsum on both sides of a wood frame) is susceptible to a resonance between the mass of the gypsum and the springiness of the air in the cavity. This "mass-air-mass resonance" can decrease STC over a certain frequency range. Also, sound waves striking the gypsum at a certain angle will produce a bending wave in the gypsum of the same wavelength as that of the sound wave in air. At the coincidence frequency, energy is very efficiently coupled from the air to the gypsum, and below the coincidence frequency, STC is decreased, producing a coincidence dip. **Figure 4** shows the effects of the mass-air-mass resonance and the coincidence dip.

Windows and doors invariably have lower STC than the walls in which they are mounted, both because of their lower surface weight and because of the likelihood of leakage around them. Special constructions have been developed to overcome these weaknesses.

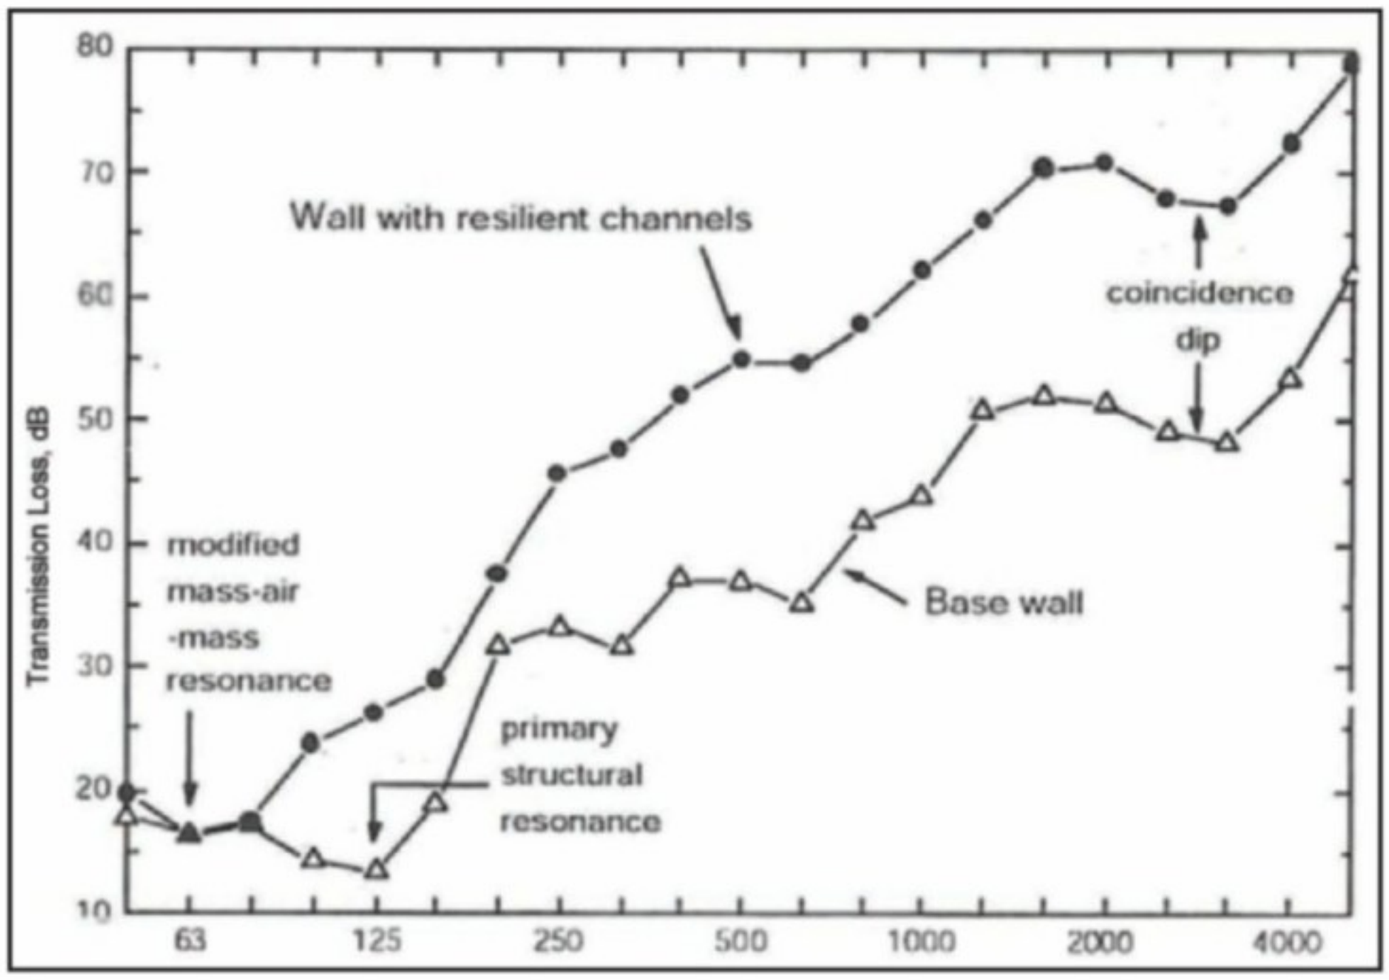

Figure 4: This graph shows the effects of the mass-air-mass resonance and the coincidence dip.

Large rooms with no ceilings and metal roofs are vulnerable to noise problems as a result of both the low surface weight of the roof and of the impact noise created by rain, sleet, and hail. Spray-on materials can increase damping of metal roof decks, reducing these effects and also increasing the STC.

One of my colleagues assisted a school client whose building was located under the flight path of military training jets that produced uncomfortably loud noise levels. Although the wall construction of the school was sufficient to provide adequate isolation, the many single-pane glass windows and the large metal roof with only thin mineral-fiber ceilings were the weak spots that had to be addressed in order to reduce noise ingress.

Garage bands hoping not to annoy neighbors may have similar issues because of the lack of sealing around the typically low-STC garage door and any windows, and minimal or absent ceiling under a shingled plywood roof. The same goes for the music listener who likes to turn his/her subwoofer up loud, if the sub is located in an outside room with poorly sealed single-pane glass windows, or in an apartment with other apartments below it.

The science of acoustical isolation is well-developed, but has unfortunately not been communicated to the "educated" public. With more widespread understanding, we could avoid many noise problems! ax

Resources

K. Gibbs, "The refraction of sound in hot and cold air," schoolphysics, www.schoolphysics.co.uk/age11-14/Sound/text/Refraction_of_sound/index.html.

R. Honeycutt, *Acoustics in Performance*, Elektor International Media B.V., London, 2015.

Combining Cascaded Stage Harmonic Distortion

While characterizing a new two-stage high-speed amplifier, our columnist took the opportunity to re-confirm a theory that cascaded amplifier + ADC harmonic distortion combine as voltage swings.

By
Michael Steffes
(United States)

The idea that cascaded amplifier + ADC harmonic distortion combine as voltage swings and not powers does strike some engineers as odd at first. In the course of characterizing a new two-stage high-speed amplifier, we had the opportunity to re-test this theory.

The Texas Instruments (TI) THS3217 was developed to ease the conversion from a high-speed complementary current output DACs to a single-ended signal (see Resource 1). This was actually intended as a stand-alone Arbitrary Waveform Generator (AWG) output stage or as a pre-driver to the most recent high voltage, high full power bandwidth, TI THS3491 current feedback op-amp (see Resource 2).

The block diagram for the THS3217 is shown in **Figure 1**. Here the high impedance input buffers isolate the DAC differential signal into a Differential to Single-ended (D2S) amplifier stage. That D2S output is directly available and also connects internally to a very high full-power bandwidth current feedback Output Power Stage (OPS).

Considerable harmonic distortion testing was performed in the course of the product release, using the standard input bandpass filter approaches to purify the source signal harmonics. In the course of testing, a single unit was tested under several conditions including:

- Just the D2S stage driving 2 V_{pp} out into a 200 Ω load
- Just the OPS stage driving a 5 V_{pp} out into a 100 Ω load set for a gain of 2.5 V/V
- The combined stages with the same output conditions in each stage. The D2S produced a 2 V_{pp} into the OPS with an external 200 Ω load where the OPS then gained that up to 5 V_{pp} into a 100 Ω load.

The original two sets of data for the second and third harmonics in a dBc sense were combined using the voltage additive form in Equation 1.

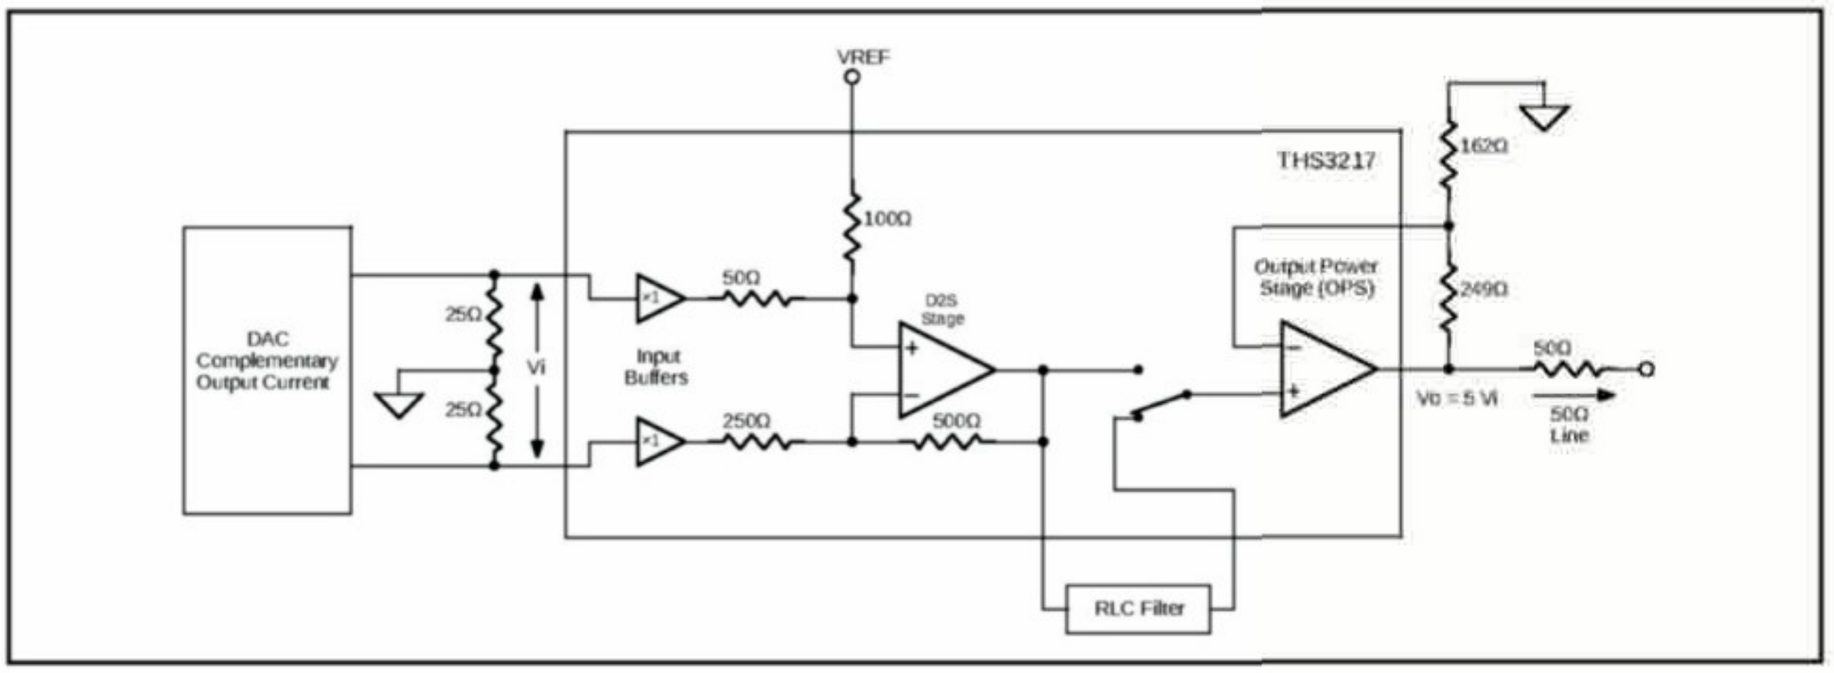

Figure 1: The block diagram for the THS3217 wideband DAC output amplifier

$$HD_{THS3217} = 20 \times \log\left(\left(10^{\frac{HD_{D2S}}{20}}\right) + \left(10^{\frac{HD_{OPS}}{20}}\right)\right) \quad [1]$$

These harmonic distortion numbers are in -dBc and where combined as shown to produce an expected cascaded stage harmonic distortion in -dBc.

Figure 2 shows the calculated and measured result for a single THS3217 device. The estimated harmonic distortion clearly follows the actual measured result for the full two-stage device very closely. Again, these distortion levels are for the D2S producing a 2 V_{pp} output with an external 200 Ω load where that swing is increased to a 5 V_{pp} at the OPS output into a 100 Ω load.

This is pretty strong evidence that the voltage additive approach for the harmonics matches reality. Having that concept clearly in mind then goes on to suggest an innovative approach to extremely low harmonic distortion measurements that became the standard new high speed amplifier measurement approach from about 2002 forward at Texas Instruments Tucson. That will be the topic for the next Audio Scope article.

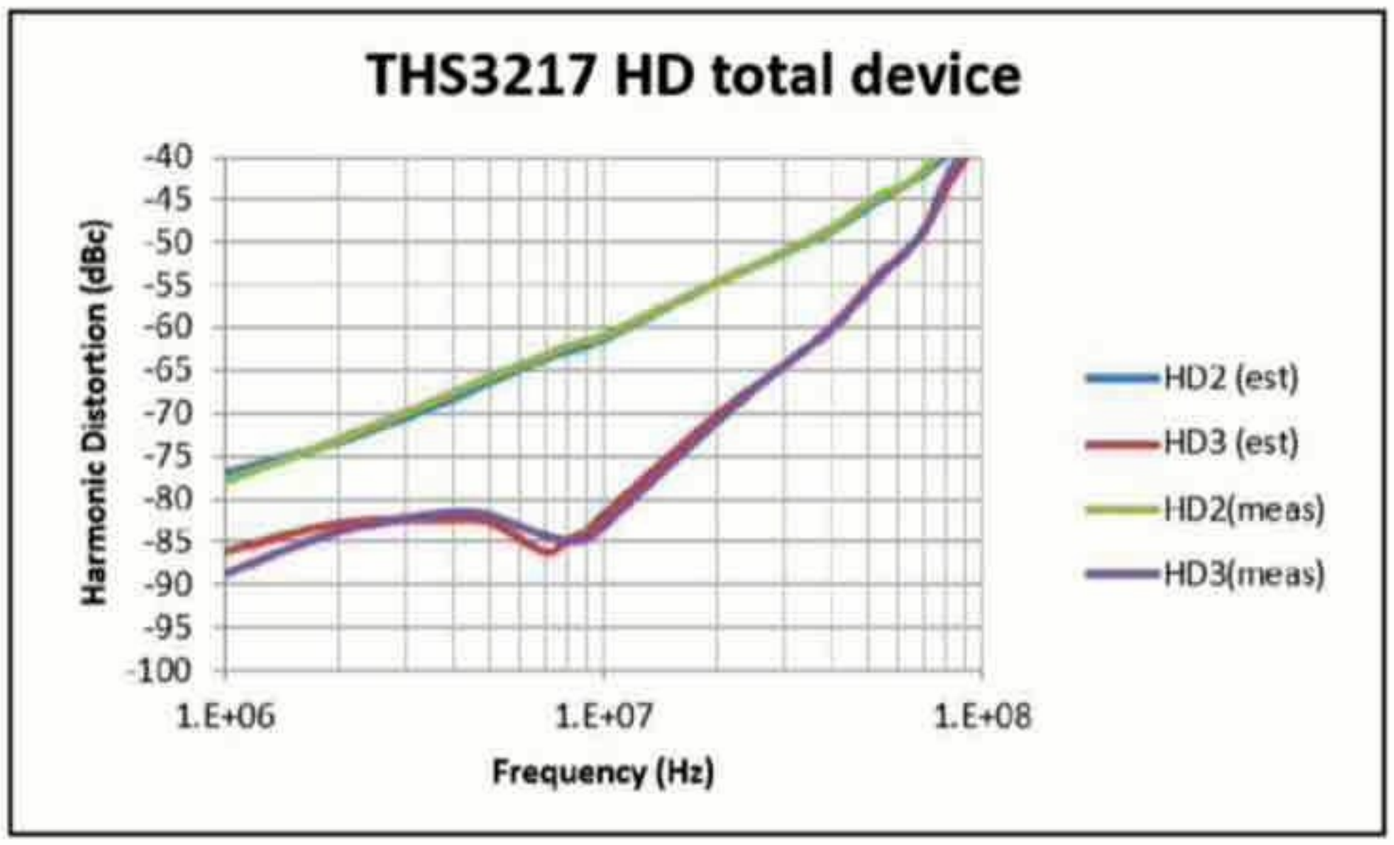

Figure 2: These are the measured and calculated second and third harmonics over frequency for a single THS3217.

Resources

Texas Instruments THS321, "DC to 800 MHz, Differential to Single-Ended, DAC Output Amplifier," http://www.ti.com/lit/ds/symlink/ths3217.pdf

Texas Instruments THS3491, "900 MHz, 500 mA High Power Output Current Feedback Amplifier," http://www.ti.com/lit/ds/symlink/ths3491.pdf

Audio Precision APx555 B Series Audio Analyzer

Expanded Capabilities and Software

Audio Precision recently loaned us one of its newest analyzers, the APx555 B Series. Stuart Yaniger had one month to test it, and in this article he provides readers with a hint of the new expanded capabilities—both of the APx555 B Series hardware and the new APx500 5.0.1 software.

With the announcement of its new APx555 B Series two-channel audio analyzer, Audio Precision also unveiled the updated version 5.X of its APx audio measurement software. Since December 2018, new APx analyzers, whether the new B Series or Legacy, ship with version 5.0 software and one year of software maintenance, effectively licensing that instrument for APx version 6.0 when it is released, as well as any minor releases that occur between versions 5.0 and 6.0.

By
Stuart Yaniger
(United States)

I was recently involved in an online discussion of the term "Gold Standard" in the context of test and measurement. Regardless of eristic quibbles, I think there would be near unanimity that in the realm of audio measurement, Audio Precision (AP) has long been the gold standard. Its earlier analyzers had spectacular performance and versatility, but were notoriously difficult to use. AP addressed this in its APx line of analyzers, which are far more intuitive and simple to use, and the company also expanded its instruments' capabilities from electronics to the realm of acoustic and transducer measurement.

Part of being the gold standard is recognizing that markets and needs for test and measurement expand and shift, and effectively respond to those changes. The growth of smart speakers over the past few years has been remarkable—this product category has gone from nonexistent to nearly 100 million units in less than five years, and the growth and proliferation of brands do not seem to be slowing down.

Smart speakers present a significant measurement challenge compared with more traditional loudspeakers. The measurement paradigm of "feed it a test signal, then measure what comes out" will usually fail, either because of the need to address the speaker via cloud storage with its variable and uncertain timing or the built-in programming that detects when signals are something other than speech or music (e.g., sine waves) and mutes the speaker.

Another change in measurement needs is the new generation of electronics: amplification; analog-to-digital converters (ADCs); digital-to-analog converters (DACs); and signal processors, which can have distortion so low that existing instruments have difficulty resolving it.

Recognizing these and other new needs, AP has introduced an updated version of its flagship APx555 analyzer, the APx555 B series. At the same time, AP has made a major update to the APx 500 software, which incorporates enough new features and operational

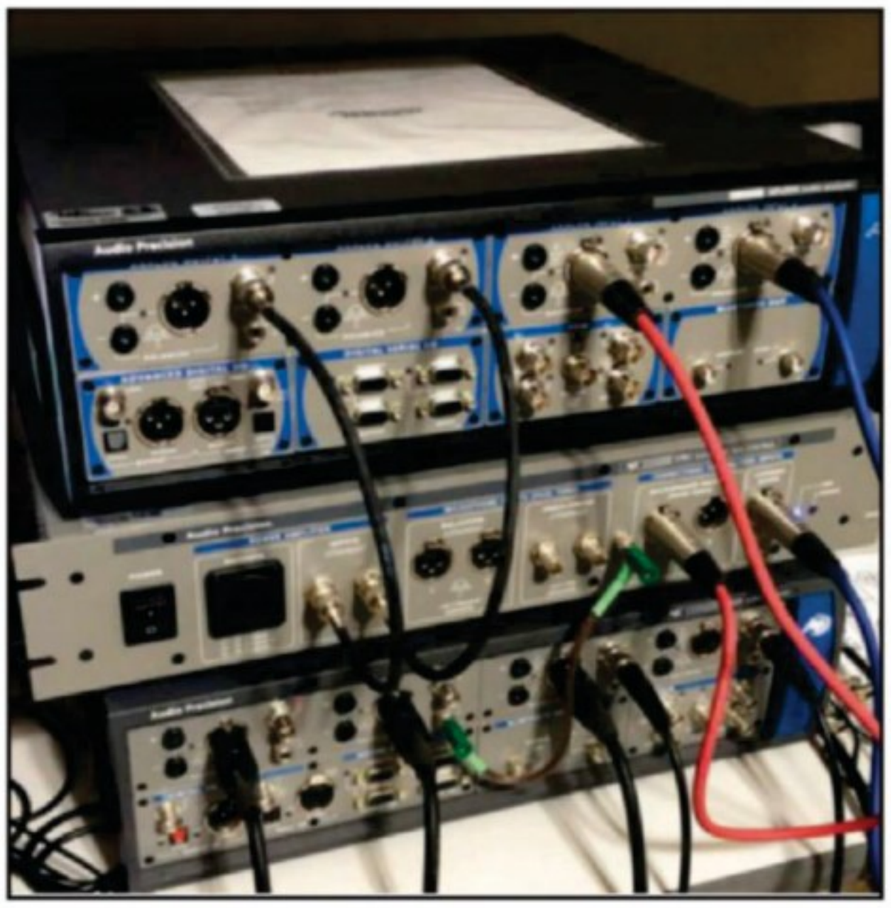

Photo 1: For four weeks, I was the coolest kid on the block. Top to bottom: APx555 B Series audio analyzer, the APx1701 transducer test interface, and the APx525 audio analyzer.

improvements to justify a new version number, 5.0.1.

After working with AP equipment and its applications people for the past several years, I was fortunate enough to be loaned one of the new APx555 B Series analyzers with a full software upgrade (see **Photo 1**). Sadly, I only was able to keep it for a month, but that allowed me time to experiment with some of the new features and see some of what the latest and greatest can do. I have only scratched the surface here, but I hope that this article will at least give a hint of the new expanded capabilities both of the APx555 B series hardware and the new APx500 5.0.1 software.

Distortion Measurements

The B series now features a new High Performance Sine Generator and a High Performance Sine Analyzer, which AP claims reduce THD+N to better than -120 dB, with a 1 MHz bandwidth and Fast Fourier Transforms (FFTs) of up to 1.2 million points. This is a significant improvement over its previous series of instruments, which were certainly not slouches in the distortion and noise department. One might reasonably ask about the practicality of this improvement, but there are certainly reasons why this could be useful. I'll provide a couple examples.

Last year, *audioXpress* Technical Editor Jan Didden was here for his annual visit and brought along a prototype of his sound card measurement interface (the Autoranger). We hooked it up to the APx525 in order to characterize its performance. We seemed to hit a limit at -115 dB third harmonic and couldn't figure out why, since simulations had predicted much lower distortion. After consideration of this observation over a glass of single cask Irish whiskey (Jan is the best house guest ever!), it occurred to us that we hadn't run a loopback measurement to determine the APx distortion floor. And... yes, it was -115 dB. To that point, the APx525 distortion was lower than any device under test I had thrown at it (and to be fair, -115 dB is pretty darn low, about 0.0002%), but we finally had something that the APx525 couldn't measure.

Likewise, I recently reviewed the superb Cambridge Audio Edge W power amplifier. I measured distortion at -115 dB, and remembering what Jan and I had seen, I wrote that the amp's distortion was likely better than that. But was it, and how much better?

Having learned my lesson, I started by performing loopback measurements between the APx's outputs and inputs. There are two ways to accomplish this, either by running a cable directly from out to in, or more conveniently, checking the "Loopback" box in the APx500 software, which automatically connects output to input. I tried it both ways and saw no differences.

I first checked distortion and noise in both balanced and unbalanced modes at 2 V and 1 kHz. **Figure 1** shows the distortion spectra at 2 V for both balanced and unbalanced; I displaced the test frequencies by 100 Hz for clarity. The largest component is the second harmonic at nearly -150 dB from the fundamental. This is astounding performance, only slightly more than 0.000003%—and no, the decimal is not misplaced, that's 3 ppm. The noise over a 24 kHz bandwidth is under 1 µV and dominates the THD+N measurement, which was slightly under 0.0001%. Clearly, AP's new methods of distortion reduction are effective! Of course, the High Performance Generator and Analyzer are slower than the standard performance option, but that's a small trade-off.

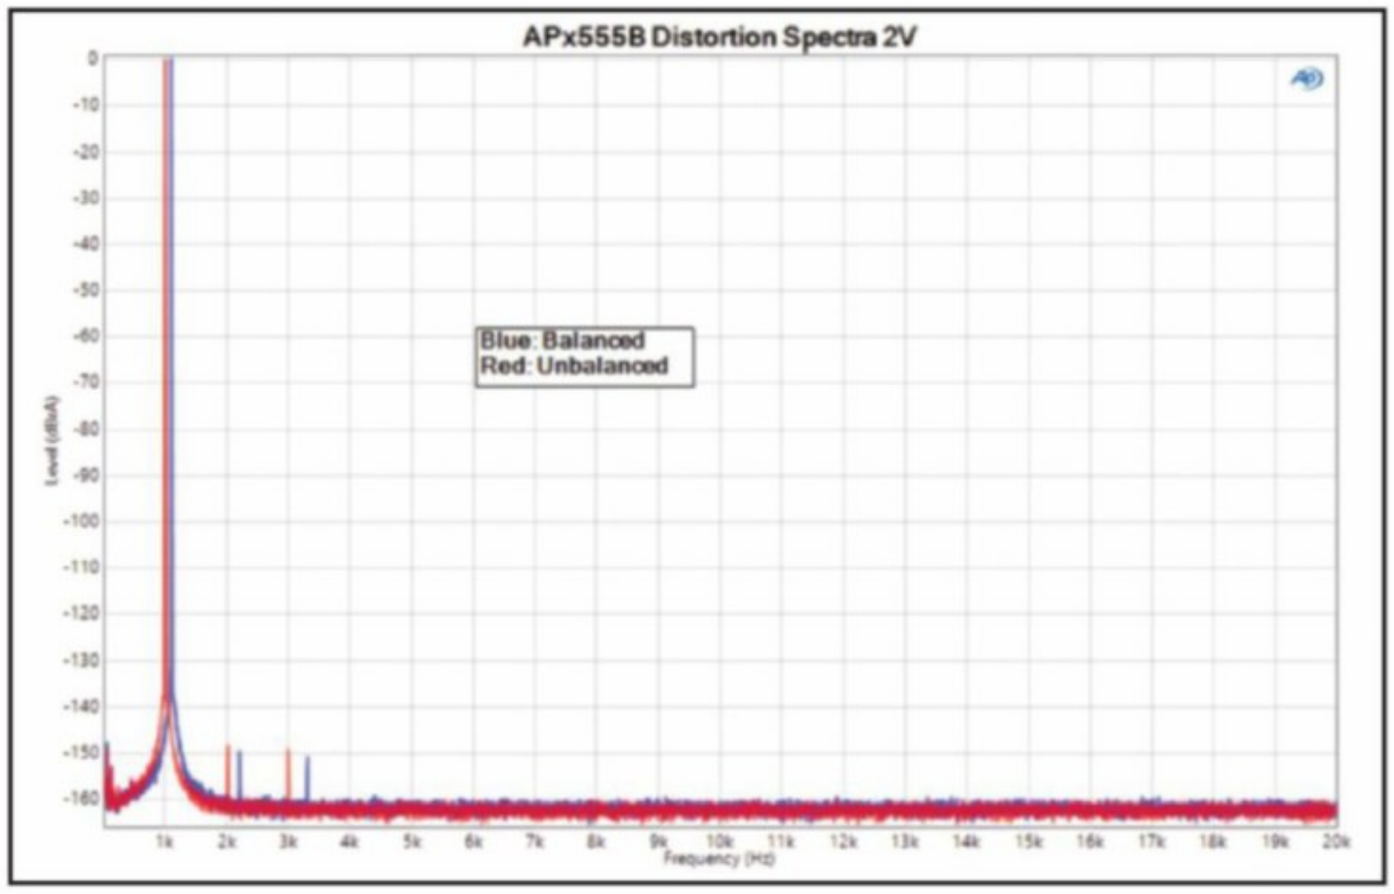

Figure 1: The distortion spectrum at 2 V output shows remarkably low harmonic content and noise floor.

Then, I heaved the Edge W up onto the test bench, connected the balanced output of the APx555B to the Edge W input, attached an 8 Ω dummy load, and measured the distortion at 1 W (2.83 V) output. I chose that low power deliberately since crossover distortion would be greater than at full power. The distortion spectrum is shown in **Figure 2** and proved my suspicions correct. This amplifier has lower distortion than the APx525 I used to measure it! I'm not sure of what impresses me more, the amp's performance or

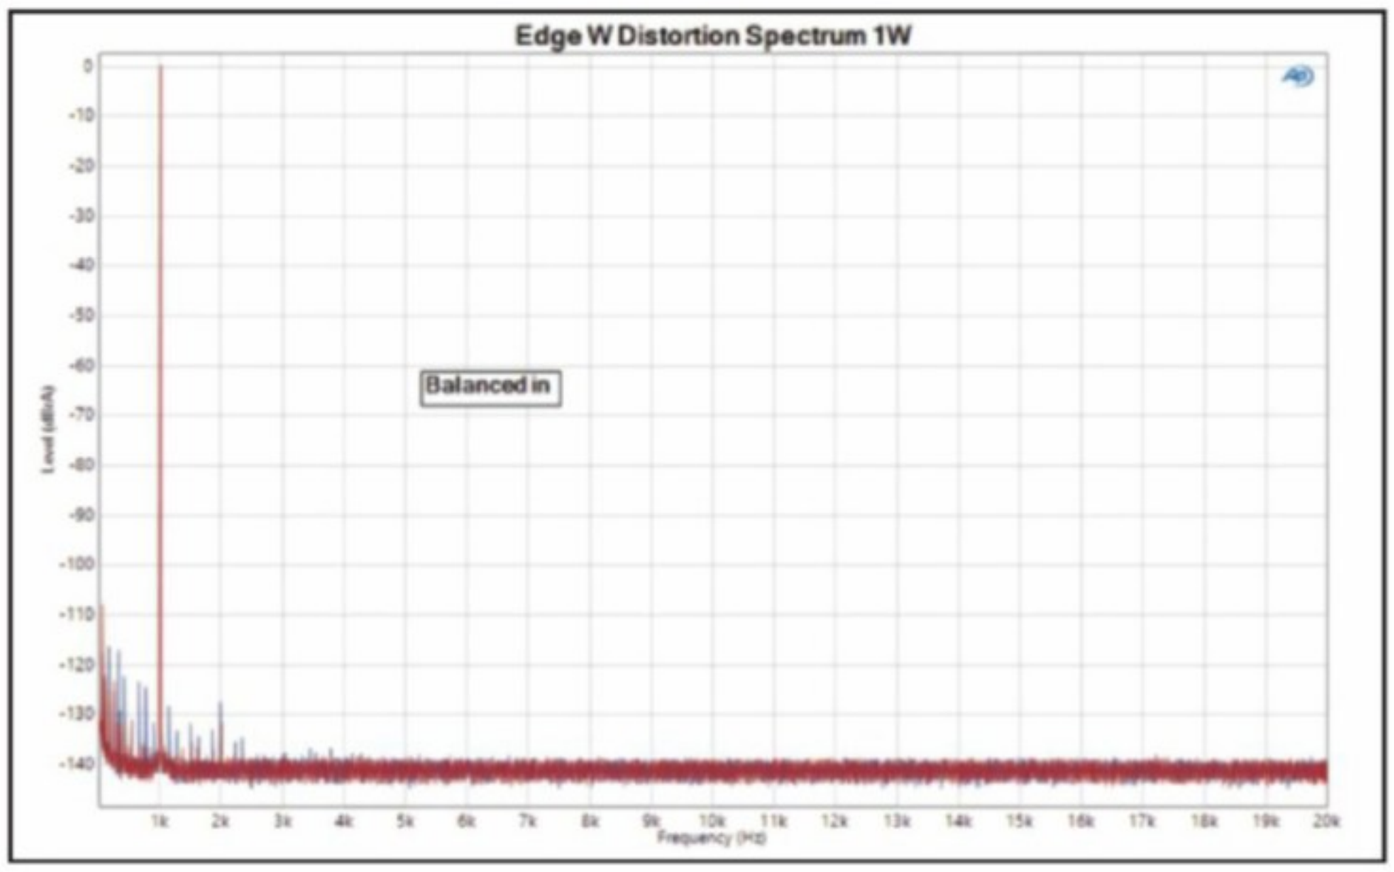

Figure 2: Re-measuring the Cambridge Edge W with the APx555 B Series showed that my previous measurements were, as I suspected, limited by the analyzer, not the amplifier.

About the Author

Stuart Yaniger has been designing and building audio equipment for nearly half a century, and currently works as a technical director for a large industrial company. His professional research interests have spanned theoretical physics, electronics, chemistry, spectroscopy, aerospace, biology, and sensory science. One day, he will figure out what he would like to be when he grows up.

the B Series APx555's ability to measure it, but either way, I'm impressed.

Of course, I had to now retest the Autoranger. And as predicted, the distortion was indeed fabulously low, about -145 dB second and -135 dB third harmonic. No wonder we couldn't reach its limits with the APx525.

Transfer Function

Traditionally, electroacoustic equipment was tested in some sort of synchronous mode: a test signal was applied and the output was measured. In the case of conventional amplification, signal processors, or transducers, there might be a constant delay, but the relation between input and output in terms of timing is pretty constant.

Devices such as Bluetooth headphones and smart speakers present a new challenge. For example, if I want to measure the frequency response of a conventional loudspeaker, I can apply a test signal like a log-frequency sweep (also called a "chirp"), derive the impulse response, and after performing a Fourier transformation, simultaneously acquire frequency response, distortion, and an energy-time curve from that single data set (the Farina method, see Resources).

Now, how do I apply the signal to a smart speaker? The only practical way is to put the desired signal in a file in the cloud, which the smart speaker can retrieve and play back. Unfortunately, this isn't in any way synchronous, and the time delay is quite variable.

Last year in version 4.6 of the APx software, AP introduced the Open Loop function, which allowed a pre-recorded chirp signal to be played back and captured by the APx analyzers for analysis. This took the variable (and long!) time delay involved in recording the chirp, uploading it to the cloud, then having the smart speaker retrieve and replay it out of the equation.

Problem solved? Well... not quite. This approach helps a great deal, but there are limitations on the sorts of test signals that can be used and how they can be accessed. For example, chirp signals can often be sensed as sine waves and suppressed by algorithms in the "smart" part of the signal processing. Likewise, multitone signals can be mistaken for noise and also suppressed. Clearly, this is also true of actual noise signals (e.g., pink or white noise). And since real-world signals such as speech and music can have high crest factors, there's always a question about correlating the measurements with test signals to actual performance in the end use.

The transfer function method allows the use of any arbitrary file (e.g., recorded speech or music) as the test signal (termed the "reference file"), with the proviso that the test signal be broadband to cover the frequency range of interest. The acquired asynchronous device under test (DUT) output, typically

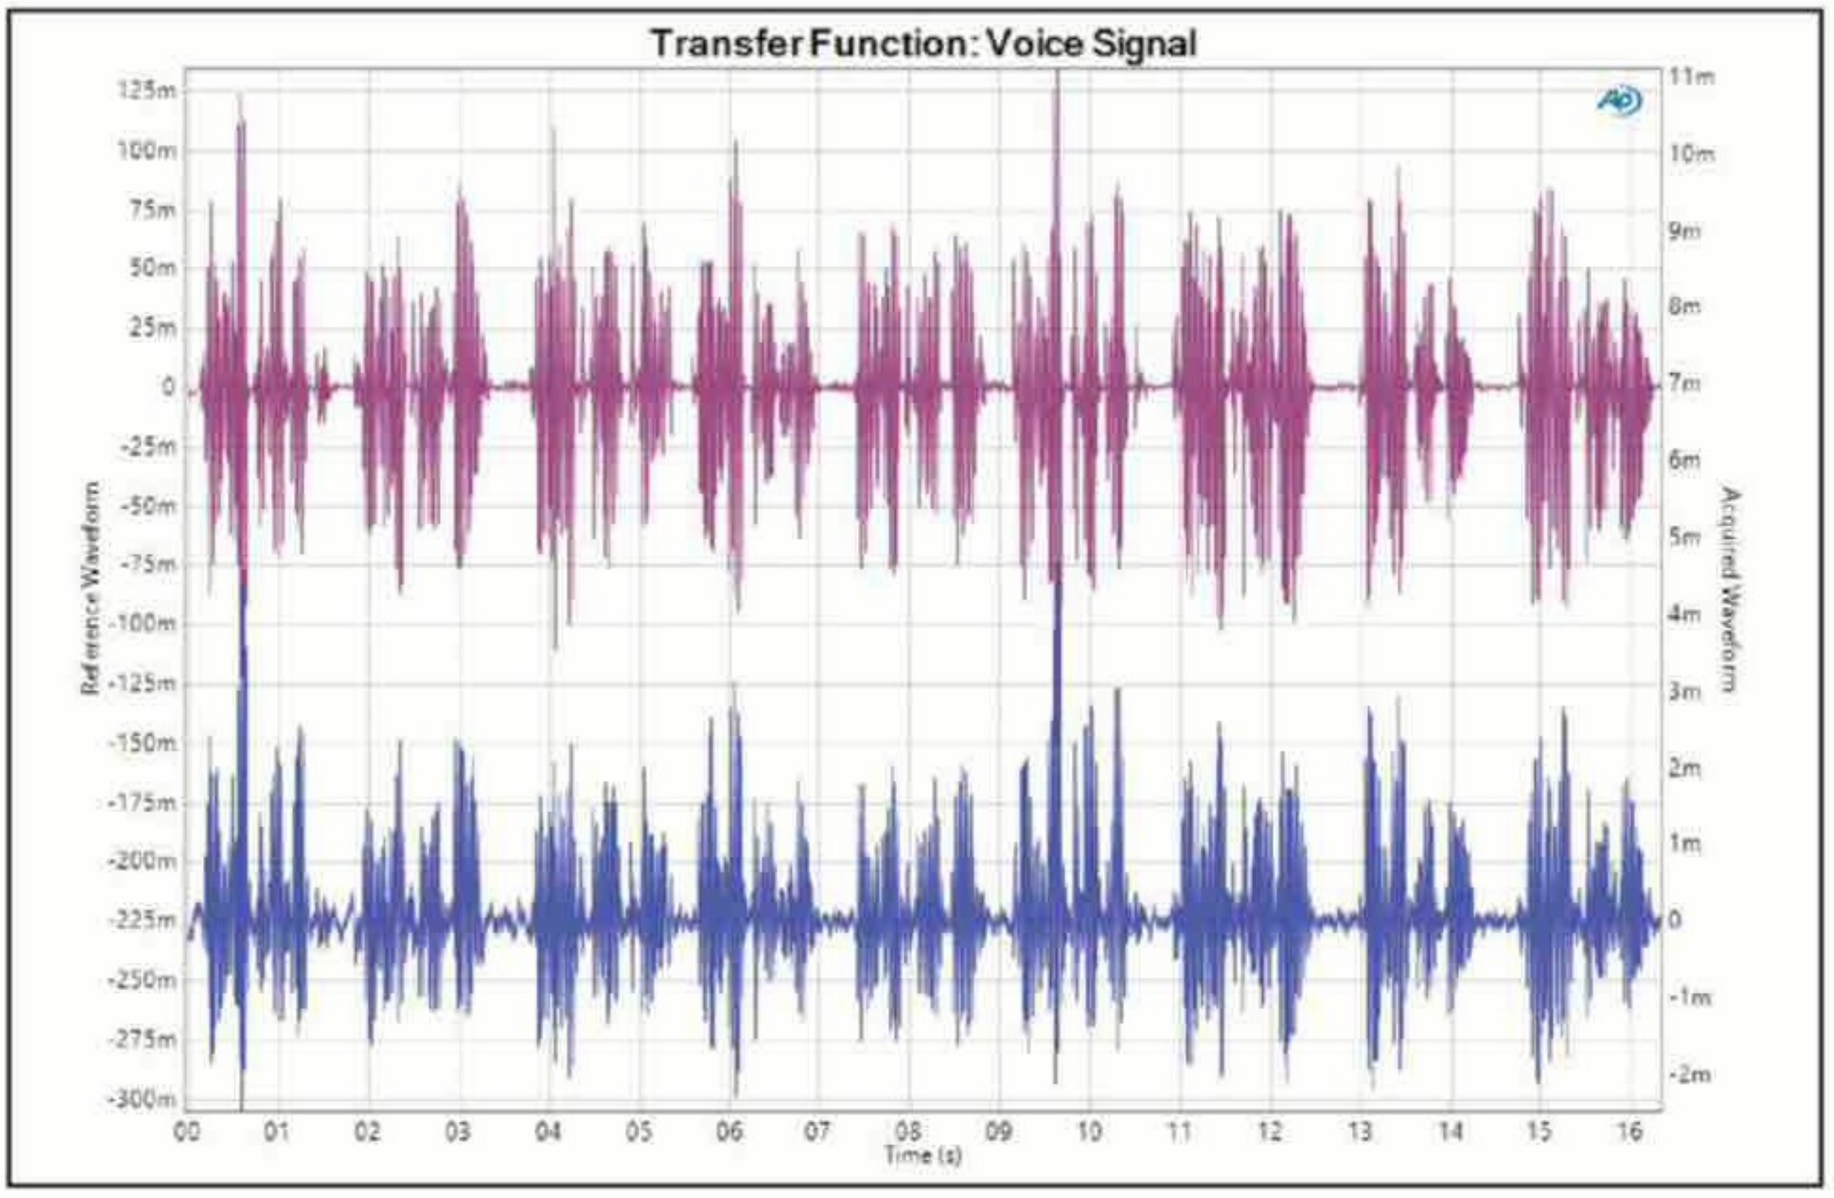

Figure 3: The transfer function measurement module automatically lines up the acquired signal (lower trace, blue) and the reference file (upper trace, red) in time.

an acoustic output measured with a microphone, is compared with the reference file. That seems simple, but the asynchronous nature requires a trick in order to match up the reference and acquired signal, namely the use of a cross-correlation function (see the sidebar article). The APx software allows the matching criteria to be set at High, Medium, Low, or None, depending on the desired degree of precision balanced against criticality of measurement conditions. Likewise, you can set the time allowed for the matching before the software "gives up" and returns an error. This can happen if there's too much noise or the reference (and/or acquired signals) have insufficient bandwidth.

Once the matching is complete and the reference and acquired waveforms are aligned in time, the acquired signal is automatically truncated to remove the portions outside of the desired output (e.g., the delay between starting the acquisition and the beginning of the actual signal output), and the two files can then be compared using a complex Fourier transform to extract DUT frequency response amplitude and phase information, as well as coherence and both amplitude and power spectral densities of the reference and acquired data.

To demonstrate the power of this method, I created a speech signal consisting of alternating male and female voices reading standard phrases (specifically, stitched together phrases from the ABC-MRT16 speech intelligibility test- see Resources). This was saved as a .wav file and used as a reference signal. The reference signal was loaded into the B Series APx555's generator module, and the balanced output was connected to a Kali LP-6 loudspeaker that I have been using in my lab for mastering. I then connected a calibrated PCB Piezotronics 376A31 microphone to the APx1701 transducer interface input and the microphone was positioned about 1 m away from the loudspeaker.

I started the acquisition, then played the reference file at about 70 dBSPL at the microphone position. The transfer function measurement module automatically lined up the acquisition and the reference signals in time as shown in **Figure 3**. Now, here's the magic of the transfer function measurement: this is a voice signal, yet the measurement yields the driver frequency response, which matches very closely with the frequency response measured by more conventional synchronous methods, in this case, a log-sweep (see **Figure 4**).

The transfer function module also allows end-to-end testing of smart speakers. This requires a calibrated acoustic source to output the "wake word" (e.g., "Hello, Siri!") and command to stimulate the DUT to play the desired file, generally a recorded speech, music, or test signal file. If you need to capture the

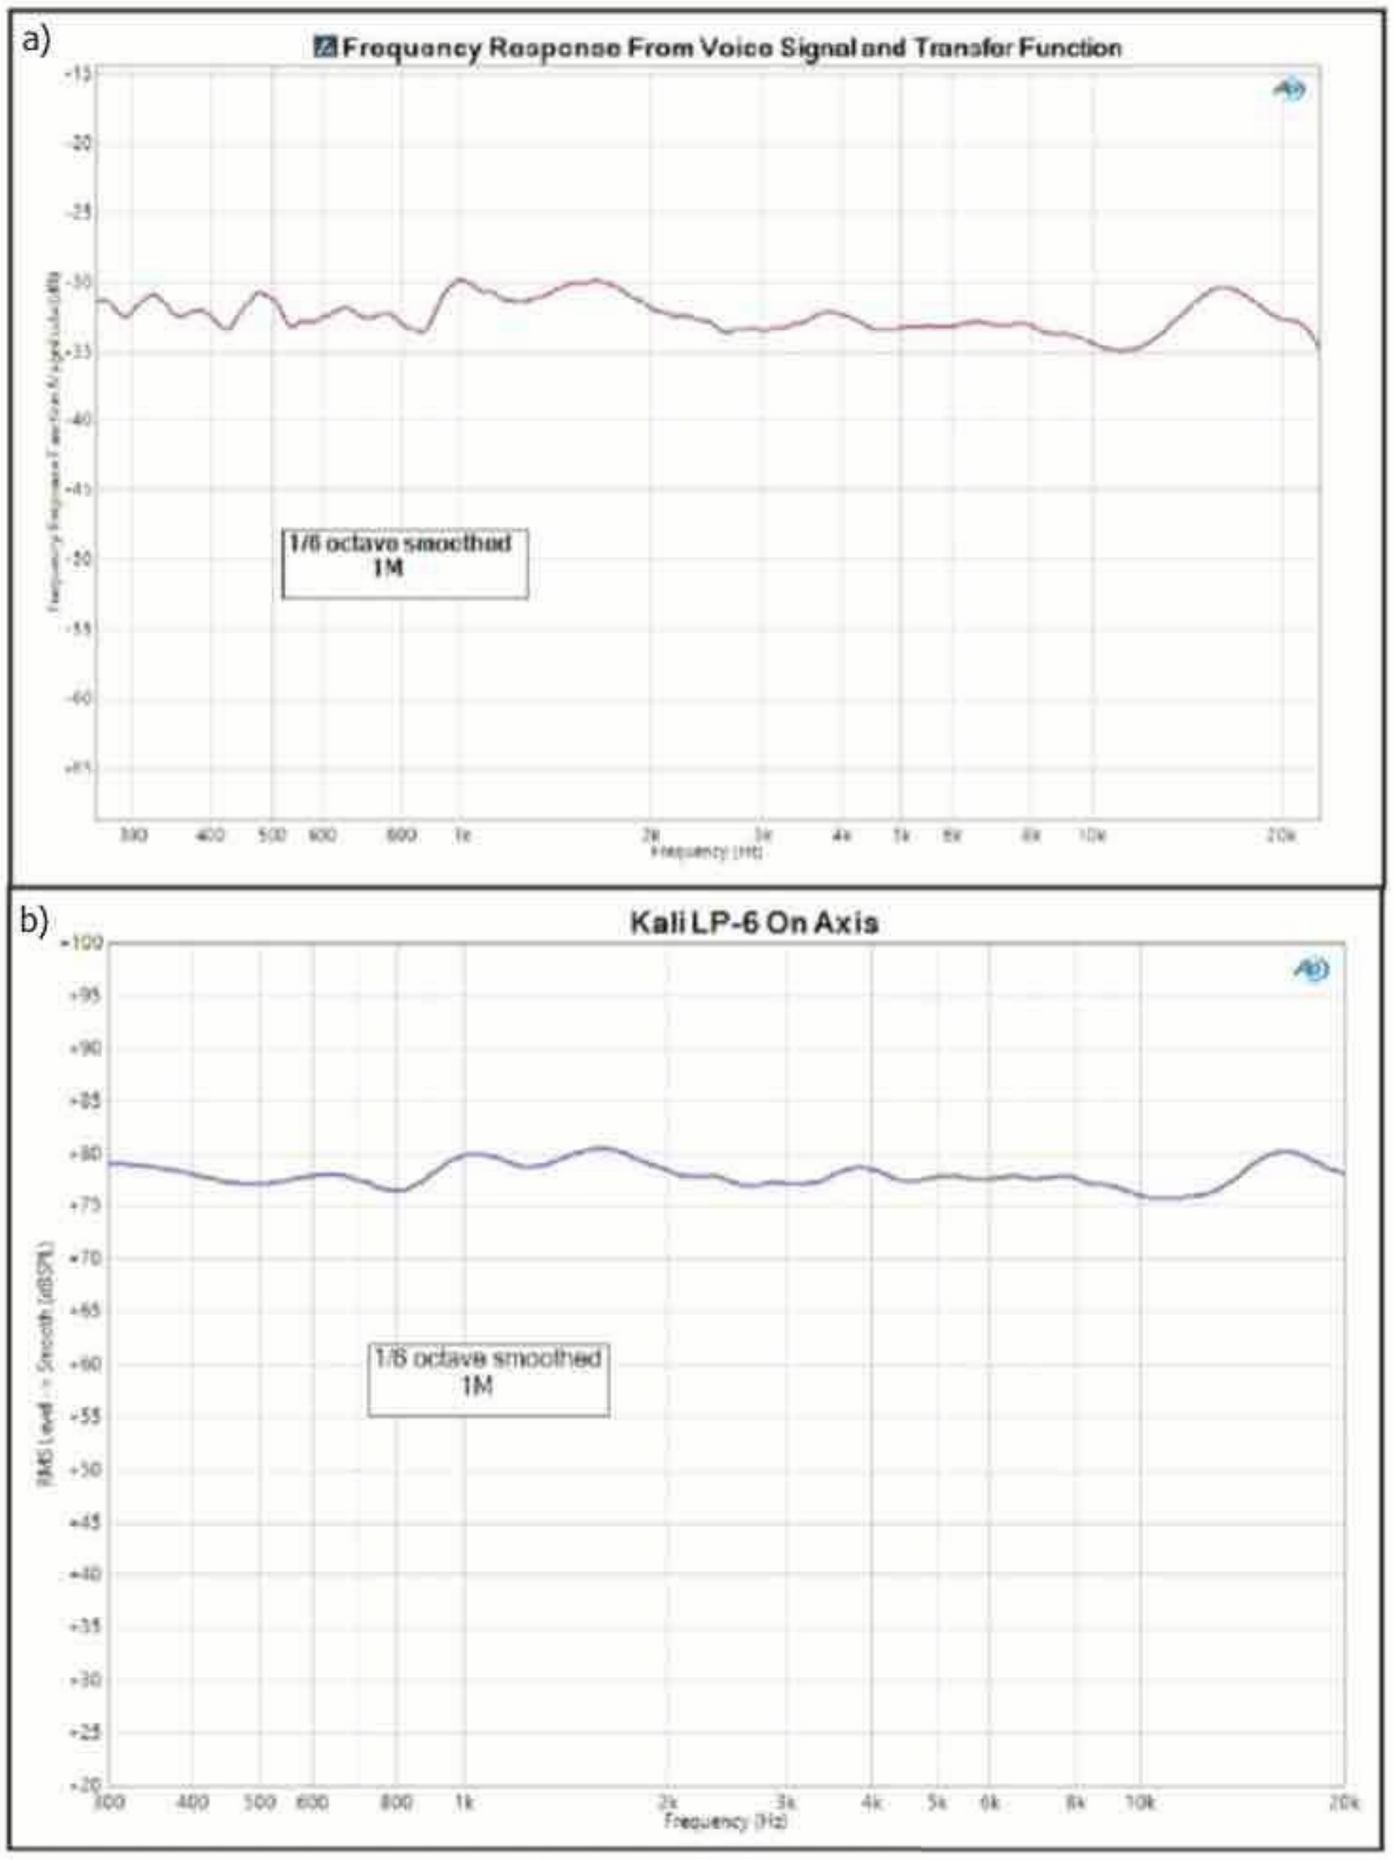

Figure 4: The loudspeaker frequency response measured using a voice signal and the transfer function module (a) compares closely with a measurement taken using a conventional log-frequency sweep (b).

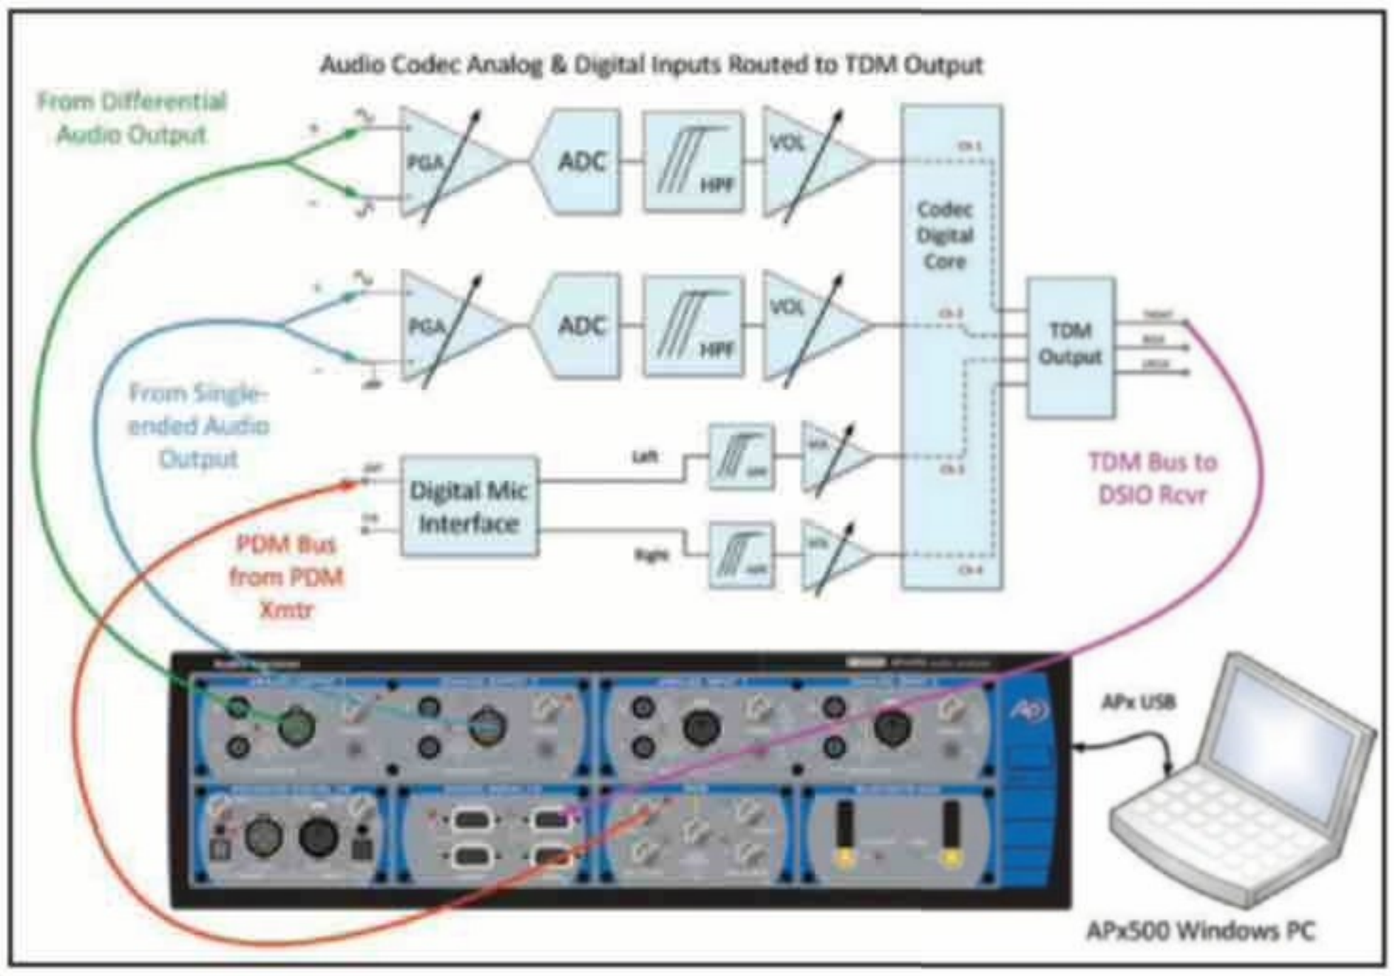

Figure 5: The APx555 B Series can be configured to do simultaneous measurements of codecs with single-ended, differential, and digital inputs. (Image courtesy of Audio Precision)

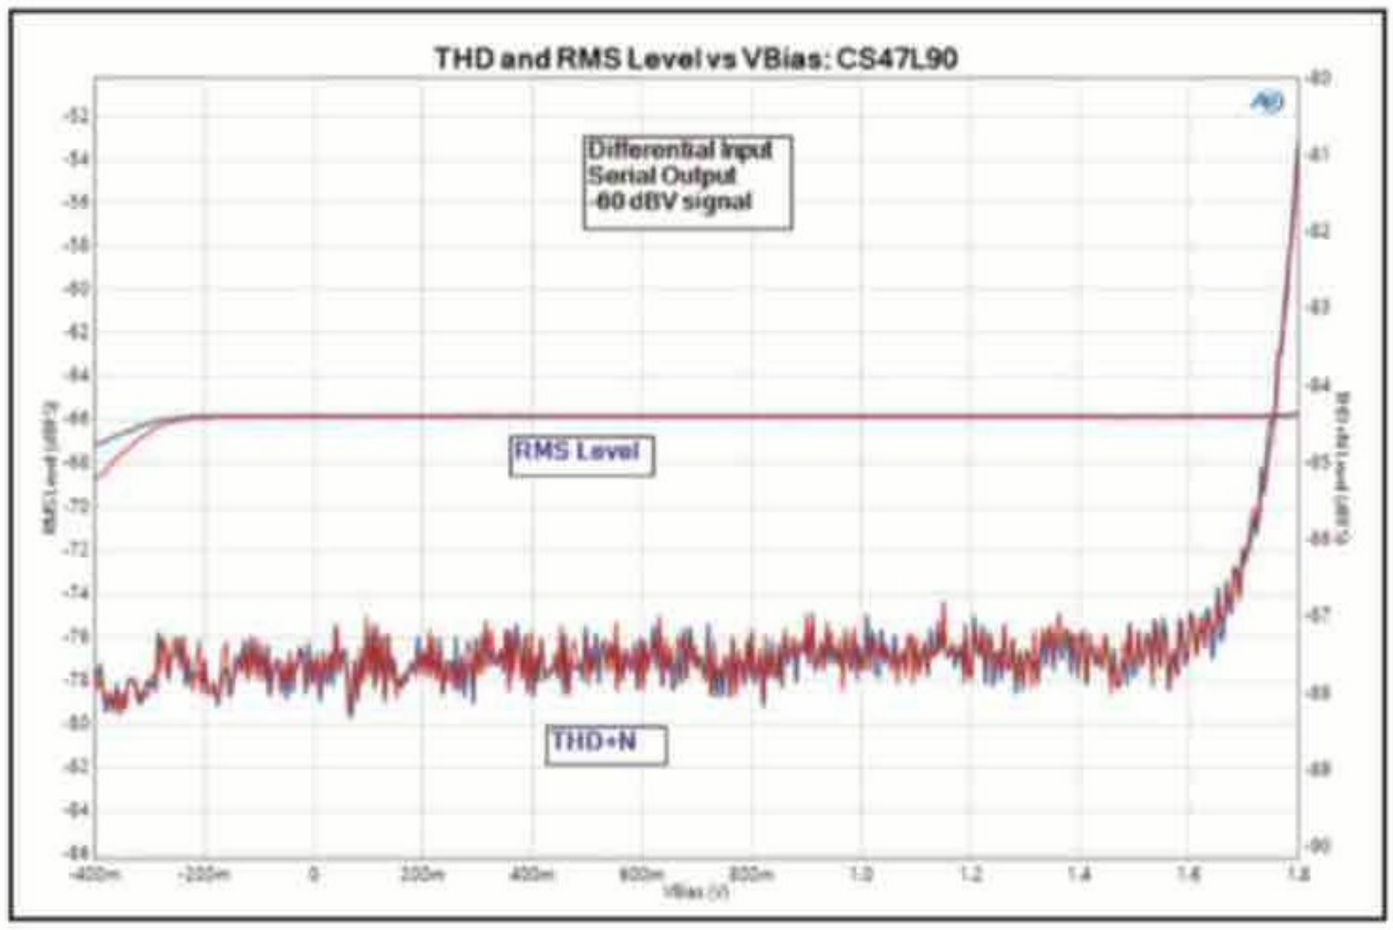

Figure 6: Using the swept DC offset function of the APx555 B Series in differential mode, the sensitivity of a codec to variations in DC offset voltage can be quantified.

speech and replay it (e.g., to characterize the beam-forming function of a microphone array), this sort of test will generally require an anechoic chamber to perform. But the issues of asynchronous capture and analysis can proceed exactly analogously to my rather simple demonstration.

ADC Test

Conventional ADC testing involves injecting signals, then examining the digital files following conversion. A good review of the traditional measurement methods appeared in *audioXpress* a few years ago (see Resources). The continuing proliferation of low power ADCs for applications such as smart phones has created a need for tests that go beyond the previous methods; for example, most ADCs in portable devices work off a single supply and need to be biased at about one-half the supply voltage (Vdd) to operate optimally.

Another complication is that, given the variety of smart devices and IoT appliances that surround us,

Some Basic Signal Processing Theory Behind the Transfer Function

Although an understanding of the basic math behind the transfer function module in the APx software is not necessary for making good use of it, it's probably worthwhile to at least peek behind the curtain, if you're the curious sort (as I am). This is all elementary to engineers, and I'm sure I'll get some significant wincing at my explanations, but for non-specialists like me, a review of the meaning of the terms used and the graphical data that the transfer function provides may be worthwhile.

The cross-correlation between two time dependent functions A and B is defined as:

$$R_{AB}(\tau) = \int_{-\infty}^{+\infty} A(t)B(t+\tau)\,dt$$

where A(t) is the reference signal and B(t) is the acquired signal. It is evident that as the time delay between the signals A and B approaches zero, the cross-correlation will reach a maximum. And this is basically how the transfer measurement aligns the reference and acquired signals; the maximum of the cross correlation function determines the time offset needed to align the reference and acquired waveforms. In the special case where A(t) and B(t) are the same signal, the function is referred to as the auto-correlation, and its maximum will of course occur at $\tau = 0$.

For the loudspeaker measurement discussed in the article, the cross-correlation function is shown in **Figure 1**, where the APx transfer function module has lined up the signals.

Coherence is a measure of causality between reference and acquired waveforms, that is, what proportion of the acquired waveform is due to the stimulus signal and what part is due to noise or other non-causal factors. It will take a couple steps to get to the actual definition, so please bear with me.

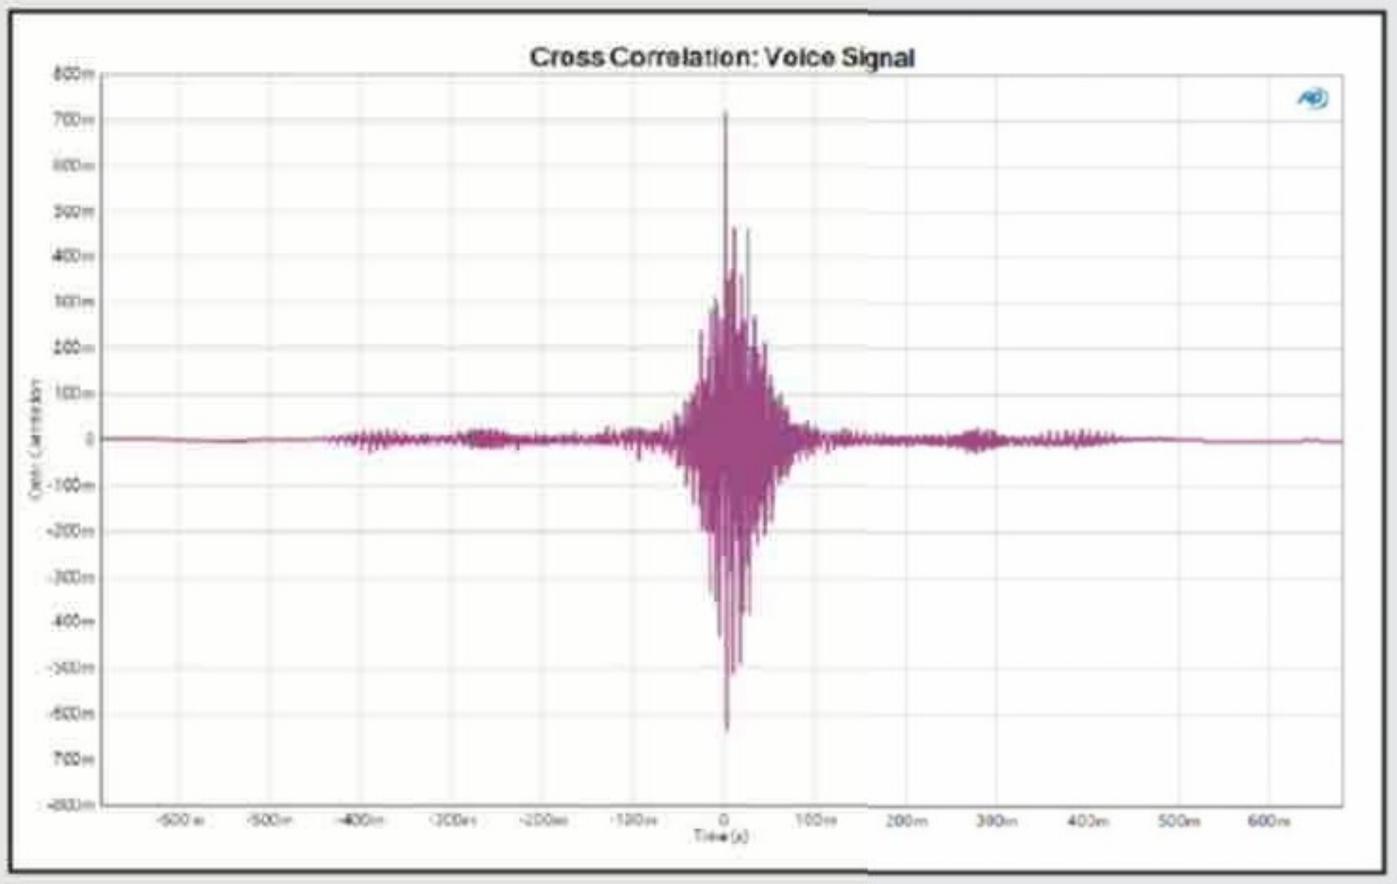

Figure 1: The cross-correlation function calculated by the APx500 software between the reference file and the acquired acoustic signal allows alignment of the two signals in time so that frequency and phase response can be determined.

ADCs are often used in either differential or single-ended mode, and often both at once in the case of multichannel chips. The testing in all of these modes can be rather lengthy and tedious, especially when throwing in the DC offset variable. Here's where the multichannel modes and flexible I/O capabilities of the APx555B come into play. **Figure 5** shows a typical (though complicated!) test setup using the APx555B to characterize a Crystal Semiconductor CS47L90 codec using the Crystal demo board. In this example, all three types of inputs are exercised, differential, single-ended, and digital microphone inputs, with the serial output of the demo board connected to the AP analyzer's digital serial input port.

AP provides several project files as part of its TN136 "Measuring Low Power ADCs" Applications Note (see Resources), so making the connections and the measurements is nearly foolproof. The particular part that is new is the ability to have the analog outputs of the APx555B generate a DC offset voltage riding on the

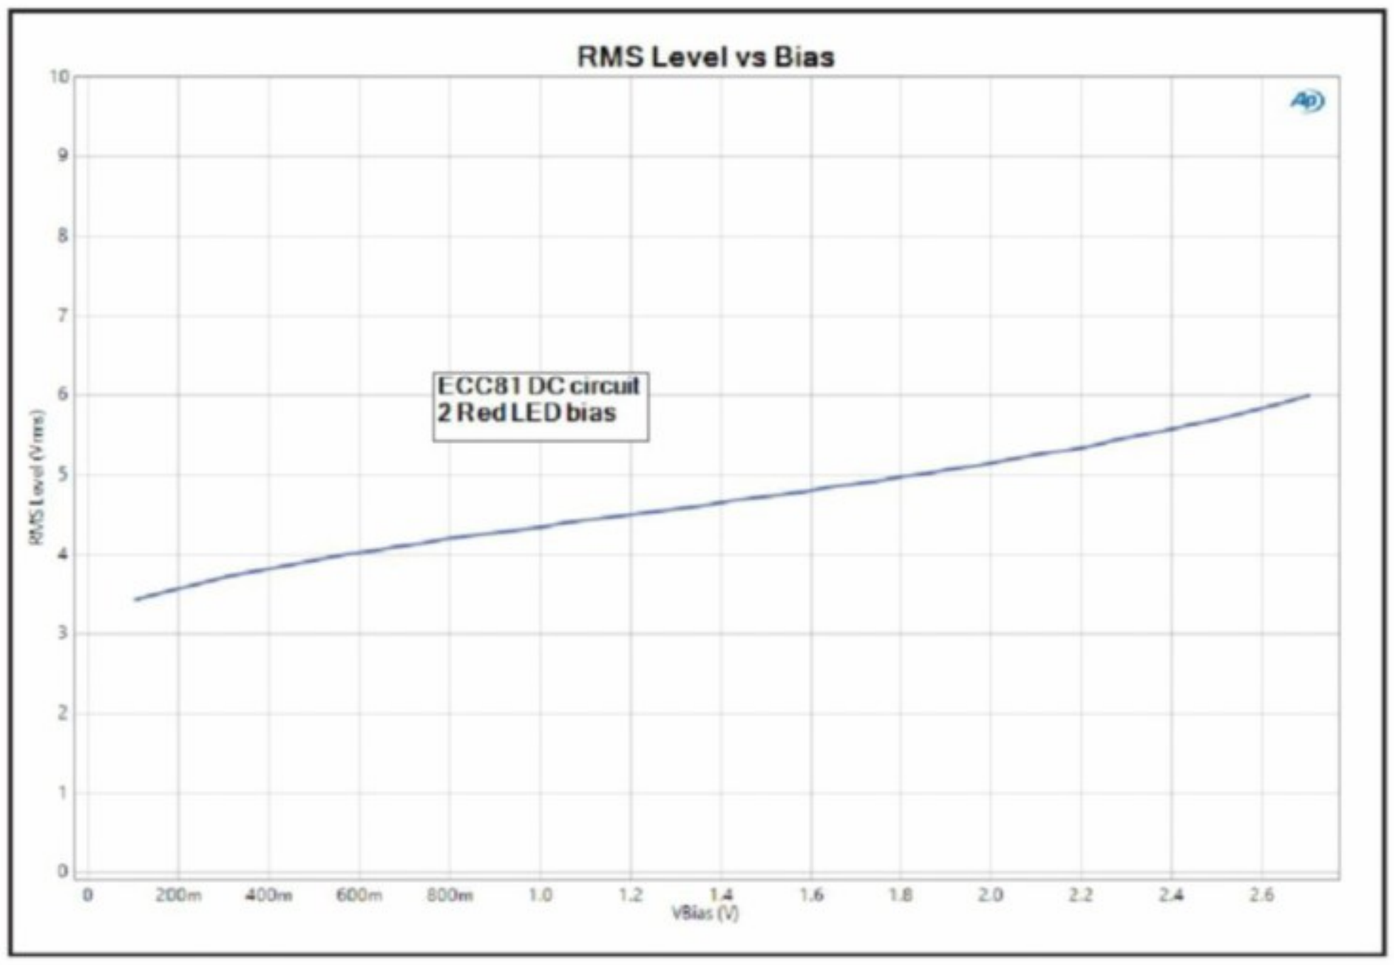

Figure 7: The DC offset sweep function also works to determine the effect of biasing on the gain of an analog circuit block (in this case, a vacuum tube circuit).

First, we need to define the cross power spectral density, S_{AB}. It is basically the Fourier transform of the cross-correlation function:

$$S_{AB}(\omega) = \int_{-\infty}^{+\infty} R_{AB}(\tau) e^{-j\omega t} dt$$

where, as usual, $\omega = 2\pi f$. The coherence is then defined as:

$$C_{AB}(f) = \frac{|S_{AB}|^2}{S_{AA} S_{BB}}$$

Clearly, if A and B are identical (perfect causality), the coherence function is 1 at all frequencies. Where A and B are completely unrelated, the numerator vanishes and the coherence is zero at all frequencies. In places where some of the acquired signal, B, differs from A because of noise or droops in frequency response, the coherence will be less than 1 but greater than or equal to zero.

For the loudspeaker measurement discussed in the article, the coherence function is shown in **Figure 2**; you can see that it's close to unity through much of the audible frequency range, but drops at low frequencies where the driver rolls off and noise dominates the measurements. It likewise has dips at the 60 Hz line frequency and its harmonics. The coherence function is very useful as a basic measurement quality check—if it's significantly less than 1 over most of the range of interest, the acquired data is suspect because of lack of causality between applied and acquired signals.

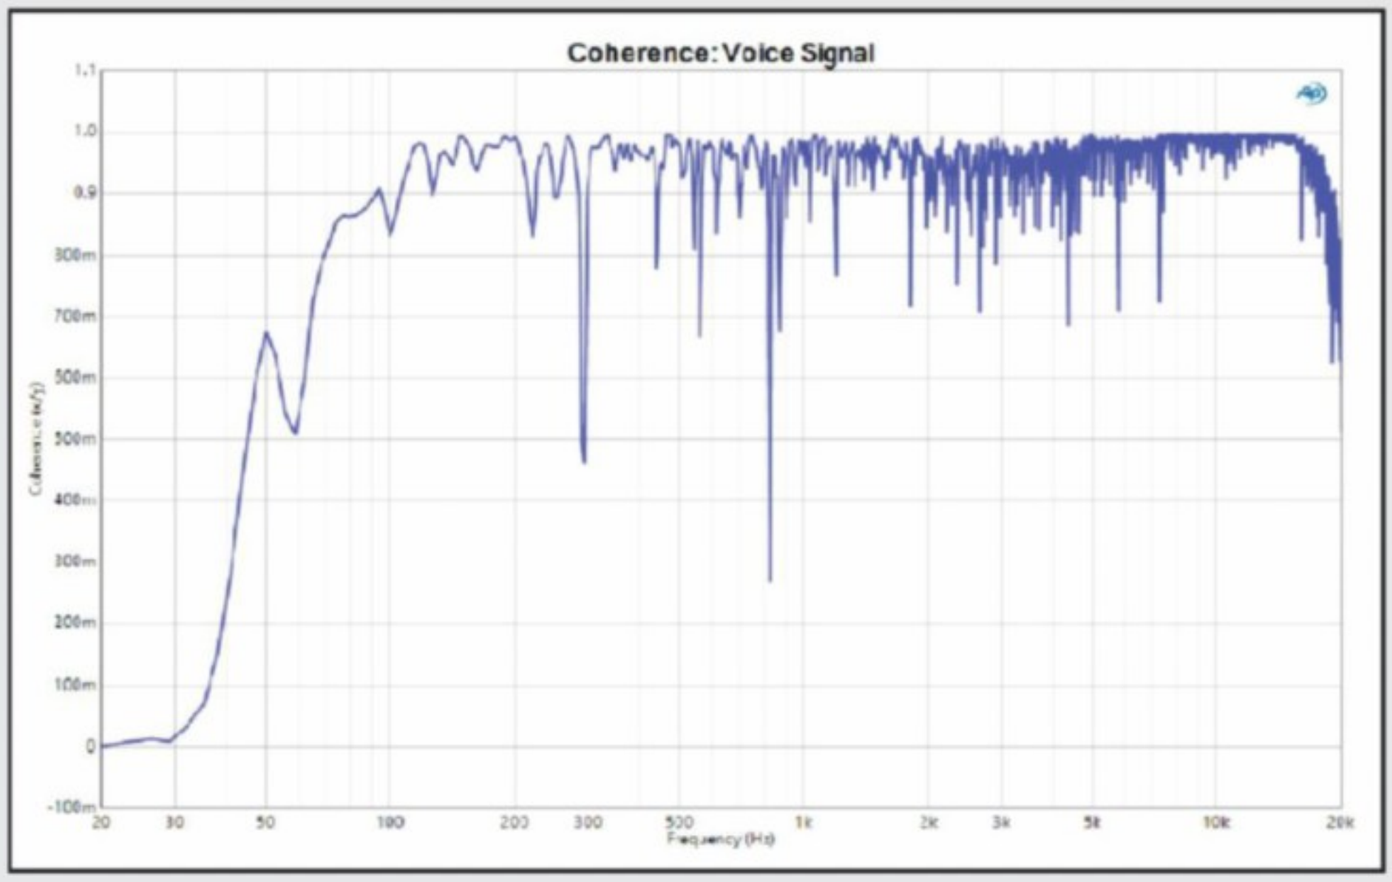

Figure 2: The coherence function calculated by the APx500 software between the reference file and the acquired acoustic signal is a measure of causality of the applied signal to the measured acoustic output.

An additional note: the APx software offers three different calculation mode options, which it terms H1, H2, and Magnitude Only. H1 is the most common choice, chosen when the reference signal has a higher signal-to-noise ratio (SNR) than the acquired signal, which is usually the case, and is S_{AA}/S_{AB}. H2 is more appropriate if the reference signal is noisy (e.g., a sound capture from the microphone of a smart speaker in a room) and is S_{BB}/S_{AB}. If the reference and acquired signals have no common clock reference, you can select the Magnitude Only option, but of course, all phase information is lost.

test signal that can be swept during a measurement. This can be particularly useful for systems using an electret or MEMS microphone that will drop a small voltage. Although 0.5 Vdd is generally the optimum offset, the transducer voltage drop will affect this, and it's important to know how much offset will cause an increase in distortion. The APx software also allows safe limits ("pin protection") for the total offset plus peak signal voltage to be set in order to avoid damaging the DUT.

Figure 6 shows the results of sweeping the offset from -0.400 V to +1.8 V with a 1 kHz sine wave test signal at 1 mV (-60 dBV) applied to the Crystal codec. As is clear, over most of the DC offset range, the total harmonic distortion (THD) and gain remain constant and low. Above about +1.6 VDC, the THD starts to rise, and below -0.250 VDC, the gain starts to drop. This sets the small signal limits for this particular ADC, although if the signal voltage were to increase, the offset limits would likewise decrease because of the sum of the peak signal voltage and offset impinging on the ADC limits.

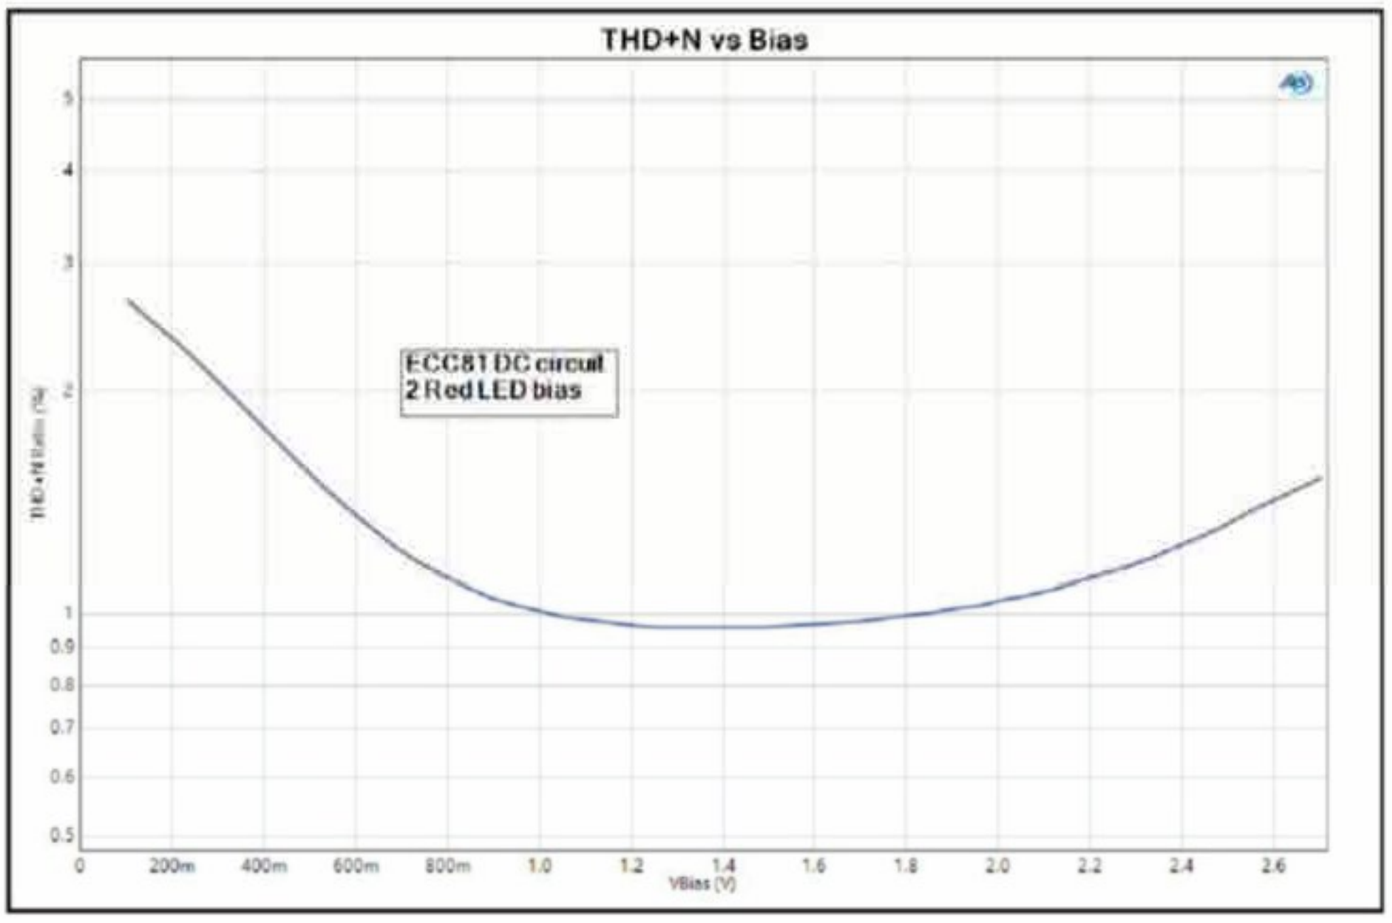

Figure 8: The DC offset sweep function can also be used to determine optimum bias point for an analog gain block to achieve minimum distortion.

Although this measurement is differential, the offset sweep can also be applied to single-ended signals. The Crystal codec in single-ended mode shows similar results compared to differential. But that got me thinking...

For those of us dinosaurs who still do discrete circuit design, one of the interesting design and verification challenges is optimizing bias. It occurred to me that this feature of the APx555 B Series could be adapted to determine the bias level for minimum distortion in an analog amplifier. To try this out, I slapped together a simple, non-optimized gain stage consisting of a 12AT7/ECC81 tube as the gain element (I use whatever is close to hand—this could equally well have been done with FETs or degenerated bipolar transistors) and cathode biased it at +3.4 V with respect to the grid using two red LEDs in series. I then applied a single-ended signal (0.141 VRMS at 1 kHz) to the grid and swept the DC offset from +0.1 V to +2.7 V, which combined with the fixed bias voltage at the cathode effectively swept the tube's overall grid-to-cathode bias from -3.3 V to -0.7 V, respectively. A graph of the output voltage versus applied bias with the drive voltage fixed is shown in **Figure 7** and one can see nearly a 5 dB gain shift over this bias range. Likewise, the THD versus bias is shown in **Figure 8** and has a broad minimum at +1.4 V, corresponding to a grid to cathode voltage of -2.0 V, which I've previously observed is a "sweet spot" for this tube type. If you're designing and optimizing any sort of analog gain stage, this is a very handy way to optimize distortion and gain.

It did not escape my attention that this feature could also be used to see the effect of offset voltage on the common mode rejection of a differential input, but unfortunately, I ran out of time before gathering this data. When I win the lottery and buy one of these for myself, I promise to show some examples!

Final Thoughts

The various new features of the APx555 B Series and the latest version of software were an absolute delight to use. Not only do they do exactly what they are intended for, their capabilities can also be extended into other uses. The upgrade from older APx555 versions is not cheap at about $6,000, but in my view, it is very much worth it for companies pushing the state-of-the-art in distortion or developing smart devices. It's unfortunate that I only had a short time to explore the new capabilities, but they have been sorely missed since returning the unit to AP. Maybe I'll get a second lottery ticket tonight...

Resources

"Application Note: Smart Speaker Acoustic Measurements," Audio Precision, 2019, www.ap.com

"CS47L90 Low Power, Highly Integrated Smart Audio Codec," Cirrus Logic, https://statics.cirrus.com/pubs/proBulletin/CS47L90_PB_20160919.pdf

D. Mathew, "Testing Audio ADCs and DACs," *audioXpress*, December 2015.

G. Munster and W. Thompson, "Smart Speaker Macro Model Update," Loup Ventures, LLC, June 2013, https://loupventures.com/smart-speaker-market-share-update

"Transfer Function Measurements with APx500 Audio Analyzers," Audio Precision Tech Note TN138, www.ap.com

S. Yaniger, Speech test file, SYclotron Audio, https://syclotron.com/sound-clip-supplement-to-audioxpress-review-of-APx555B

Simulation + testing = optimized loudspeaker designs

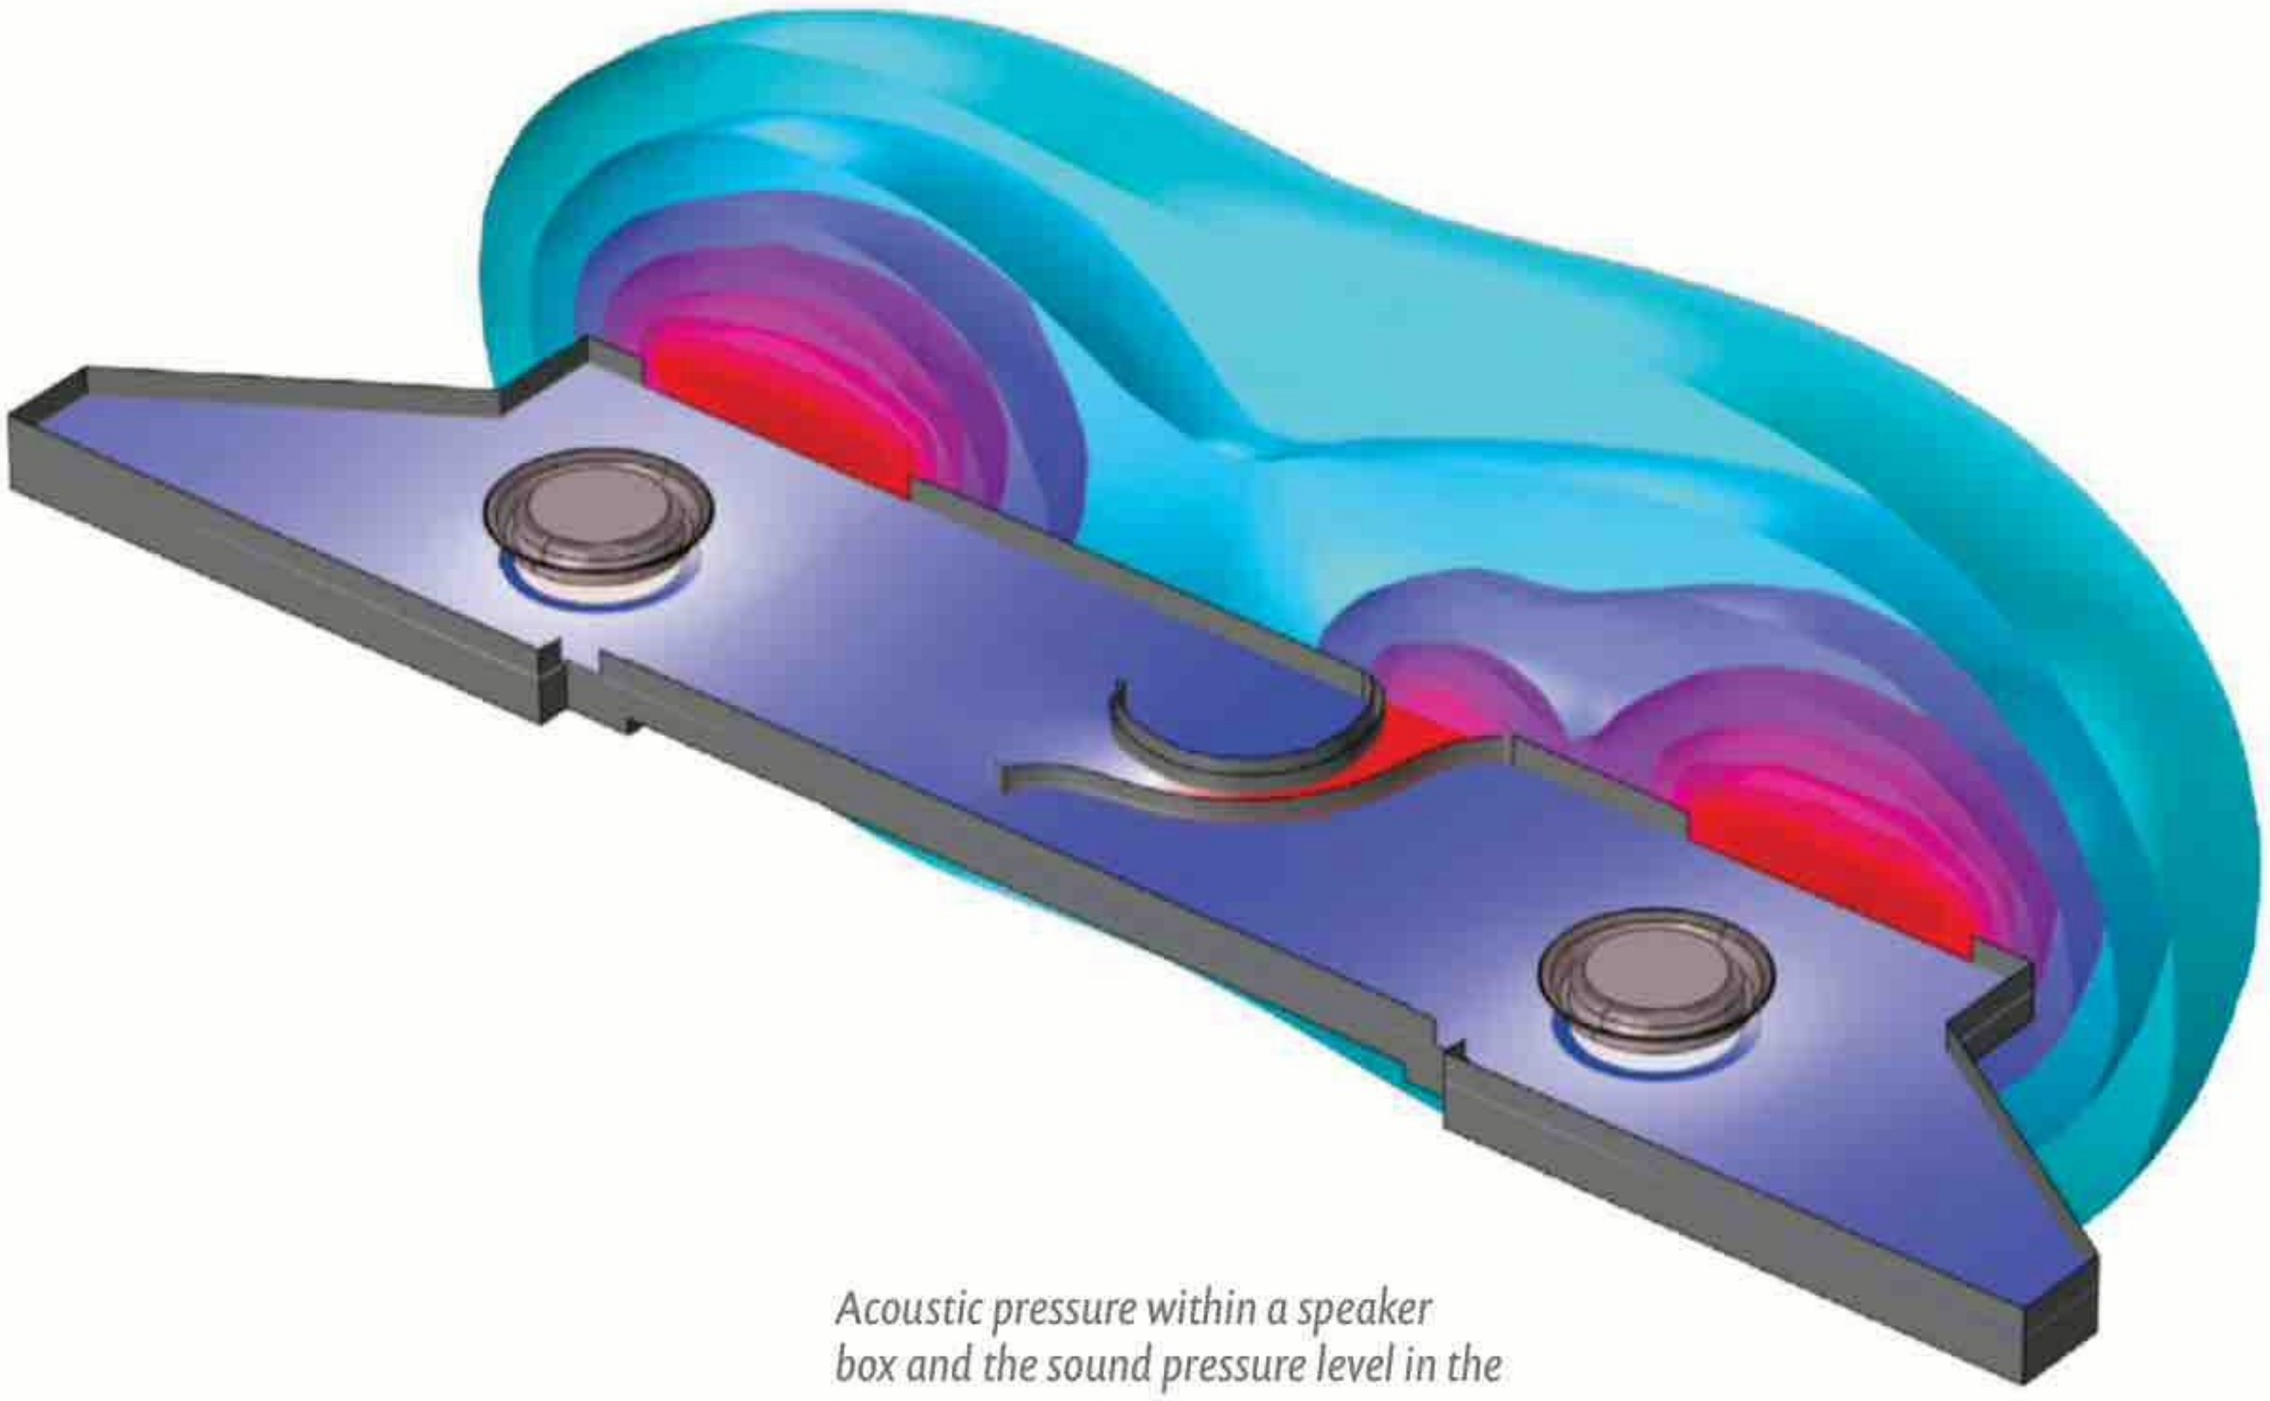

Acoustic pressure within a speaker box and the sound pressure level in the surrounding domain.

A global leader in electronics rose to the top of the audio industry by adding multiphysics simulation to their design workflow. COMSOL Multiphysics® enables audio engineers to couple acoustics analyses and other physical phenomena to address design challenges inherent to loudspeaker and soundbar designs.

The COMSOL Multiphysics® software is used for simulating designs, devices, and processes in all fields of engineering, manufacturing, and scientific research. See how you can apply it to your loudspeaker designs.

comsol.blog/loudspeaker-design

COMSOL

The xCORE VocalFusion Speaker Linear Kit by XMOS

How Many Microphones Does It Take to Run Your Life?

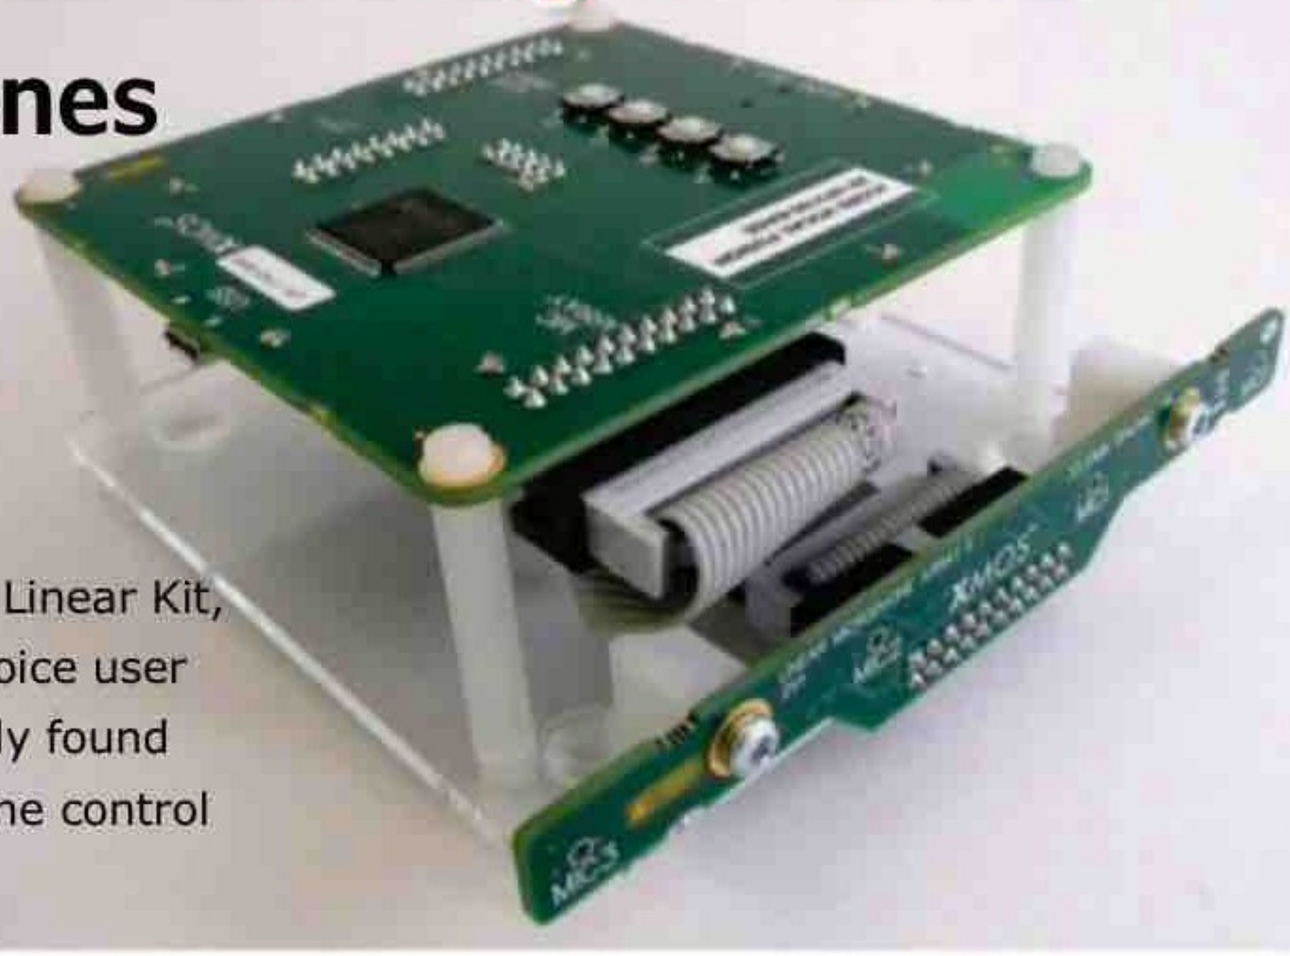

Photo 1: VocalFusion is an XMOS platform evaluation kit for far-field voice capture and voice processing using linear arrays of microphones.

Learn more about the VocalFusion Speaker Linear Kit, designed for those who want to integrate voice user interfaces into consumer electronics typically found at the edge of the room, such as smart home control panels and washing machines.

By
David Logvin
(United States)

In my home, I have 49 lightbulbs, which is only slightly higher than the average number of 45 found in the typical US household. The majority are LED, with a small minority of CFL, and just a handful of incandescent bulbs. How did we end up with so many lighting options and sources in today's home? What other aspects of technology are increasing in our houses without us noticing? With the advent of smart devices—everything from our phones to remote controls to household appliances—we probably all have more microphones around us than we realize.

Let's assume that most households have at least five mobile devices and that each device is equipped with one or two microphones. Add a few more if the household uses a voice-assisted remote or dictation software application.

This number dramatically increases with the rise of Smart Speakers and Voice User Interfaces (VUI) in the past few years. While each of these smart audio devices can be utilized with just a single microphone, today they are being built with as many as seven microphones per device for optimal performance. With a smart speaker in every bedroom and a voice-controlled assistant in the kitchen or living room, a future with 45-plus microphones per household is rapidly approaching.

Enter VocalFusion

VocalFusion is an XMOS platform that aims to fast-track development of smart speaker and voice-assisted communication products that are specifically aimed at the far-field user (see **Photo 1**). The use cases are quite varied, but it is generally targeted at solving the many typical acoustic problems that exist in the real world. XMOS likes to divide the market into center-of-room applications (Smart Home—where incoming voice is coming from 360°) and edge of room (Smart TV or Video—where the microphones are typically at a room boundary where incoming voice commands originate at 180° of coverage).

I have worked on consumer digital room correction systems and Smart Home IoT products in the past, so when offered the opportunity to explore the xCORE VocalFusion Speaker Linear Kit by XMOS (XK-VF3100-L33), I was intrigued and excited to learn about it and share my first-hand impressions of the developer capabilities this product enables.

Briefly, the challenges include users who are actively moving and interacting with voice commands in the far field, with less than ideal passive acoustic conditions such as high levels of reverberation and early reflections, double talking via other "distractor" humans, and various noise sources that bleed into

the listening space and reduce speech intelligibility.

There are readily available standard firmware libraries of DSP functions that can be broadly divided into two categories: Acoustic Echo Cancellation (AEC) and Adaptive Beamforming/Post-Processing.

AEC cleans up the measured input signal by suppressing the playback signal that the voice processor is actually playing (as well as its derivative reflections). Adaptive beamforming steers the microphone pickup pattern toward the user who is actively speaking to better reject background noise interferers. De-reverberation and various noise reduction algorithms help improve audio quality and speech intelligibility after ~10 seconds of sampling the steady-state conditions that exist in the listening space. Automatic gain control allows the host processor to more easily adapt to a sound source that is often changing in both position and amplitude. Keyword Detection allows interfacing with third-party voice solutions, creating an impressive library of natural language processing that is only limited by the powers of the cloud.

Additional post-processing is configured depending upon whether the speech is transmitted to other humans for a VOIP call, or, if the intended recipient is a machine for Automatic Speech Recognition (ASR), then voice quality can be sacrificed in order to improve accuracy. It is perfectly acceptable to create multiple outputs, too, if perhaps you want to support a person-to-person call and simultaneously generate a highly accurate voice transcription.

The VocalFusion Speaker Linear Kit

The particular development board that I am exploring is the xCORE VocalFusion Speaker Linear Kit (XK-VF3100-L33). This model comprises a BaseBoard and a separate linear microphone array (see **Figure 1**).

The BaseBoard is built around the XVF3100-TQ128 device. Besides this processor and a four-channel digital microphone interface, inspection of this board reveals that there are 13 LEDs mounted on the top surface of the baseboard that are controlled by the GPIO, as well as four programmable general purpose push-button switches (see **Photo 2**). The development kit also includes the xTAG debugger, which is a must-have tool since that is how you will load and flash firmware onto the BaseBoard.

The linear microphone array is mounted on a separate PCB with four omni-directional PDM Infineon IM69D130 MEMS microphones with 33.3 mm spacing between elements. These microphones feature 69 dB(A) SNR, -3 dB at 28 Hz, 1% distortion at AOP 130 dB SPL, tight sensitivity matching (±1 dB), and

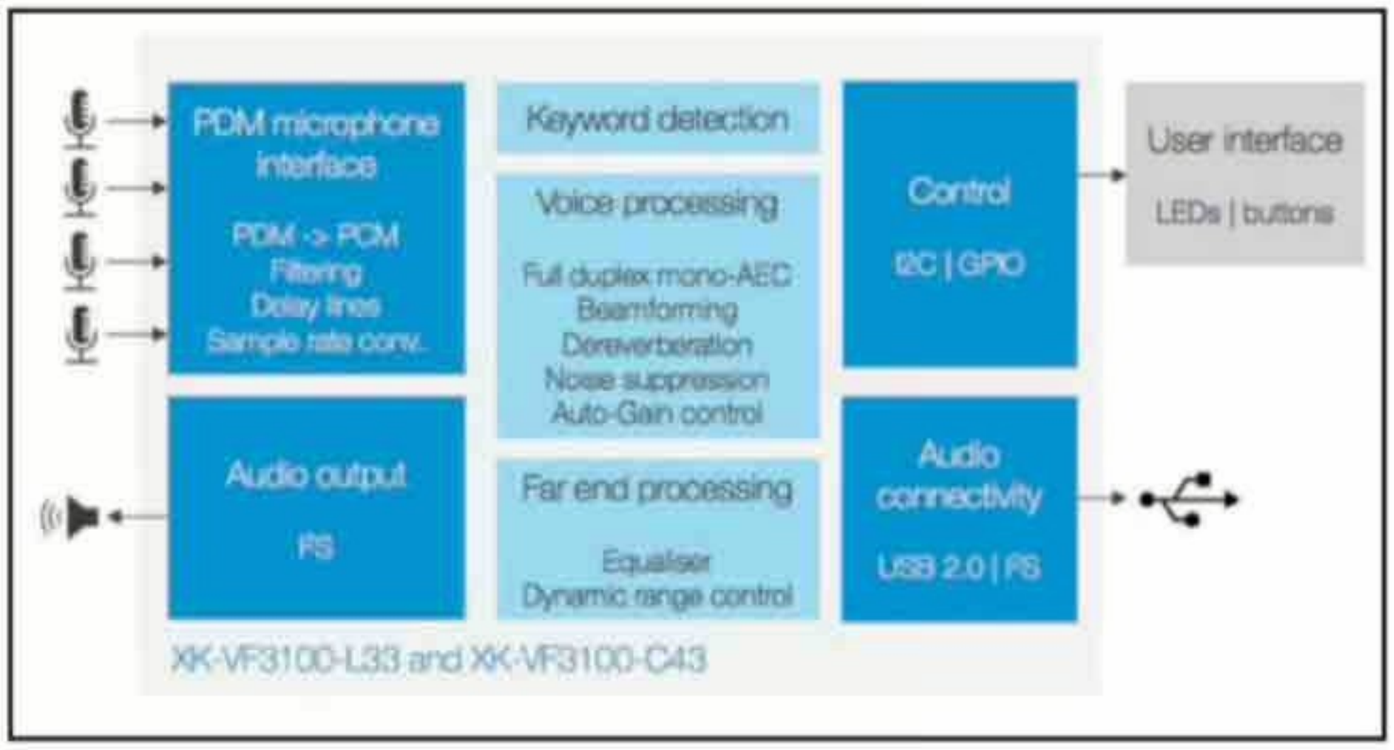

Figure 1: This is the block diagram for the XK-VF3100 development board.

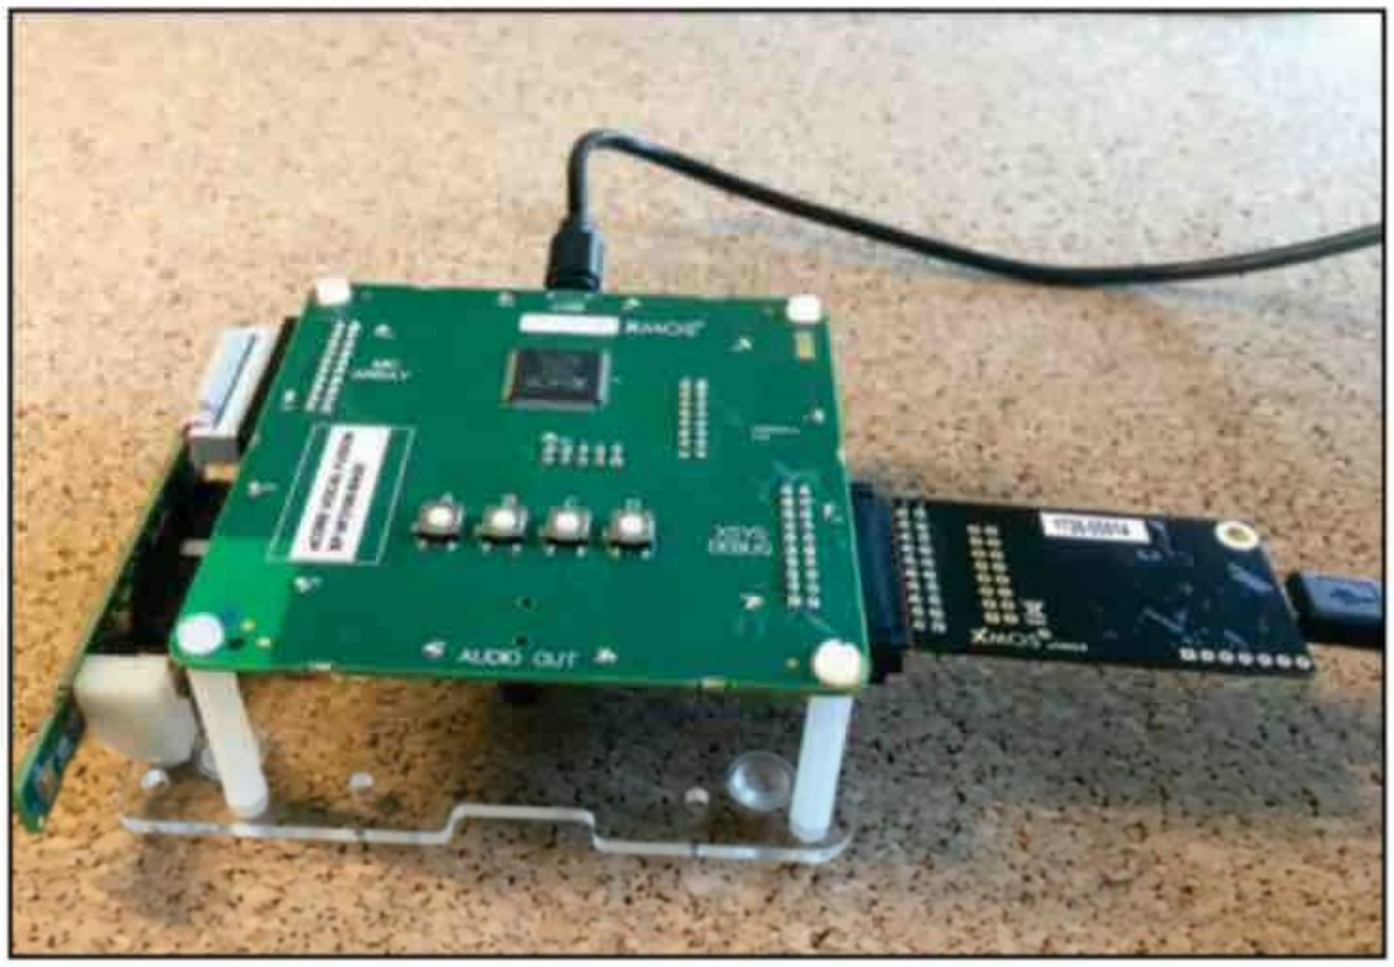

Photo 2: The BaseBoard contains the USB connections for the microphone array and the debug tool.

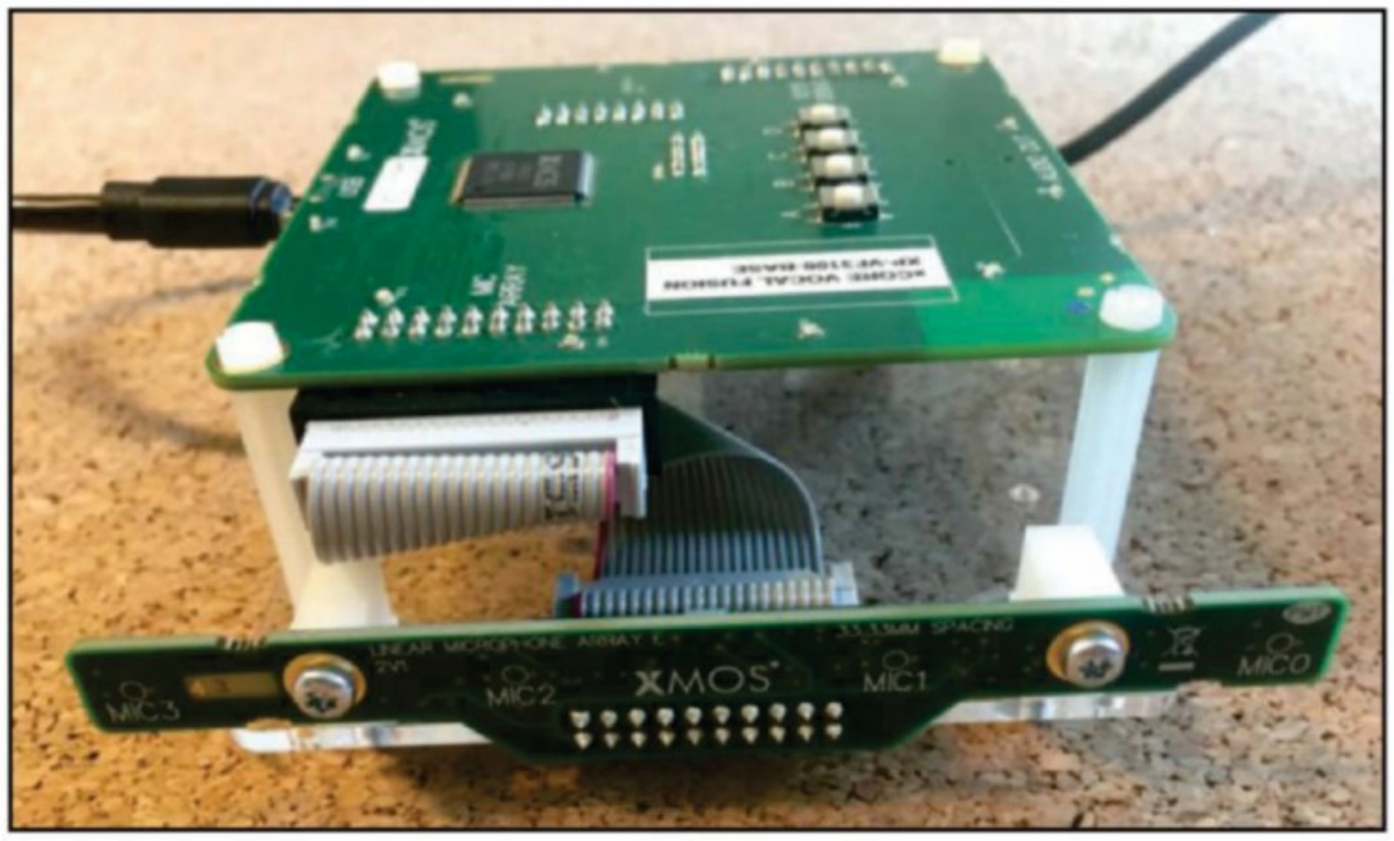

Photo 3: This close-up image shows the linear microphone array daughterboard.

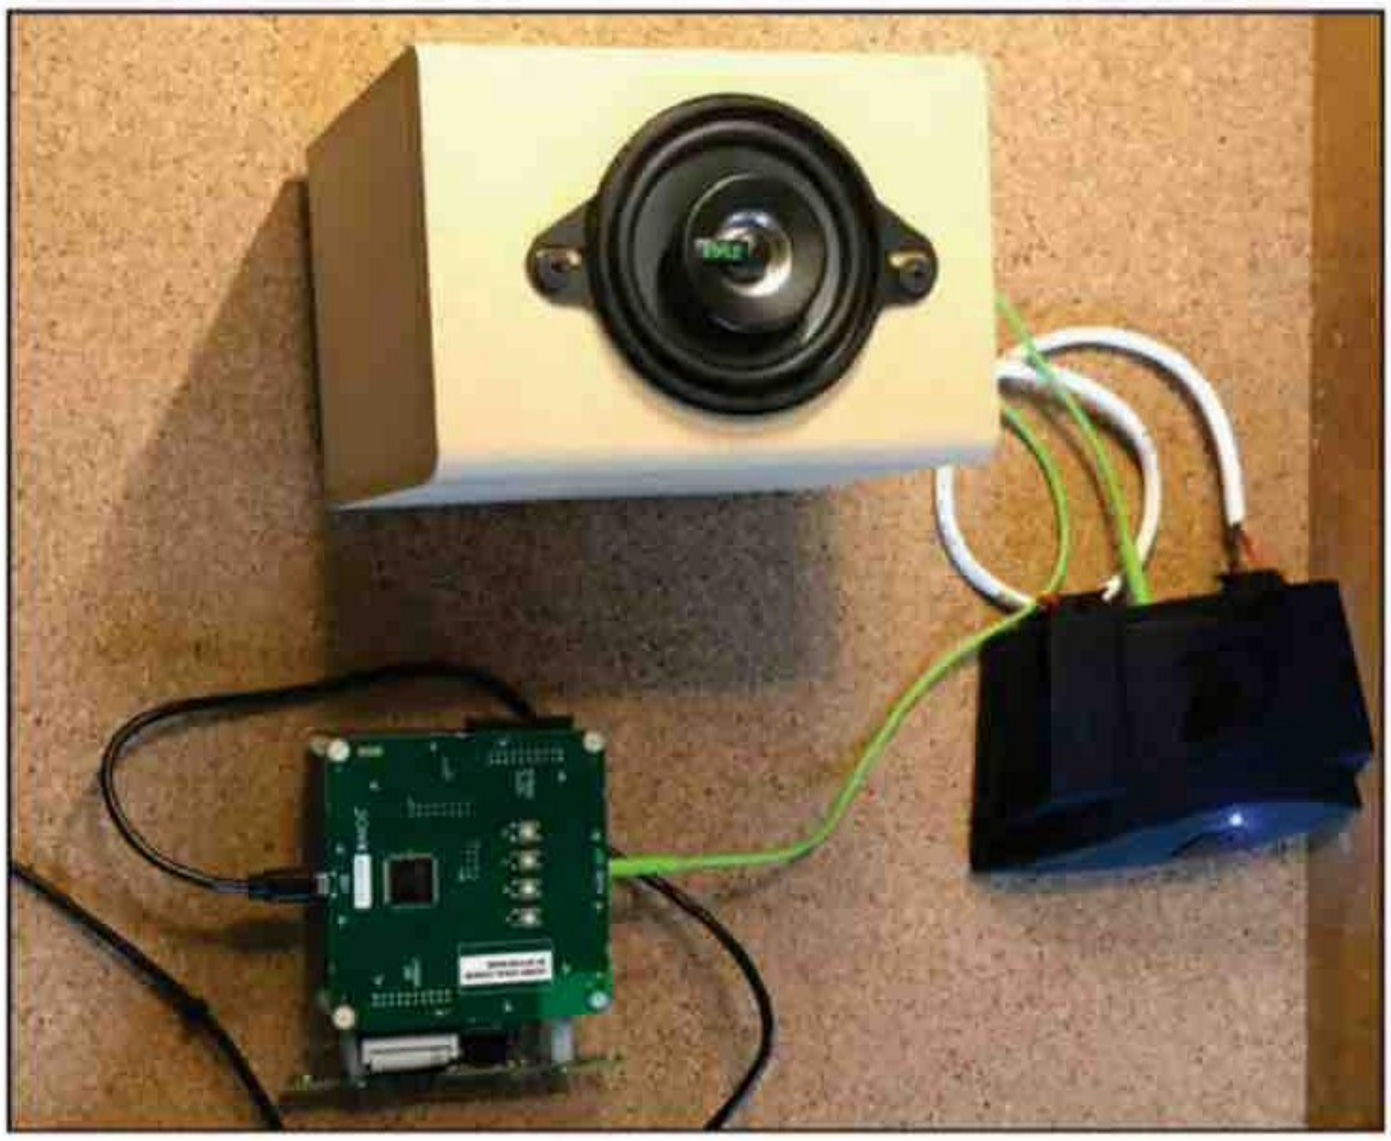

Photo 4: The test setup consists of a monophonic full-range speaker, an amplifier, a XMOS development board, and a host PC.

phase matching (±2°), which are critical parameters for microphone arrays.

The linear microphone array is intended for use applications at the edge of a room and provides 180° of microphone coverage (see **Photo 3**).

If you intend to create a product that would be used in the center of a room, XMOS offers the same kit, but with an optional circular four or seven microphone board. The only difference would be coverage would then extend to 360°.

There are additional permutations available for this product; notably, there are stereo versions available (built upon the xCORE XVF3500 voice processor) as well as versions built with Amazon Alexa Voice Service (AVS) natively supported.

Setting Up the Development Kit

The heart of this operation is built around a powerful Eclipse-based (Java) Integrated Development Environment (IDE) called xTIME Composer Studio. The Composer Studio is available for Windows, MacOS, and Linux and is provided to users at no additional charge. This package includes a full GUI and command line interface, as well as valuable debugging, scope, cycle simulation, explorer and context sensitive help utilities.

While digging into such a feature-rich product, the technical team at XMOS has dramatically minimized the learning curve thanks to a plethora of online documents. You will need to familiarize yourself with the Evaluation Kit Hardware Manual, the Quick Start Guide, and the XVF 3100 datasheet. The Quick Start Guide, in particular, is quite good, and will have you up and running and exploring product features in short order.

The initial task is to load an appropriate firmware binary onto the development board using the debug tool in order for your host computer to enumerate this hardware. I recommend using a modern OS such as Windows 10, and avoiding the use of USB 3.0 ports and USB hubs in order to easily load the firmware quickly and without difficulties.

Initial Explorations

XMOS recommends using the development board with a monophonic point source and no onboard sound processing. I used a simple 3.5" (88.9 mm) car audio coaxial driver in a small sealed box enclosure (1 ltr) and a small Class T digital amplifier with no onboard compressors or limiters in order to avoid introducing any nonlinearities into the test platform (see **Photo 4**).

For the initial inquiry, I recommend working through the AEC demonstration highlighted in the Quick Start Guide. First, I started VideoLan VLC as my music player and selected XMOS VocalFusion Spk as the playback device and played some content. Next, I opened up Audacity and enabled communication with the evaluation kit and created a new project. The project sample rate should be set to 16 kHz. This microphone sample will then yield a Nyquist frequency of 8 kHz, which is a perfectly reasonable upper limit for the task of recording human speech. I pressed the Record button, and

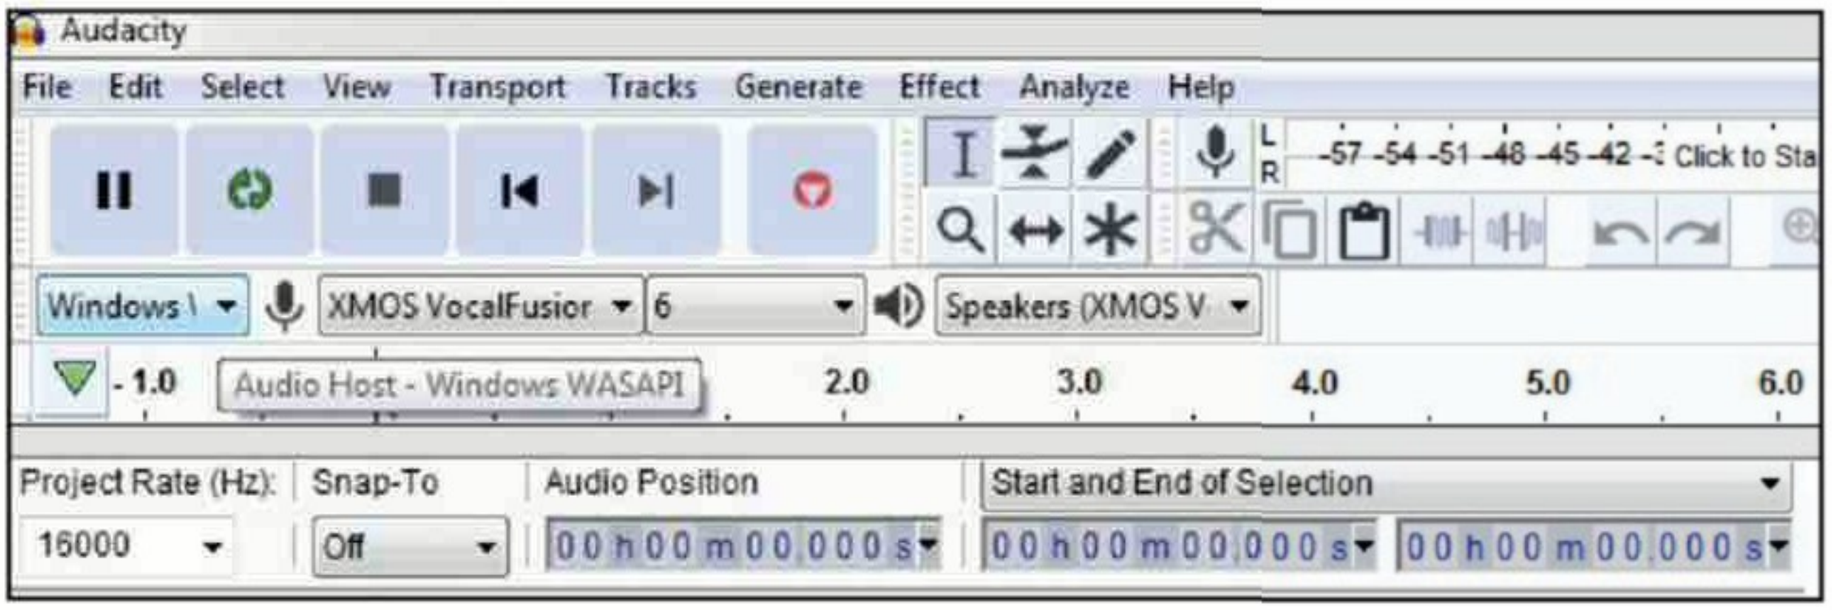

Photo 5: These are the Audacity settings for a six-channel audio capture.

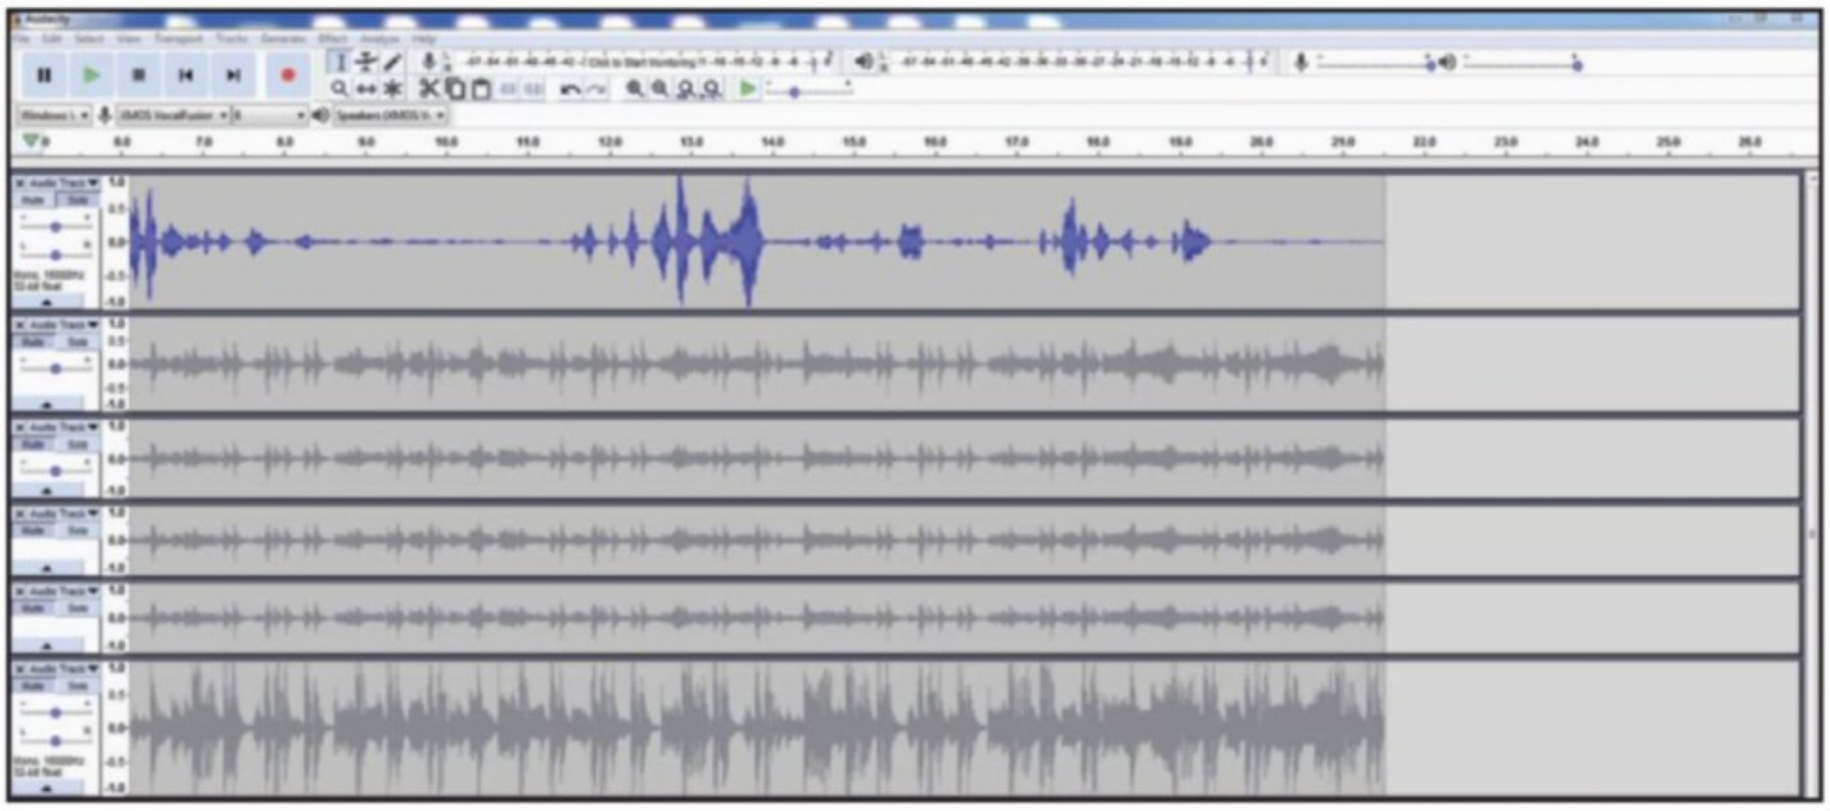

Photo 6: This screen shows the extracted, post-processed voice signal, which is highlighted in Track 1.

began speaking at various amplitude levels at various locations from several meters away while the audio playback continued (see **Photo 5**).

After I stopped recording, I then reviewed the six channels of audio that were captured. Audio tracks 2, 3, 4, and 5 correspond to microphone inputs 0, 1, 2, and 3 on the linear microphone array board. Track 6 is a capture of the mono playback stream that the VocalFusion was outputting. Track 1 is the real prize: this is the extracted voice signal after AEC, adaptive beamforming, and additional DSP. In this particular application, the Track 1 extracted signal has been optimized for human perception.

Parameters can be modified in the firmware using xTIME Composer to create an extracted signal that would be optimized for Automatic Speech Recognition (ASR), or separate outputs for both human perception and ASR (see **Photo 6**).

Another key feature is the Voice Activity Detector (VAD). When the XVF3100 detects voice activity over a particular threshold, it will briefly illuminate the center LED (D14), regardless of whether music is playing or not (see **Photo 7**).

A related feature to VAD is the Direction of Arrival

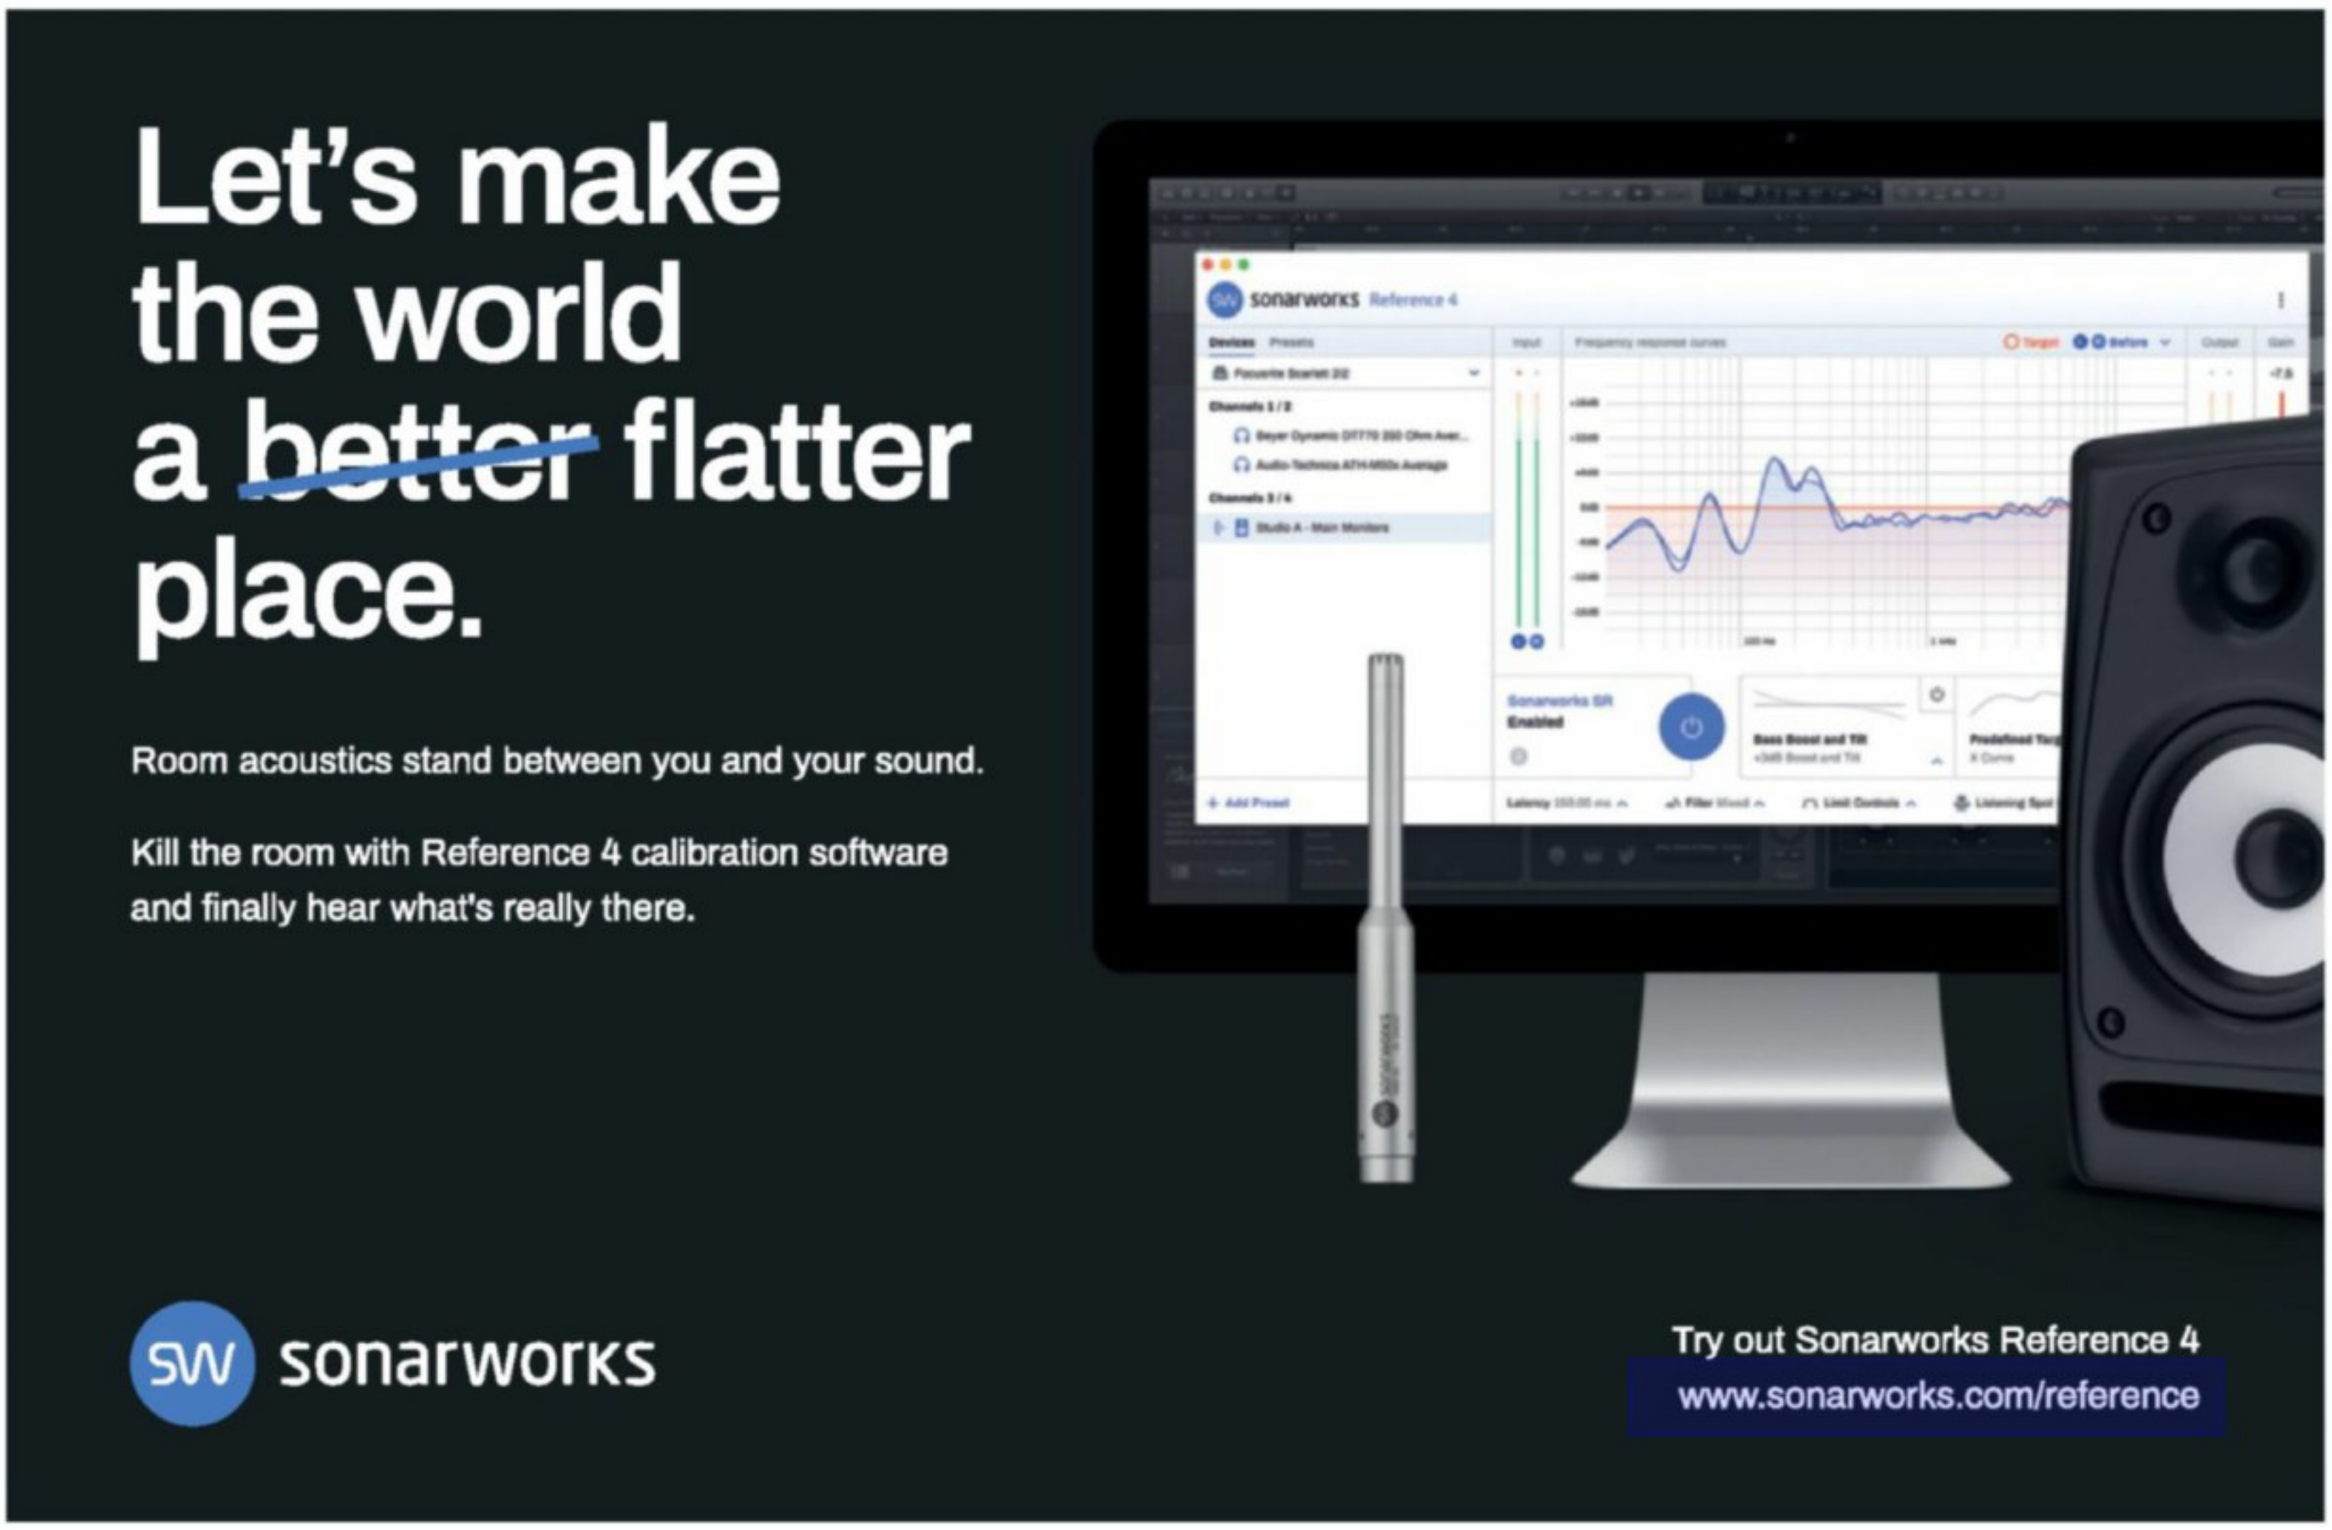

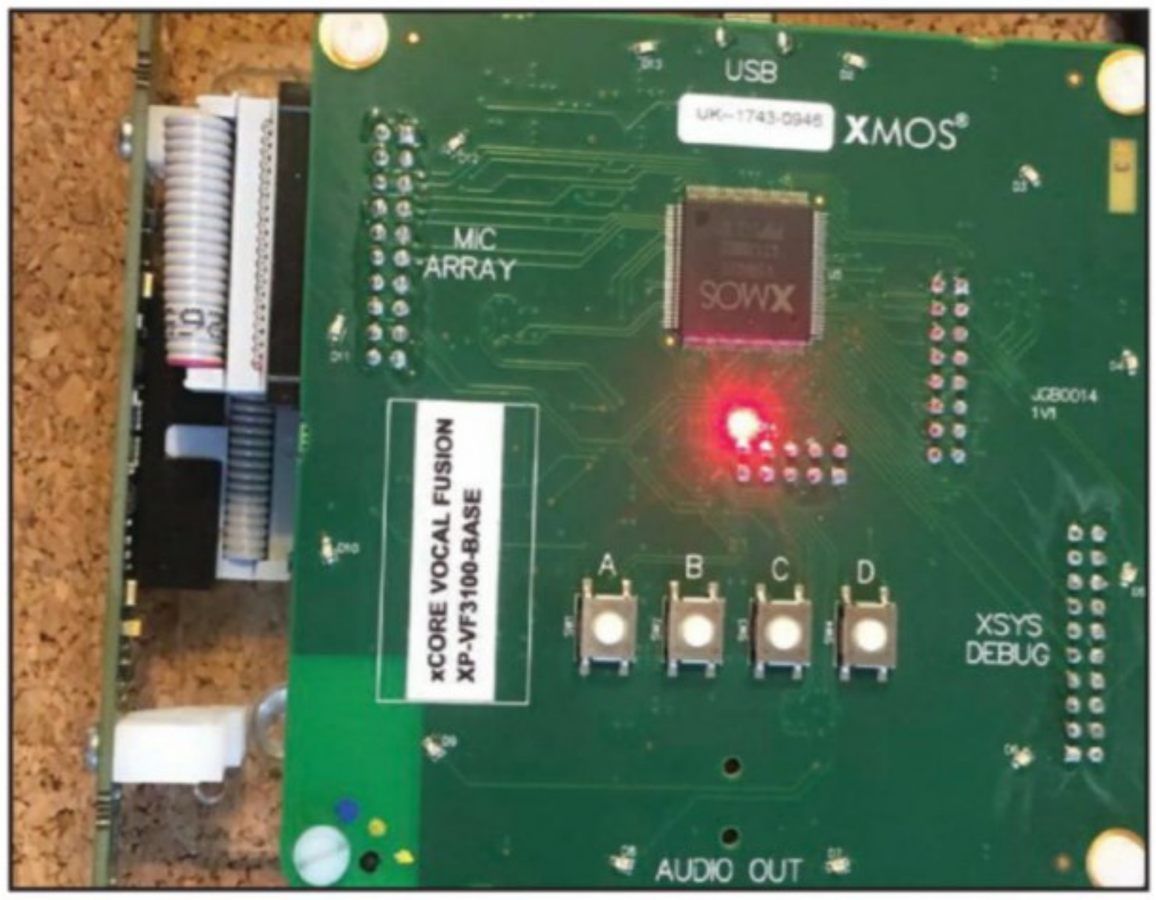

Photo 7: The center LED illuminates when voice activity is detected by the linear array, even with music playing.

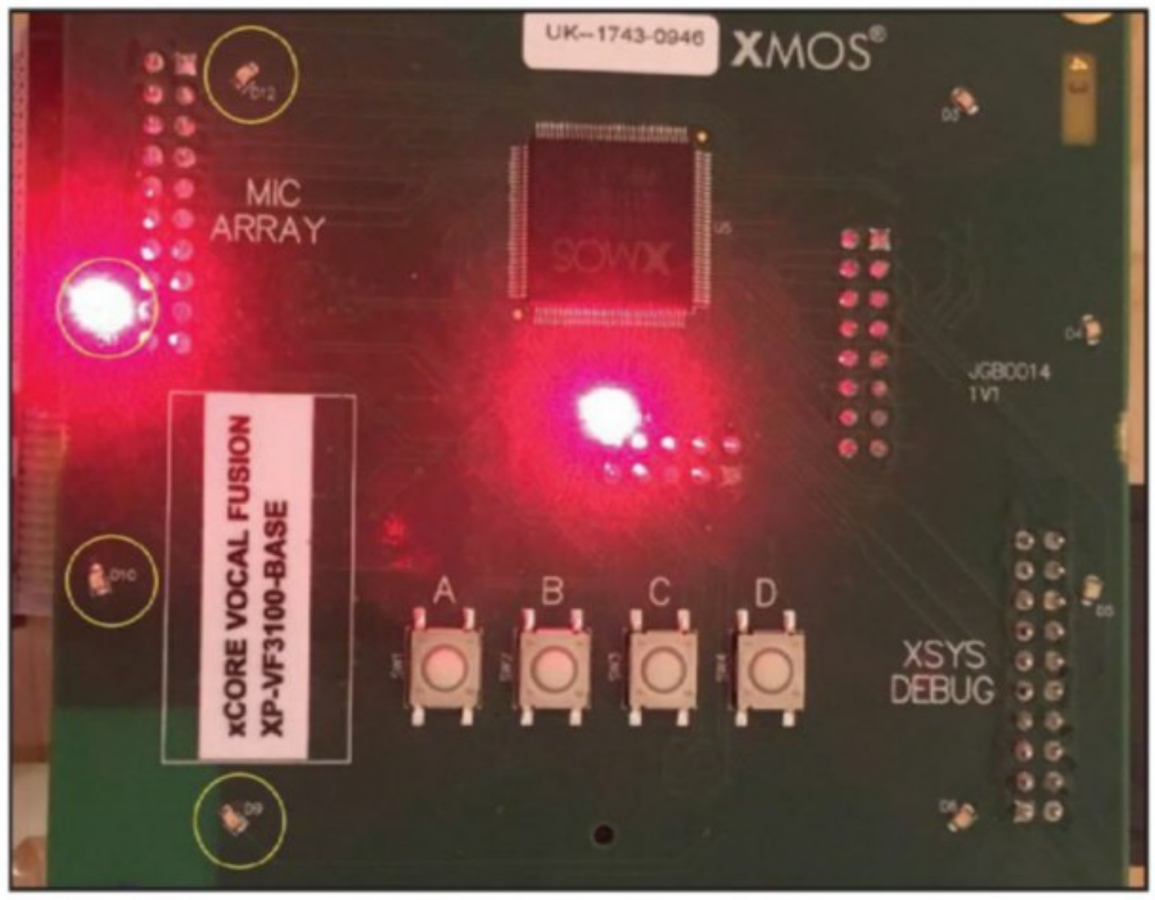

Photo 8: The four circled LEDs (D9, D10, D11, and D12) are used to indicate the direction of arrival of spoken voice as picked up by the linear array.

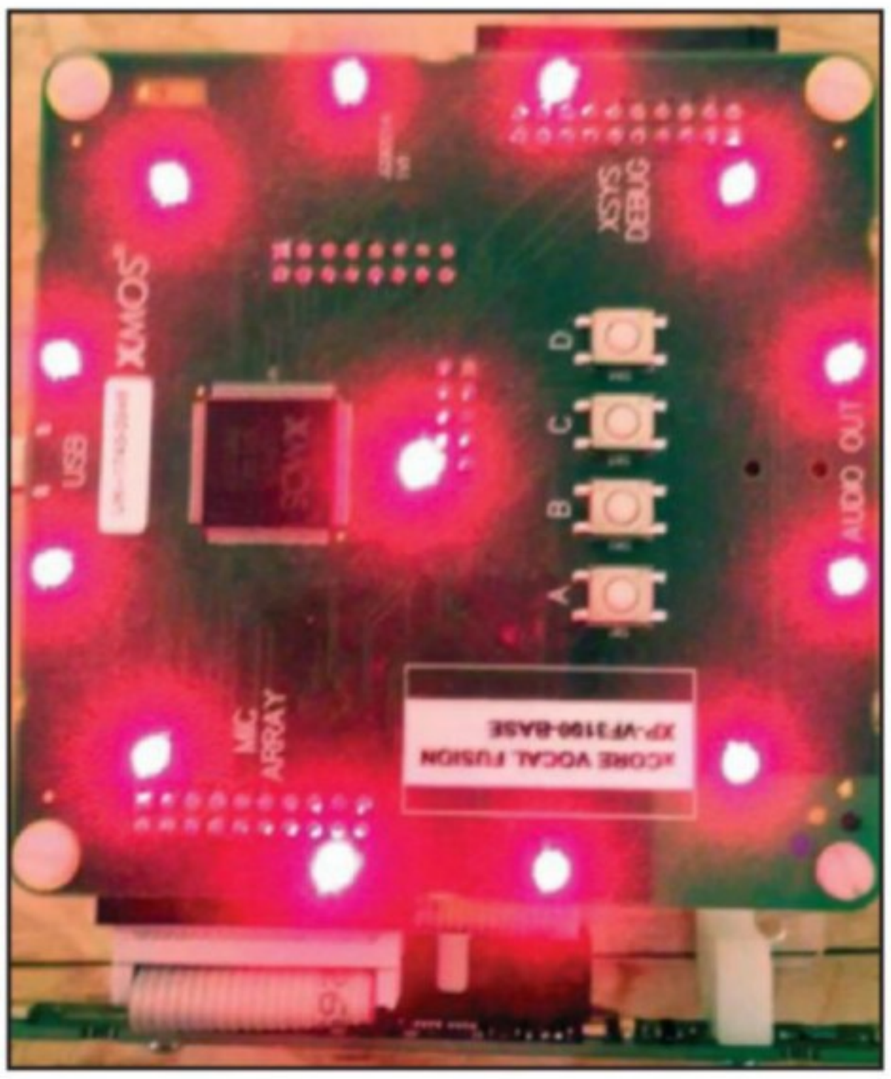

Photo 9: The ring of red LEDs illuminate after the keyword is spoken to indicate that the system is ready and waiting for additional commands.

notification. After the Keyword (aka "wake word;" see further down in the article for more details) is spoken, one of four LEDs on the top of the BaseBoard will illuminate to indicate the direction of arrival of the spoken voice. The expected behavior is that particular LED will then stay illuminated as long as that person continues speaking (see **Photo 8**).

This particular feature is not well-implemented with the linear mic array and is limited both by the coverage of the linear mic array for highly specific directional cues and potentially by the code itself. In testing, I found that only two (D10 and D11) of the four LEDs (D9, D10, D11, D12), would ever illuminate and the indicated direction was quite random.

I reached out to XMOS technical staff to discuss this, and they confirmed my observations with the linear microphone board, as well as my suspicion that this feature works quite well with the circular microphone array.

One of the prime functionalities required in a Smart Speaker is Keyword recognition. XMOS has

Resources

Getting Started with xCORE VocalFusion Speaker, XMOS, https://www.xmos.com/developer/download/private/sw_vocalfusion:-sw_vocalfusion-%5bGetting-Started-Guide%5d(1.1.2rc4).pdf

VocalFusion Evaluation Kit Hardware Manual, XMOS, https://www.xmos.com/developer/cn/documents/vocalfusion/kits

VocalFusion Fine Tuning Guide, XMOS, https://www.xmos.com/file/vocalfusion-fine-tuning-guide

VocalFusion Software Design Guide, XMOS, https://www.xmos.com/file/vocalfusion-stereo-software-design-guide

XVF 3100 Datasheet, XMOS, https://www.xmos.com/file/xvf3000_3100-tq128-datasheet

xCORE VocalFusion Evaluation Kit Quick Start Guide, XMOS, https://www.xmos.com/file/xcore-vocalfusion-speaker-evaluation-kit-quick-start-guide

licensed TrulyHandsfree voice control from Sensory. When the board detects the phrase "Hello Blue Genie," the system wakes up so that it is actively listening for additional VUI commands. This wake up can occur either with music playing or no music, thanks to AEC.

This is confirmed with a visual cue on the board, where the ring of 13 red LEDs illuminate, confirming that the system is ready. In a finished product, additional commands would be loaded into the system (for track and volume changes, mute, accepting a phone call, etc.). Adding these additional commands would be completed by directly working with Sensory. Alternatively, voice commands could be sent to a third-party voice service to leverage natural language cloud processing (see **Photo 9**).

Adapting the Design

With a functional platform, the next steps are to understand how to make software modifications needed to evolve from development hardware to a fully-customized project.

The XK-VF3100-L33 kit is built around a four-microphone, 100 mm linear array, with 33 mm inter-mic spacing. To mathematically describe the positions of the four kit microphones, the following Cartesian coordinates (X, Y, Z) would be entered into the XBECLEAR_MIC_GEO command (see **Table 1**), as described in the *VocalFusion Fine Tuning Design Guide*.

Accurate coordinates are required for the AEC and Beamforming algorithm functionality and are used to perform all computations based on fixed microphone delays and positions.

The actual spacing and number of microphone elements is obviously dependent on your system's industrial design (case and housing) and the PCB layout to determine the final microphone locations. Alternative positions and number of microphone elements are simple to update.

About the Author

David Logvin is an acoustics engineer at TDK InvenSense, working in the MEMS microphone division. With more than 20 years of experience, he has designed loudspeakers and other products for Snell Acoustics, Outlaw Audio, ClearView Audio, Cambridge Sound Management, and others. During his 14 years at Snell Acoustics, he developed and designed dozens of high-end loudspeakers and amplifiers, as well as helped to create loudspeaker standards in close collaboration with THX. David then focused his efforts on R&D, creating innovative, award-winning piezo-driven membrane loudspeakers with ClearView Audio (formerly known as Emo Labs). At Cambridge Sound Management, he put his skills in room acoustics to use in the field of sound masking for corporate applications. Most recently, he rejoined his fellow Snell Acoustics alumni, Dr. Joseph D'Appolito, to design several new loudspeakers for Outlaw Audio. In addition to his lifelong love of music and acoustics, David is an active participant in the Boston tabletop game design community. He is a member of the Audio Engineering Society (AES) and the Boston AES Section and lives in Lowell, MA with his family.

Table 1: To mathematically describe the positions of the four kit microphones, the following Cartesian coordinates (X, Y, Z) would be entered into the XBECLEAR_MIC_GEO command

Microphone Number	X (mm)	Y (mm)	Z (mm)
0	0	0	0
1	100 mm	0	0
2	33 mm	0	0
3	66 mm	0	0

To work through the needed customizations for a real product, it is necessary to use the control tuning guide to determine the optimal parameters. The end goal is to build a customized firmware binary, based on the desired application requirements and physical geometry and installation use case.

The first step is to build a host control tool in order to read and write control parameters. XMOS makes available source code, so you can build the tool in Window, Linux, Raspbian or Mac OS.

If your preferred operating system is Windows, use Microsoft Visual Studio to build a USB host control. Once this is completed, you will now have an executable command line utility. From the xTIME Composer Studio CLI (Command Line Interface), you can launch the tool vfctrl_usb.exe.

Control Parameter Examples

High-Pass Filters

In certain applications, it may be critical to select the proper high-pass filter on the mic array to control the amount of low-frequency energy that is recorded. Control parameters provide four options: Flat, 70 Hz (default), 125 Hz, and 180 Hz (see **Photo 10**). For the non-flat options, the high-pass filter is a fixed fourth order (24 dB/octave slope).

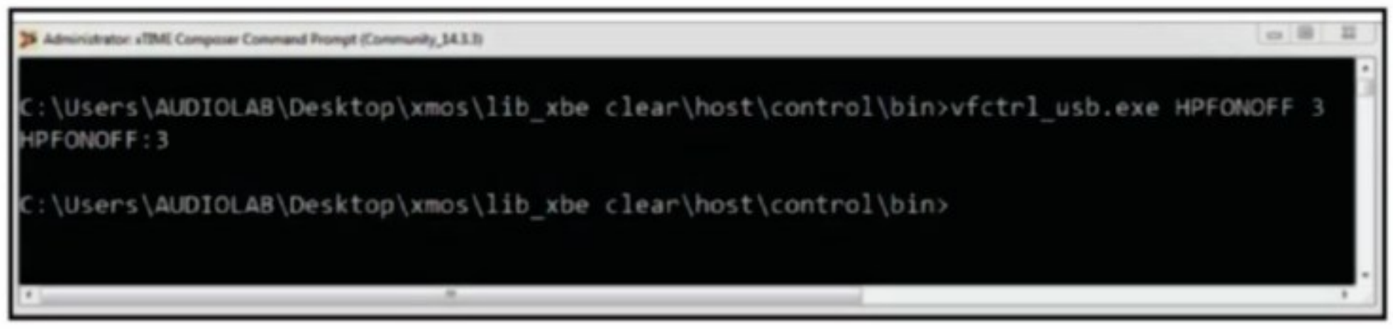

Photo 10: This CLI example demonstrates how to change the high-pass filter frequency from the 70 Hz default to 180 Hz.

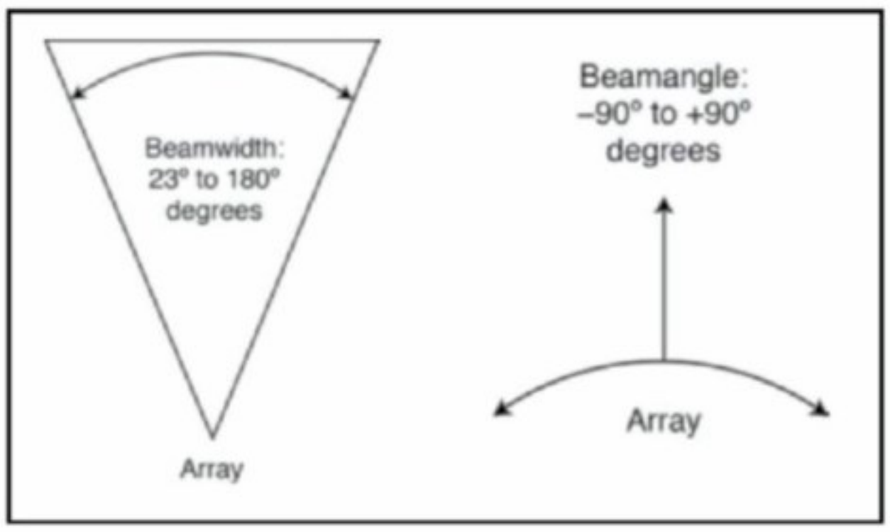

Figure 2: Here are the BEAMANGLE and BEAMWIDTH diagrams for the Beam Forming Array.

Adjusting BFA

The default width of the BFA is 60° with a beam angle of 0°. Beam width can be varied from 23° to 180°. Beam angle can be adjusted to any angular measure between -90° to +90°. BEAMANGLE and BEAMWIDTH parameters might need adjustment if the installed product is meant for a corner mount as opposed to a centrally located location (see **Figure 2**).

Fine Tuning a Design

There are 32 control parameters for VocalFusion, along with 15 parameters that are specific to the linear microphone array.

After the desired settings for any control parameters have been determined, the next steps are to build a customized firmware, erase the existing firmware, and flash this new code onto the XMOS processor. This is a necessary requirement in order to preserve control parameter settings in non-volatile memory.

After determining optimal parameter settings, the custom firmware should be built and loaded onto a working prototype with microphones installed in a representative housing. Then the system should be thoroughly tested in a variety of room types, with varying levels of background noise. This final round of testing will determine if any additional adjustments to the control parameters are needed to improve far field performance of the AEC and beam forming algorithms.

During this part of the development, the objective is to determine if further refinements to firmware are required for the AEC & Beamforming algorithms for optimal performance. Further details on this interesting engineering process are detailed and described in the *VocalFusion Fine Tuning Guide* and *VocalFusion Software Design Guide*.

Observations

In a world of constant monitoring and adaptive audio, the XMOS voice processors provide robust and flexible algorithms for enhanced performance in typical audio spaces. With a variety of control support, as well extensive tools for Windows, MacOS, Raspbian, and Linux, the command line and IDE software tools are extensive and powerful. XMOS VocalFusion offers a real opportunity to fully realize your smart audio project's true potential.

A future article will take a deeper look at the DSP algorithms, acoustical tuning, and field measurements. This will allow for a further exploration of the VocalFusion voice processors' capabilities in enhancing speech intelligibility in real-world environments.

DESIGNATED DRIVER

Leading Pro Audio brands choose Celestion

Innovation, performance and reliability make Celestion the drivers of choice for a growing number of big-name PA brands. So when you need superior quality compression drivers or professional loudspeakers, check out about the speakers the pros use.

CELESTION

Reggie-Phil 2.0: A Current-Regulated Filament Supply

Inspired by previous approaches published by *audioXpress*, this article explores alternative options to the design of a linear current regulator for audio tube amplifiers, including employing feedback to sense load and other design strategies.

By
Christopher Paul
(United States)

audioXpress recently reprinted Bill Reeve's noteworthy article, "Constant Current for Tube Heaters. Extend the Life of Your Amplifier's Vacuum Tubes," (https://www.audioxpress.com/article/constant-current-for-heater-tubes-extend-the-life-of-your-amplifier-s-vacuum-tubes). The title pretty much describes what it's all about. Reeve presents a simple but effective 150 mA linear current regulator for a pair of series-connected 12AX7 heaters, which are themselves connected in series.

He points out several advantages to this approach. Perhaps chief among them is the elimination of turn-on surge currents through cold filaments, with the expectation of extending tube life. Another is that the slowly increasing heater voltage can be used to trigger a comparator to enable a B+ supply. Furthermore, if one of the tubes were removed, or its filament opened or shorted, the remaining tube's filament would be undamaged (although a comparator-controlled B+ supply could be turned on immediately. Activating B+ should also require that a second comparator senses that current is being drawn.)

To those cited benefits, I would add additional line voltage ripple rejection—in this specific case almost 30 dB thereof. This is in comparison to what the heaters would see if the regulator were eliminated and the unregulated DC voltage reduced from the 30 V the article cites to 2 × 12.6 V = 25.6 V.

As useful as this approach is for a pair of low power filaments, the load on small signal devices such as we would find in a preamp is only about 3.8 W. This means that the regulator circuit dissipation is only (30 V - 25.2 V) × 0.15 A = 0.68 W. If we were to extend this approach to a power amp, we would really begin to "feel the heat."

KT88s and 6550s each need 1.6 A. Stereo push-pull amplifiers composed of either demand at least 6.4 A. Adding one 6DJ8 per channel would require a total of 7.13 A. Two would bring the demand to 7.86 A—a 50 W load! If the relative power required of the linear regulator scales, it would dissipate 9.5 W! Even if this much dissipation were tolerated, the filaments of such different tubes could not operate in series. We need to consider an alternate approach for power amps. (By the way, a filament is a directly heated cathode, while a heater is electrically isolated from the cathode it heats. The regulator discussed here can be used in both cases, and I'll use these terms interchangeably.)

Such was considered in the October 2016 issue

of *audioXpress*: "Reggie-Phil: A Regulated Filament Supply." The design employed an under $20 Meanwell 75 W 12 VDC output off-line switcher. It drove a $25 Texas Instruments (TI) TPS54821 DC-DC (switching) converter Evaluation Module (EVM), reconfigured for a 6.3 V output and capable of supplying 8 A to a load. In that design, the (differing current-rated) filaments are in parallel. While Reggie-Phil offered the advantages Reeve listed, it had a knock against it—it took almost two minutes for the filaments to warm up! This was a consequence of regulation by load voltage rather than current and the use of a slow start-up, constant volts-per-second rise rate. This prevented activation of the DC-DC converter's short circuit protection as well as current flow in excess of the filaments' ratings during start-up.

Updating the Design

Suppose instead, the design employed feedback that sensed load current instead of voltage. That would be a good first step, but it's not the complete solution. If a tube's filament or its contacts with its socket were to open, the remaining active filaments would share an increase in current that the "open" filament no longer consumed. Obviously, this would exceed the active filament ratings. Feedback needs to take into account and exceed neither the specified current nor voltage ratings of the filament assembly. There is a class of commercially available ICs that does just that—lithium ion battery chargers. They employ what is called constant current/constant voltage (CCCV) control. It would have been nice to be able to use one of these, but: (1) these ICs supply only a small current until the batteries reach a minimum voltage. The heaters would never reach this voltage at such low currents. (2) The cost of EVMs for these ICs is in the $100 rather than the $25 range.

Fortunately, we can implement CCCV control as these ICs do. Please refer to **Figure 1**. We place a small (to limit dissipation) sense resistor in series with the heater. One differential amplifier senses the voltage across the heaters and applies an attenuation factor such that when the heaters are warmed up, its output voltage is V_f. (V_f is the regulator's feedback voltage when its loop is closed and it's operating properly.) A second amplifier amplifies the voltage across the sense resistor. This voltage is the product of the intended heater current and the sense resistor's resistance. The gain is chosen so that again, the amplifier output is V_f. Each amplifier drives an "ideal" diode, one with no voltage drop across it. The amplifier producing the higher output and its diode "win" and control the loop. (In a practical circuit, we enclose the diodes in feedback loops so that they act in an ideal fashion.)

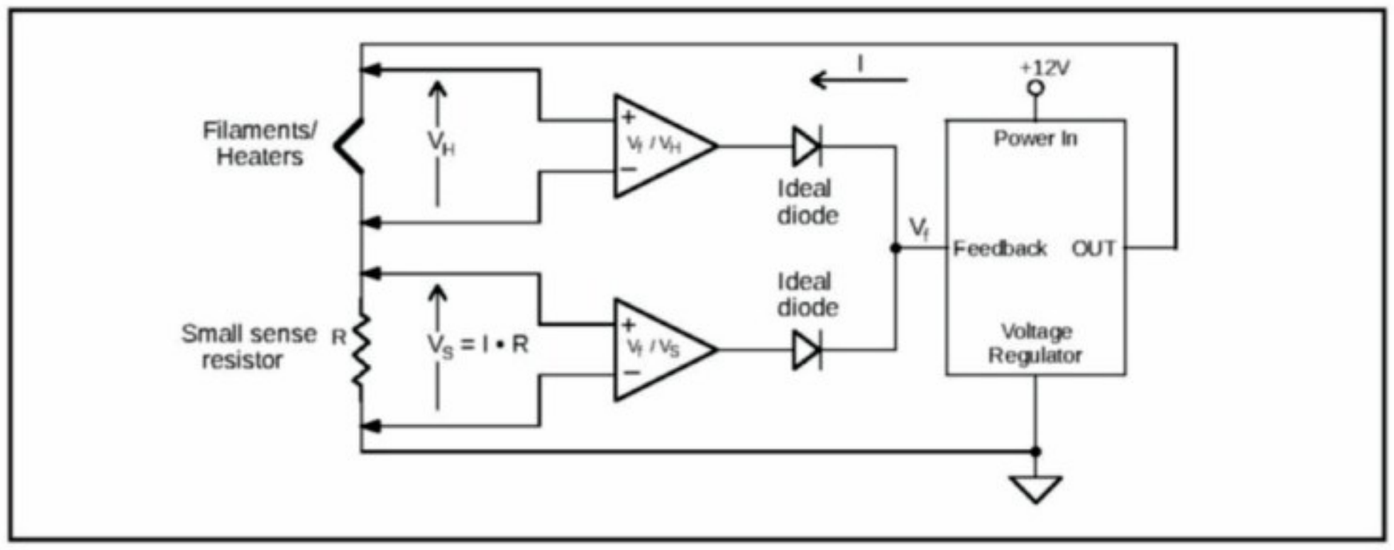

Figure 1: The block diagram shows the function of a constant current/constant voltage (CCCV) regulator.

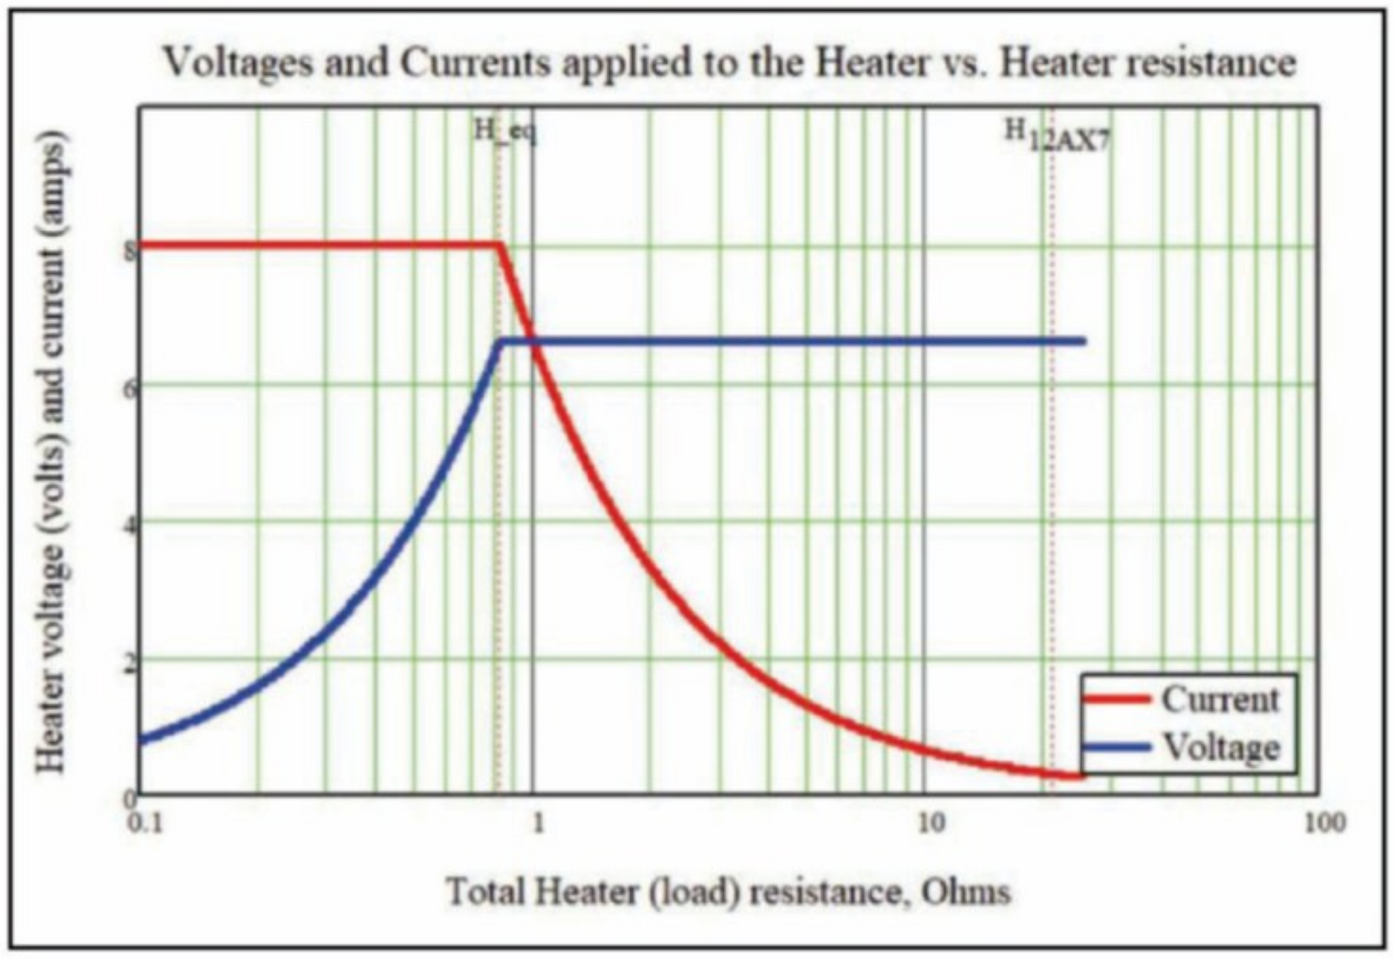

Figure 2: An example of how an 8 A constant current, 6.3 V constant voltage heater supply works with various heater resistances. Resistance typically varies from a little more than 0.1 Ω when cold to about 0.8 Ω when warmed-up.

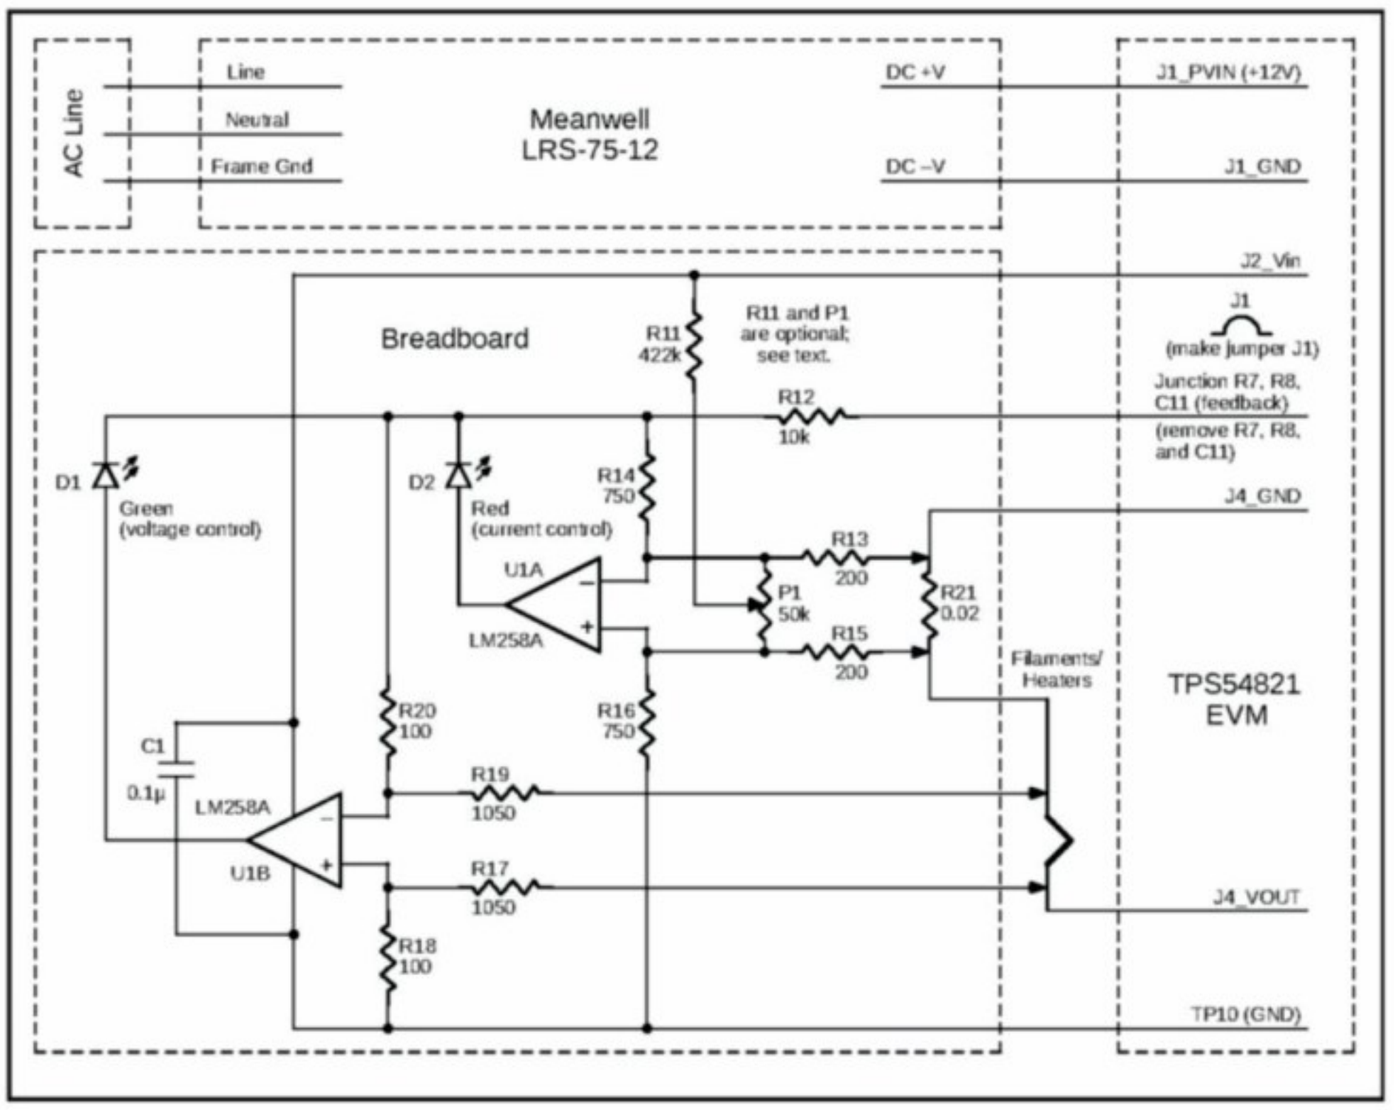

Figure 3: A schematic of a breadboard for an 8 A constant current, 6.3 V constant voltage heater supply, the connections between it, a modified TPS54821 Texas Instruments EVM, and a Meanwell 12 V supply

When the heaters are cold or a filament is shorted, heater resistance is low and therefore so is the heater voltage, so the scaled current sample controls the loop and the regulator supplies the designed-for current. If a filament opens or is disconnected when the remaining filaments are fully heated, the resistance becomes more than expected and the voltage sample limits the voltage. When operating as intended and fully heated, the current and voltage feedback signals "tie" (or are nearly the same) and it doesn't matter which one wins. We can call this heater resistance H_eq. **Figure 2** shows how voltage and current vary with heater resistance. An H_eq of 6.3 V/8 A and the resistance of a single 12AX7 heater (6.3 V/300 mA) are shown on the graph for reference.

This approach ensures that when fully heated, no filament will be damaged due to the absence of another. It's not perfect—during start-up, if another designed-for filament is absent, it does not prevent these cold filaments from receiving more than their rated currents. However, this is no worse than the standard approach, which depends on transformer saturation to limit current. **Figure 3** shows one implementation of this scheme that can be assembled on a breadboard and how it can be connected to a slightly modified TI TPS54821 EVM regulator (see **Figure 4**). The circuit works as follows:

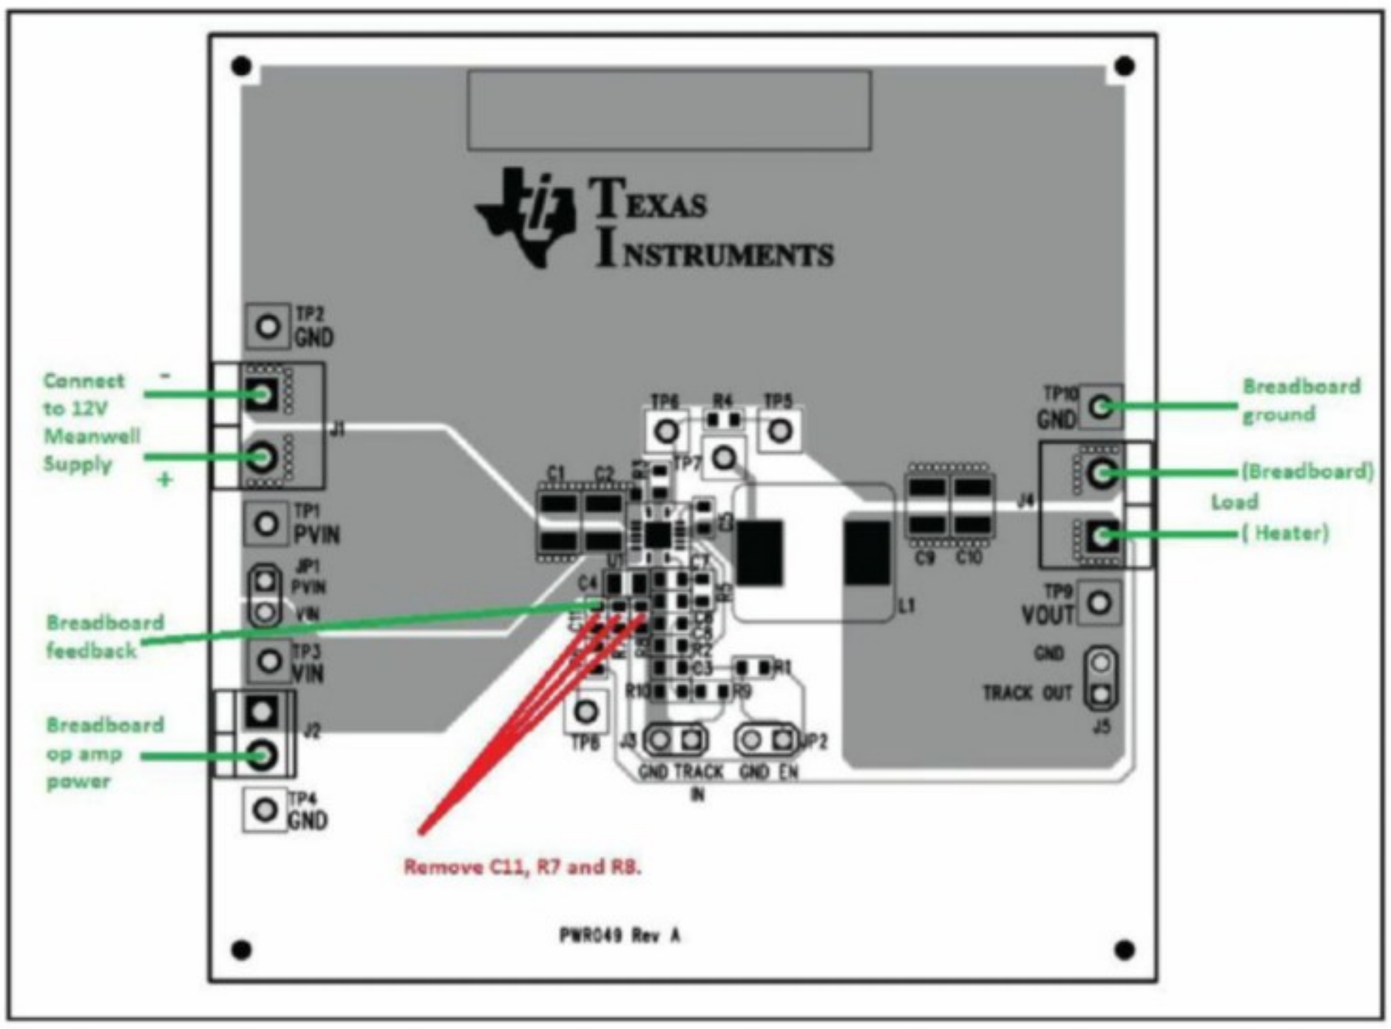

Figure 4: The TPS54821 EVM (courtesy of Texas Instruments), connections to it, and how it is modified for use as described in this article

The Meanwell off-line switcher supplies 12 V to the EVM. The EVM drives the Load (the series combination of sense resistor R_{21} and the paralleled filaments) through connector J4 (Vout and GND). The EVM's TP_{10} ground is also employed, but it's reserved for connection to low current paths on the breadboard. The breadboard's U_1 is powered through the EVM's J2_Vin, from the Meanwell's 12 V output. Breadboard R_{17} through R_{20} and U_{1b} sample, scale and attenuate the heater voltage. U_{1a} and R_{13} through R_{16} sample the heater current flowing

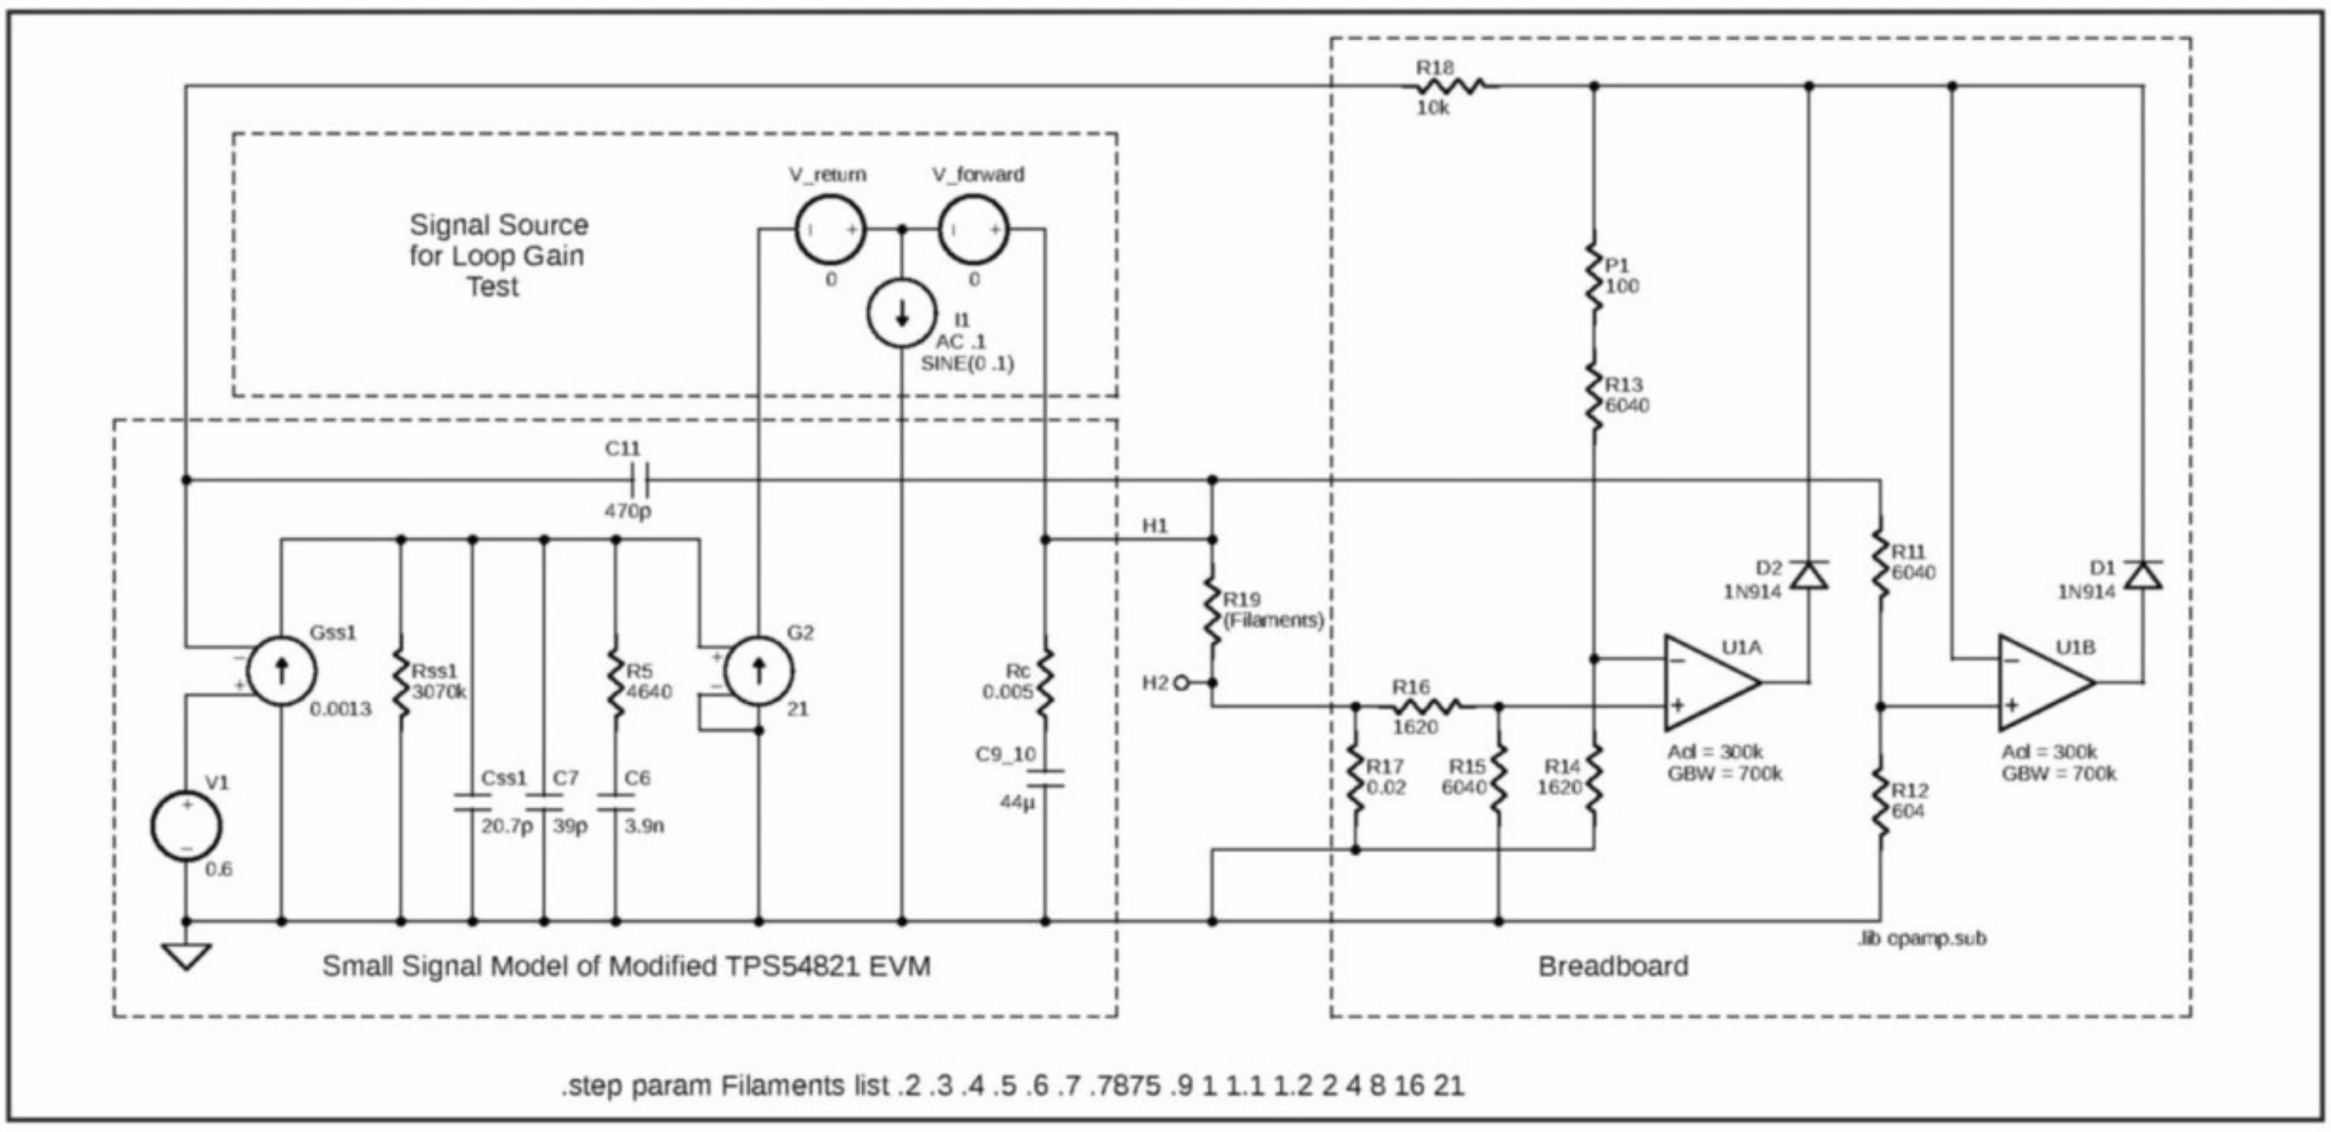

Figure 5a: The schematic of a simulation of a Reggie Phil 2.0 regulator

through R_{21} and amplify the voltage across it. U_1 drives D_1 and D_2, which vie to apply feedback to the regulator through R_{12} from the diode anodes. Enclosing them in the op-amp feedback loop implements "ideal" diode functions. Diodes could have been used for D_1 and D_2, but the bicolor LEDs supply a simple visual indication of which sampled signal has "won."

The venerable LM258A is an excellent choice for the op-amp. It's cheap, easily tolerates the 12 V provided by the Meanwell supply, and is specified for a maximum input offset voltage of 4 mV across its industrial temperature range at voltages near ground, where it's employed. (Many modern op-amps specify offset at half their supply voltage.) R_{11} and P_1 offer the opportunity to null out this offset.

The modification of the TPS54821 EVM feedback loop brings up concerns about stability and oscillation. The load varies as heater resistance increases from cold to warmed up conditions. Further changes could result if one or more of the intended heaters were not present. This wide range of heater resistance leads to notable variations in loop gain. I tested the parallel combination of two 6CD6GAs, a 6AU4GTA, and a 6BQ6GTB. When warmed up, it drew 8 A at 6.3 V. There was no evidence of instability from start-up to full warm-up. I then repeated the test with a single 6DJ8 heater as the load. Again, everything looked stable. But how much margin is there?

I used the small signal model of the EVM thoughtfully provided by TI, which I modified and augmented with the breadboard circuit, to run a simulation (see **Figure 5a**). Middlebrook's dual-injection technique for loop gain measurement was employed. Only the injection of a current source was required, since the regulator IC output was modelled as a pure current source. The two voltage sources are set to zero and are present so that there is something convenient to measure current through. Heater resistances were examined from 0.2 Ω through the nominal 0.8 Ω (6.3 V at 8 A) up through the 21 Ω of a single 12AX7 6.3 V-configurated heater. **Figure 5b** shows that the phase margin is a very comfortable 90° for all values.

A Design Strategy

We start the design process by selecting the aggregate filaments' heating voltage (V_H) and current I_H, 6.3 V and 8 A in this example. From the TPS54821's datasheet, the feedback voltage is 600 ±6 mV. An equation can be written from inspection to determine the voltage feedback path components:

$$\left[\frac{V_H \times R_{18}}{R_{18}+R_{17}} \pm V_{OS}\right] \times \left[\frac{R_{20}+R_{19}}{R_{20}}\right] = 600 \pm 6\ mV \qquad [1]$$

Solving for V_H,

$$V_H = \left[(600 \pm 6\ mV) \times \left(\frac{R_{20}}{R_{20}+R_{19}}\right) \pm V_{OS}\right] \times \left[\frac{R_{18}+R_{17}}{R_{18}}\right] \qquad [2]$$

Here 6 mV and V_{OS} = 4 mV are the regulator feedback and op-amp input offset voltages, respectively. Ignoring offsets and resistor tolerances:

$$\frac{R_{20}}{R_{19}} = \frac{R_{18}}{R_{17}} = \frac{600mV}{V_H} \qquad [3]$$

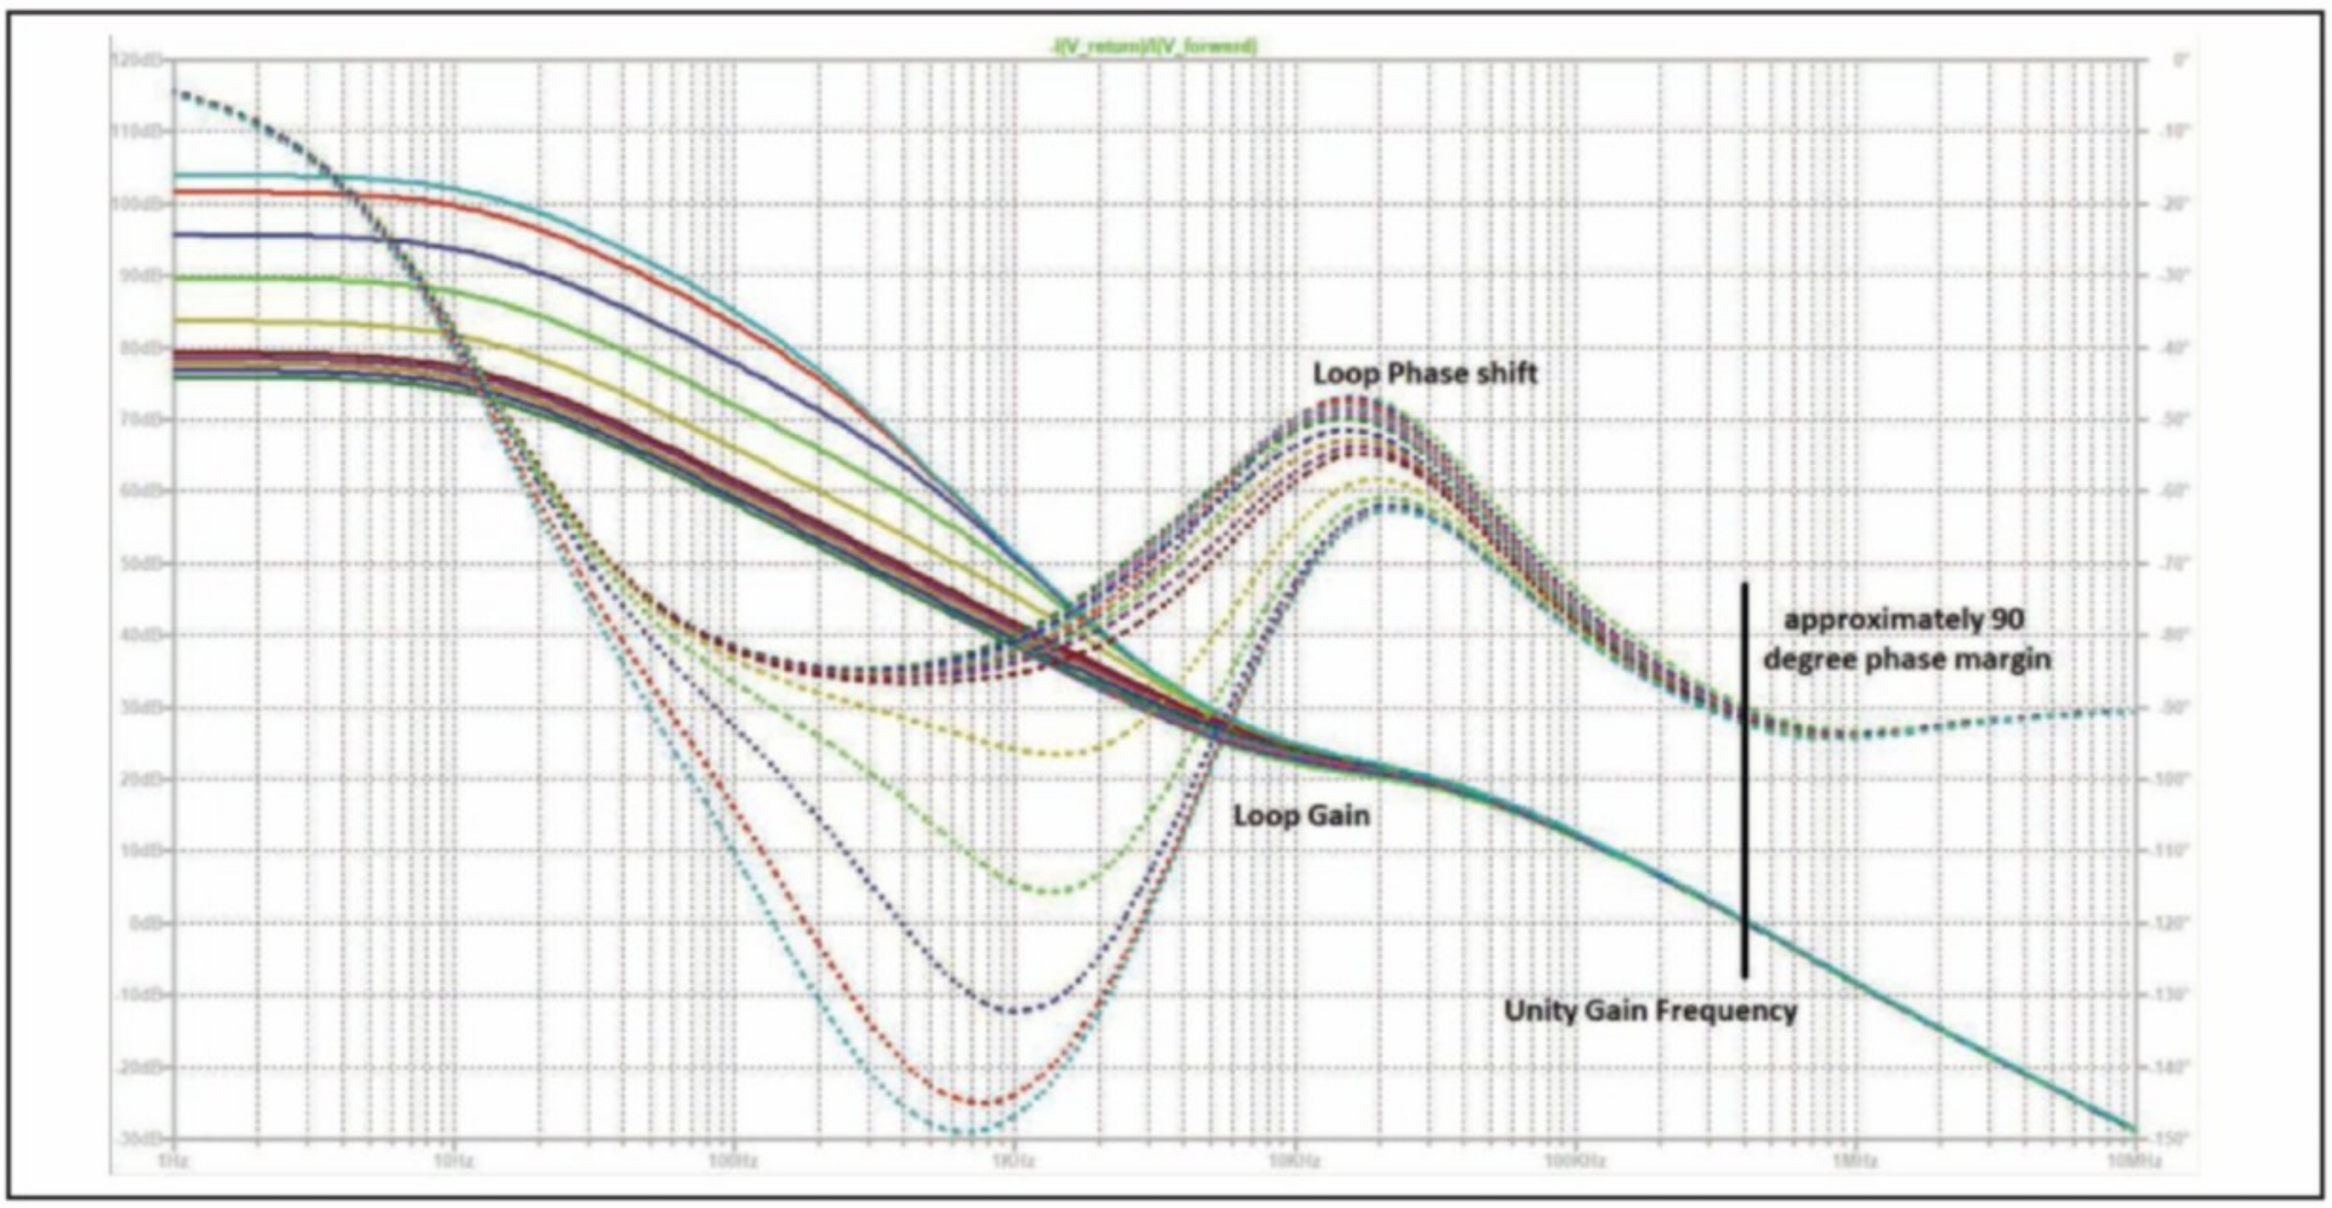

Figure 5b: The results of a simulation of the Reggie Phil 2.0 regulator shown in Figure 5a

We see that standard 1% values of 100 Ω for R_{18} and R_{20} and 1050 Ω for R_{17} and R_{19} work perfectly for 6.3 V and 100 and 2100 are perfect for 12.6 V. But we must consider offsets. Suppose we used 1% resistors. Applying the worst combinations of offsets and resistor tolerances to equation [2], we find that $6.068 < V_H < 6.538$. So at the absolute worst, V_H is within 3.8% of 6.3 V. This is adequate if the limit for heater voltage variation is 5%. But 0.1% resistors would be even better. In this case, $6.178 < V_H < 6.422$, better than 2.0%. Digi-Key has 100 Ω available in 0.1% tolerances. You can buy 100 1050 Ω resistors for a small cost and select the best at test, if you are so motivated.

Sampling current is more challenging. If we wish to limit dissipation to no more than 2 W at 8 A, then the current sense resistor R_{21} can be no more than about 30 mΩ. The cost of anything better than 1% tolerance for such low resistances is unreasonable. But we can buy a 1% resistor under 30 mΩ for less than $1 USD. We should tend toward the larger values, since the op-amp inputs will see voltages no higher than 240 mV, and with 4 mV offsets, the error will exceed 1%. I selected a 20 mΩ 3 W resistor. Again, we can write an equation from inspection and solve it, this time for heater current I_H:

$$I_H \times R_{21} \times \left[\frac{R_{16}}{R_{15}+R_{16}} \pm V_{OS} \right] \times \left[\frac{R_{14}+R_{13}}{R_{13}} \right] = 600 \pm 6mV \quad [4]$$

$$I_H = \left[(600 \pm 6mV) \times \frac{R_{13}}{R_{14}+R_{13}} \pm V_{OS} \right] \times \frac{\left[\frac{R_{15}+R_{16}}{R_{16}} \right]}{R_{21}} \quad [5]$$

Ignoring offsets and tolerances, we can require that R13 = R15 and R14 = R16 and conclude that:

$$\frac{V_{R21}}{600mV} = \frac{(R_{21} \times I_H)}{600mV} = \frac{1}{3.75} = \frac{R_{13}}{R_{16}} \quad [6]$$

where V_{R21} is the voltage across R_{21}.

According to my calculations, 750 Ω for R_{14} and R_{16} and 200 Ω for R_{13} and R_{15} work perfectly in this case. But considering tolerances with 1% resistors, we find that $7.436 < I_H < 8.585$ A; the error could be as high as 7.3%. This requires mitigation. If you could find 0.1% resistors (Digi-Key presently has them) it would mean that $7.575 < I_H < 8.434$ amps; the error could then be 5.4%. In practice, it is very unlikely that all offsets and tolerances would line up to give you an error greater than 5%. Still, some form of mitigation might have value, and it is available from R_{11} and P_1. P_1 can be adjusted to compensate for ±4 mV of op-amp offset. With that offset nulled, even with 1% resistors, the worst error barely exceeds 4%, and P_1 can even over-compensate in almost all cases to further reduce the error. But how to calculate the values of P_1 and R_{11}?

The design-center voltages at U_{1a}'s inputs are

$$V_{1a_in} = \frac{R_{21} \times I_H \times R_{16}}{R_{15}+R_{16}}$$

With P_1 turned to an extreme position, the current through R_{11} must cause a 4 mV change at one of the inputs. And so:

$$\frac{R_{15} \times R_{16}}{R_{15}+R_{16}} \times \frac{12V - V_{1a_in}}{R_{11}} = 4mV \quad [7]$$

$$R_{11} = \frac{R_{15} \times R_{16}}{R_{15}+R_{16}} \times \frac{12V - R_{21} \times I_H \times \frac{R_{16}}{R_{15}+R_{16}}}{4mV} = 469k\Omega \quad [8]$$

About the Author

Christopher Paul was born in the US the year before the Nobel prize was awarded for the invention of the transistor, the development of which he claims no credit for. His love of music from an early age led to some puzzling experiments with musical instrument amplifiers, and from there to an interest in electronics. Chris finds it interesting to delve into relatively simple circuits and to investigate and write about aspects of their performance that are little known or appreciated. He enjoys deriving equations that can be used to evaluate and design them, although he admits to being in therapy for this. Chris has worked for companies whose products are in the fields of communications, electronic article surveillance, and enterprise hand-held computers. He holds two M.Sc. degrees in electrical engineering from Brooklyn Polytechnic (now part of New York) University.

Resources

S. Franco, "Loop Gain Measurements," EDN Network, September 2014, https://www.edn.com/electronics-blogs/analog-bytes/4434609/Loop-gain-measurements

"Industry-Standard Dual Operational Amplifiers," Texas Instruments LM258A datasheet, June 1976 (Revised October 2019)http://www.ti.com/lit/ds/symlink/lm258a.pdf

Linear Technology LT1013CP datasheet, https://www.analog.com/media/en/technical-documentation/data-sheets/LT1013-LT1014.pdf

Meanwell LRS-75-12, 75 W Single Output Switching Power Supply https://www.meanwell-web.com/content/files/pdfs/productPdfs/MW/LRS-75/LRS-75-spec.pdf

R. D. Middlebrook, "Measurement of Loop Gain in Feedback Systems," International Journal of Electronics, Vol. 38, no. 4, pp. 485-512, 1975.

Riedon, MT 20 mΩ, 3 W resistor, MT-2B series, https://riedon.com/media/pdf/MT.pdf

"TPS54821 4.5 V to 17 V Input, 8 A Synchronous Step Down Converter With Hiccup," Texas Instruments datasheet, October 2011 (Revised February 2016) http://www.ti.com/lit/ds/symlink/tps54821.pdf

"TPS54821EVM-049, 8-A, SWIFT Regulator Evaluation Module," Texas Instruments Users' Guide, October 2011, http://www.ti.com/lit/ug/slvu479/slvu479.pdf

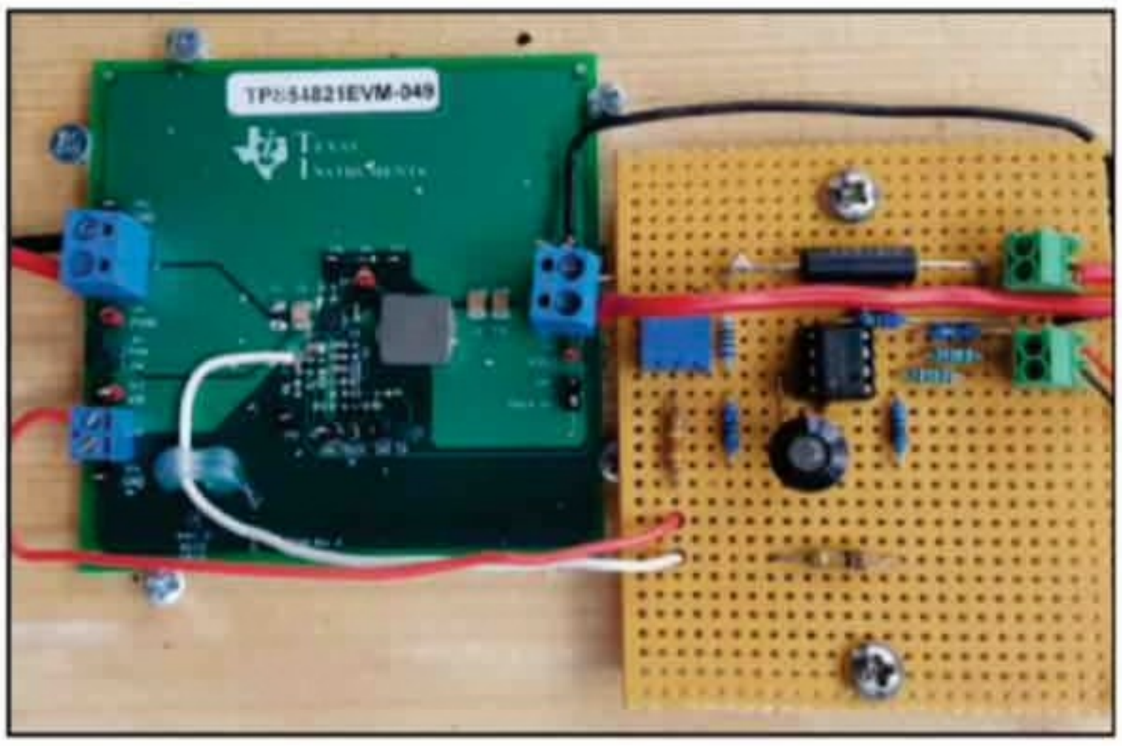

Photo 1: The breadboard and the modified Texas Instruments TPS54821 EVM are shown with interconnections. (The Meanwell 12 V supply is not shown.)

A smaller value of 422 kΩ for R_{11} was selected to accommodate 12 V supply tolerance. It's important that the R_{11} – P_1 network have a negligible effect on the impedances seen by U_{1a}'s inputs. We can evaluate this effect by approximating 12 V - V_{1a_in} as 12 V. From the equation [7], R_{11} is approximately 12 V/4 mV = 3000 times R_{15} || R_{16} and R_{13} || R_{14}. That is, these paralleled resistors are perturbed a mere 0.033%. Since P_1's end terminals are connected to points of virtually equal voltage (U_{1a} input terminals), it has no effect on circuit operation other than to conduct the current travelling through R_{11}. Choose values in the 50 kΩ to 200 kΩ range for P_1. Keep the values of R_{13}, R_{15}, R_{18}, and R_{20} between 100 Ω and 1000 Ω to avoid added input offset voltages due to input offset currents.

Alternate Design Strategies

Your design might want different values of V_H, I_H, and R_{21}. These could require different resistor ratios from those obtained from earlier evaluations of equations [3] and [6]. Suppose a ratio R were required. Calculate $k = 96 \cdot \log_{10}(R)$. Round the absolute value of k to the nearest integer. Note that there are 96 different mantissas of E96 series resistors. Starting with a value of 100 Ω, for each of next 95 higher values R_D, calculate the fractions R_{k-1}/R_D, R_k/R_D, and R_{k+1}/R_D, where the numerators are k-1, k and k+1 resistor values above R_D in the E96 table. Now divide each of these 3 × 96 values by R. The result(s) which are closest to unity give you the resistor pair(s) most closely matching R.

You might want to avoid using P_1 and R_{11}. One choice would be to consider V_H and calculate an expected aggregate filament resistance H from the filament specs. Evaluate the following considering tolerances: equation [2] gives the minimum value V_H, V_{H_min}. I_H should be chosen so that the maximum value I_{H_max} from equation [5] is such that $I_{H_max} \times H < V_{H_min}$. An advantage of this approach is that the red "current-controlling" LED will always extinguish and the green light as voltage control takes over when the filaments are (nearly) fully warmed up. A minor disadvantage is that the heaters will be subjected to a

hypex electronics
www.hypex.nl

NCAS500MP

400 + 100W in 8 ohm

Compatible with our MP-DSP
to create the perfect setup
for your active speaker

New family Ncore modules

NCAS500MP is the first to be released

Future members:
NCAS250MP: 200 + 50W
NCAS1000MP: 800 + 200W

For further information:
website: www.hypex.nl
mail: sales@hypex.nl

Ncore
audio power conversion

© Hypex Electronics B.V. All Rights Reserved

little more than their maximum rated current during warm-up. A second choice would be to substitute a $4.62 LT1013CP (maximum input offset voltage of 0.57 mV rather than 4 mV!) for the $0.44 LM258A.

These approaches are applicable if you intend to power a specific assembly of heaters of known resistance. But this is not the case for a lab supply. Here I recommend the original Reggie Phil, which used the TPS54821's slow start voltage feature: Add a 330 µF capacitor across the EVM's Track IN jumper and stuff a 0 Ω resistor in the R9 position. In this case, you don't need the breadboard; the only EVM modification is to set its R_7 equal to:

$$R_8 \times \left[\frac{V_H}{600mV} - 1\right]$$

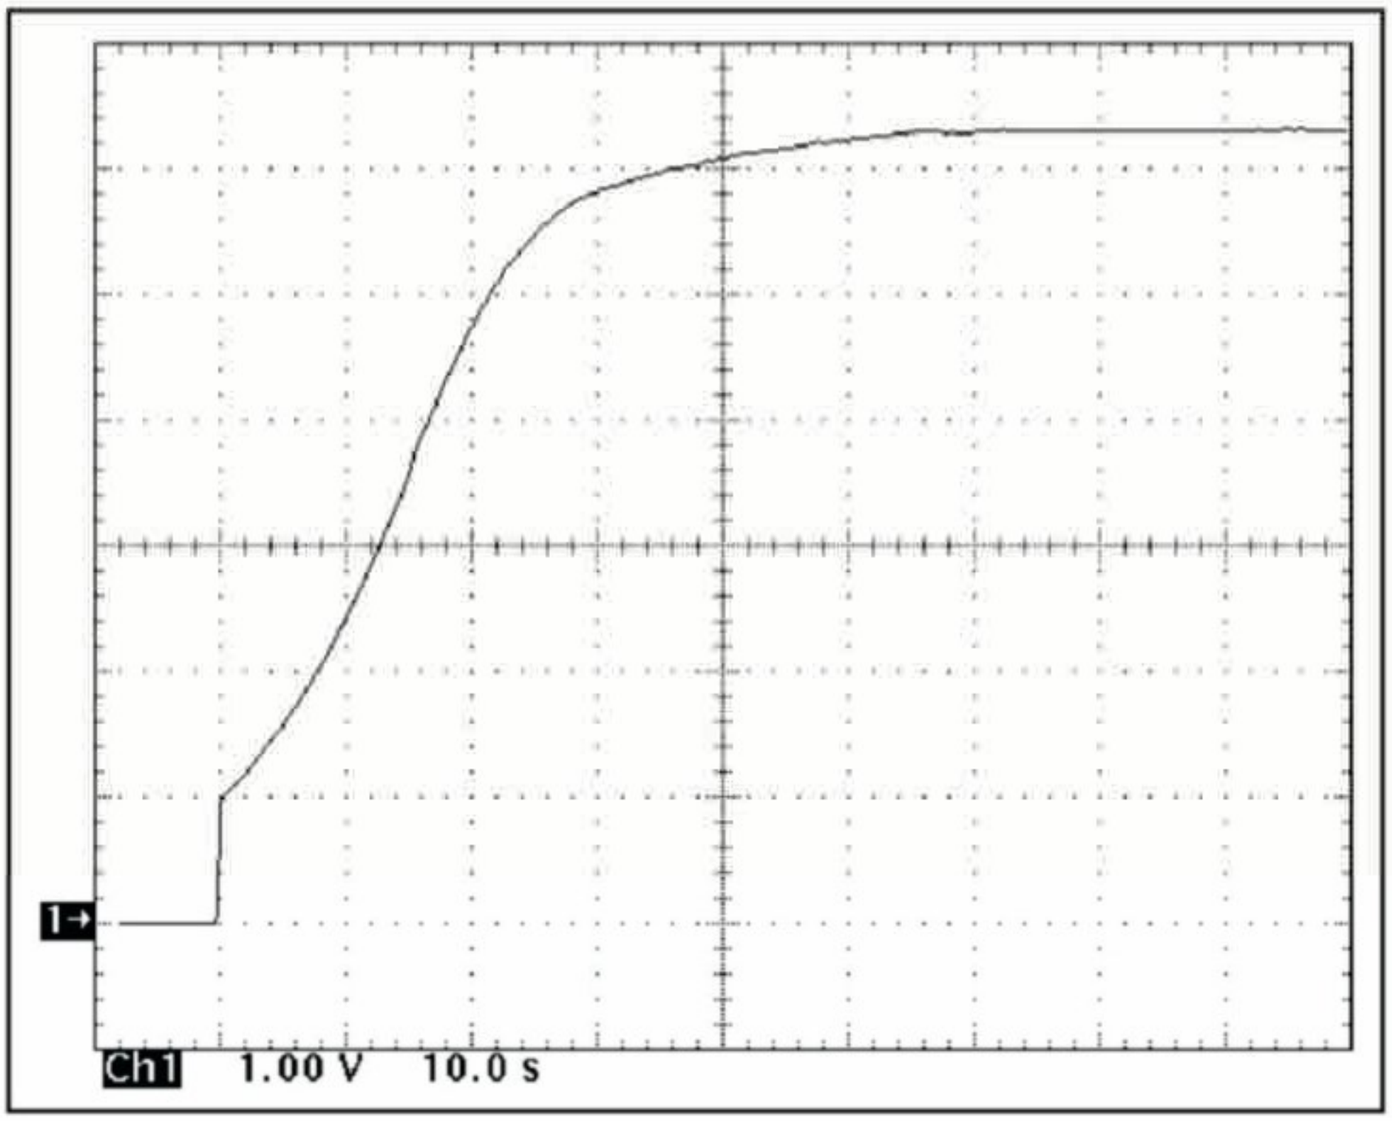

Figure 6: Heater voltage over time from cold turn-on of the parallel combination of two 6CD6GAs, a 6AU4GTA, and a 6BQ6GTB driven by the Reggie Phil 2.0 regulator. Note the jump in voltage at turn-on time due to the immediate application of an 8 A current to these 6.3 V-rated filaments.

Bill of Materials

(All parts available through Digi-Key)

Quantity	Description
1	AC powered, 12 VDC, 75 W output power supply, Meanwell LRS-75-12
1	DC-DC converter, TPS54821EVM-049
1	AC line cord, male plug to wire leads, three connector
1	LM258A op-amp, DIP package or LT013CP
2	Terminal blocks, 2 pin, (any through hole mount part able to accept 16 AWG wire)
1	Bicolor red/green LED, common cathode (any through hole mount part)
As needed	Hook up wire, twisted pair, 16 AWG
As needed	Hook up wire, twisted pair, 22 AWG
1	Current sense resistor, 1%, R21, (see text)
1	Resistor, R12, 10 kΩ 5%, 0.125 W or 0.25 W
1	Capacitor, C1, 0.1 µF X7R, 25 V
1	Potentiometer, multi-turn suggested, P1, 50 kΩ to 200 kΩ, (optional, see text)
1	Resistor, R11, 1%, 0.125 W or 0.25 W (optional, see text)
1 each	Resistors, R13 – R20, 1%, 0.125 W or 0.25 W

Construction and Layout

Consider the voltage drops in the wiring in the high current paths. The arrows in the schematic are meant to highlight that R_{13} and R_{15} must be connected to the leads of R_{21} 0.665" (16.9 mm) from the part's center to maintain the 1% tolerance (see Riedon MT2B-0R02F1 datasheet.) The leads are long enough to connect one end directly to GND on the EVM's J4 terminal block. The other lead should be connected to a terminal block on the breadboard, which will accept a high current connection to the heaters. The other end of the heater connection should be to the VOUT terminal of the J4 EVM terminal block. Another pair of schematic arrows indicate that R_{17} and R_{19} must each be connected to the terminals of a second breadboard terminal block whose other connections are to a pair of low current-carrying wires from a point representing the average of all filament voltages. Finally, connect C_1 as close to U_1's power pins as possible. Other than these considerations, layout and construction are non-critical.

Photo 1 shows the completed design, which I've mounted on a board that I use as a test bed for various tube circuits. R_{21} is the large black component on the breadboard. The sensing connections are at the ends of the visible portions of that component's leads.

Calibration

Calibration is appropriate only if P_1 and R_{11} are employed. Connect a voltmeter across the points on R_{21} sensed by R_{13} and R_{15}. (Alternatively, if you want better than the 1% accuracy promised by sensing across R_{21}, you could insert an ammeter in series with the heaters before you energize them.) As the heaters are warming up (red LED on), adjust P_1 for the designed-for heater current.

Performance

During turn-on, I adjusted the current to 160.6 mV across R_{21}, which would be between 8.011 A (+1.01%) and 7.95 A (-0.6%). When the green LED illuminated, indicating voltage control, I measured 6.236V, an error of -1.02%. **Figure 6** shows a 90% of final voltage turn-on time of about 25 seconds, a notable improvement over Reggie Phil 1.0's 110 seconds.

Enjoy your Reggie-Phil 2.0!

Engineering Consulting for the Audio Community

Celebrating 36 YEARS 1983-2019

MENLO SCIENTIFIC

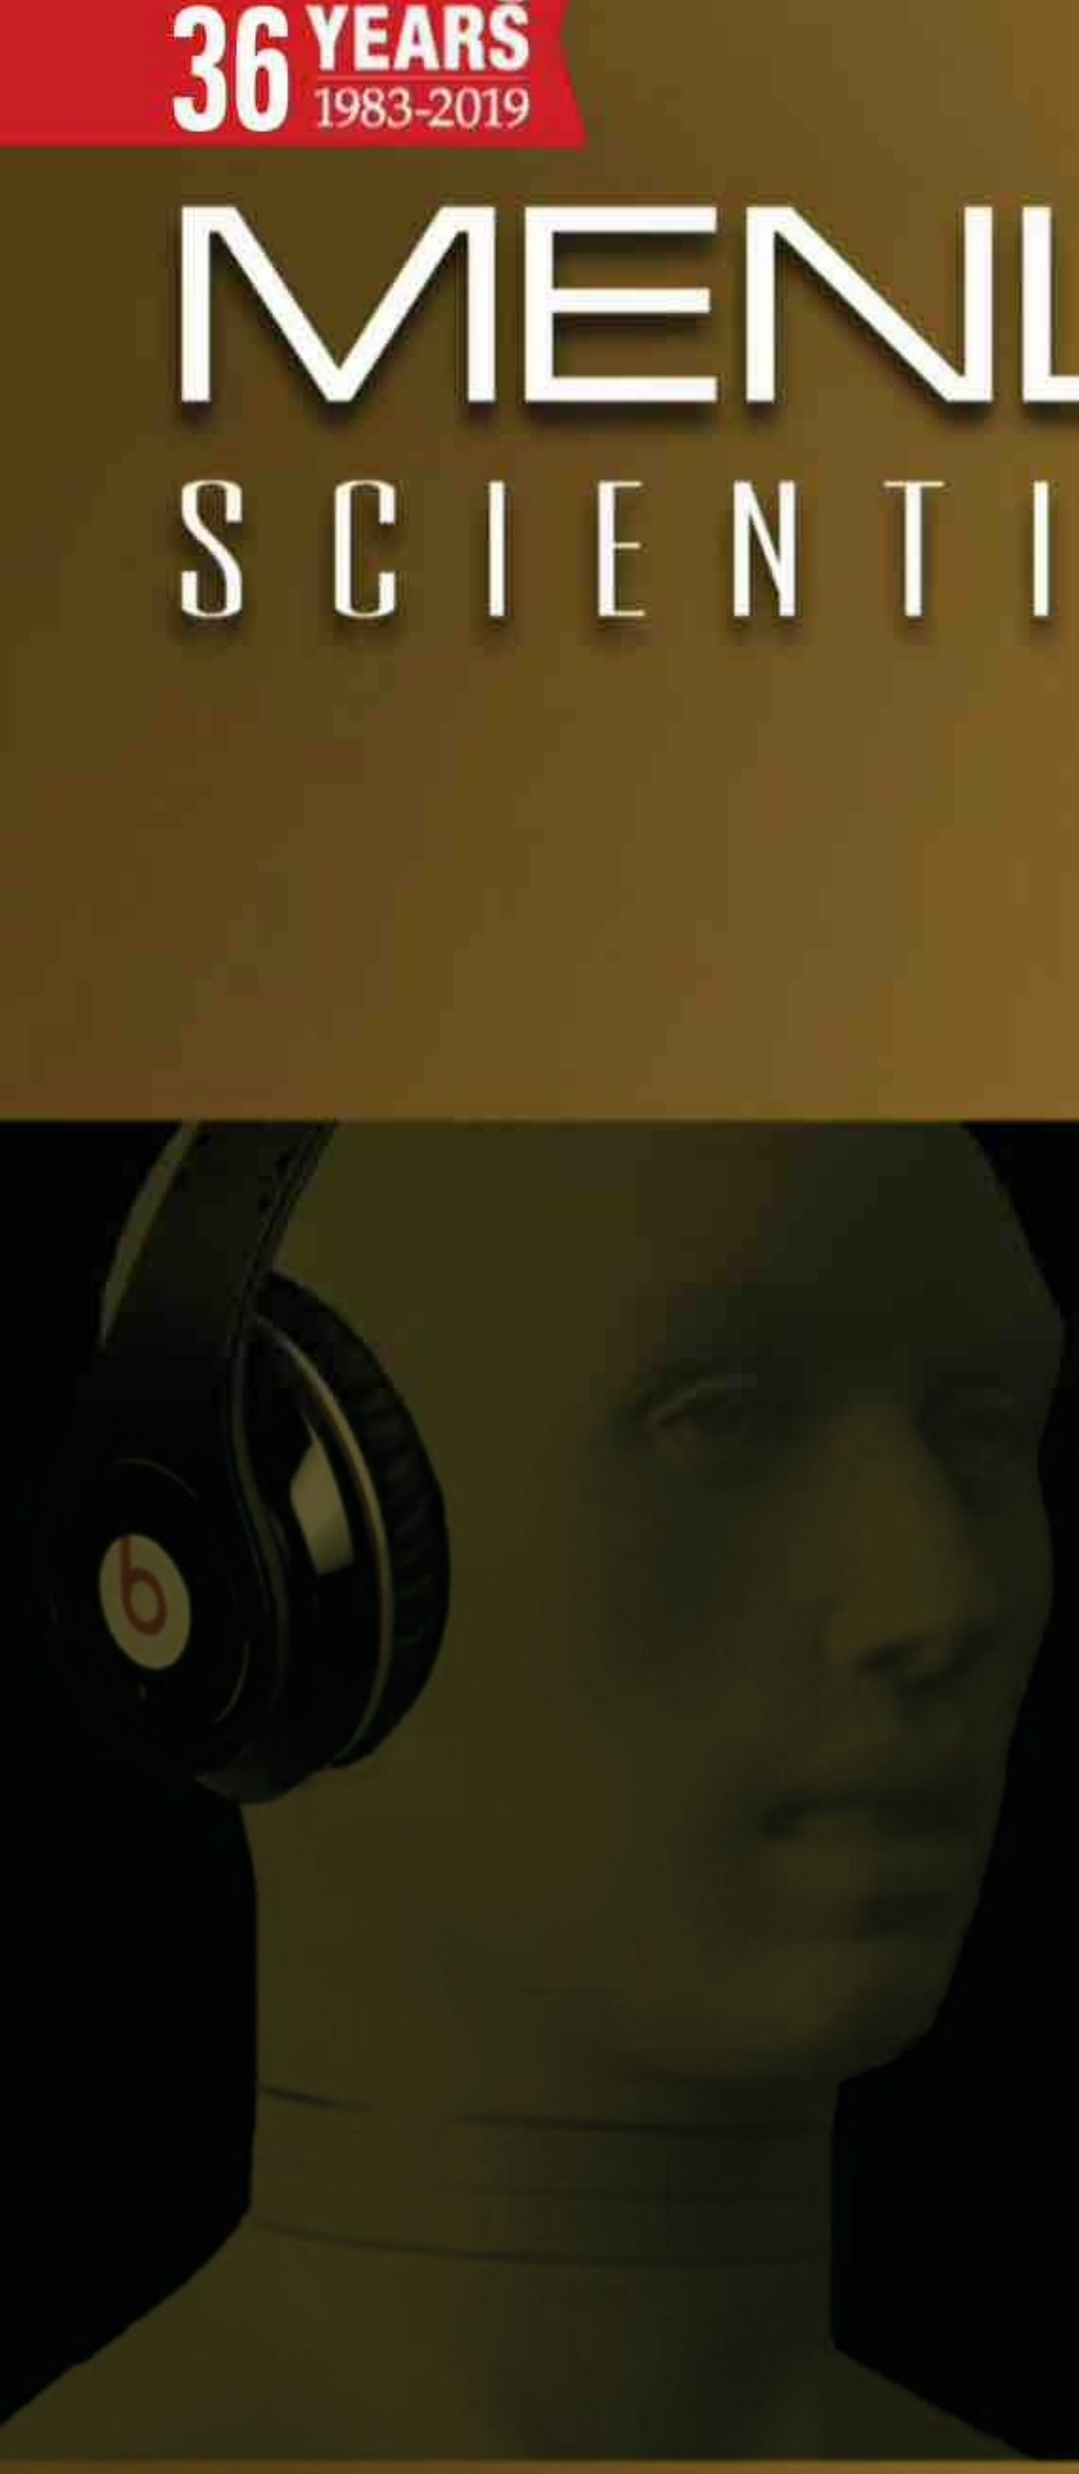

- Liaison Between Materials, Suppliers, & Audio Manufacturers
- Technical Support for Headset, Mic, Speaker & Manufacturers
- Product Development
- Test & Measurement Instrumentation
- Sourcing Guidance

MENLO SCIENTIFIC

USA • tel 510.758.9014 • Michael Klasco • info@menloscientific.com • www.menloscientific.com

Build an 18 W Single-Ended Valve Amplifier with Trans Technology

This article shows an easy to build 18 W single-ended valve amplifier using Trans-technology with pentodes, current sources, and toroidal transformers (see **Photo 1**). The amplifier was created by the author to support his TubeSociety students in the Netherlands.

Photo 1: This is the top view of the Vanderveen Trans-SE18 amplifier.

By
Menno van der Veen
(The Netherlands)

Years ago I rediscovered a special technology, called "Transconductance," first published around 1940. I wrote a book[1] and *audioXpress* articles[2],[3] about it, where I explained the technology, which I termed "Trans." Please refer to the publications listed in Resources to learn what Trans implies in more detail.

For now: Trans is local feedback only around the output power tube. Preferably this tube is a pentode. Trans converts the tube into a new Trans-tube with little harmonic distortion and low plate resistance, which minimizes transient distortions inside the output transformer. Consequently the Trans-tube and OPT marriage distorts minimally. An extra bonus is the conservation of pentode output power, which is twice the power of a triode.

About 12 years ago, I started my TubeSociety academy where I try to teach my TubeSociety students to become the best designers and constructors of tube amplifiers. For the TubeSociety students' 2018 project, I shared the design steps of the Trans-SE18 (see **Photo 2**). The design project can also be found on my website (www.mennovanderveen.nl) where I shared information with my students about their TS2018-project.[4]

To apply Trans, a new voltage controlled current source (VCCS) is introduced, called the "Menno-cell." An extra Menno-cell is used to linearize this VCCS. Inside the Menno-cell there are some field-effect transducers (FETs), which perform the current source function and linearization. For those who dislike semiconductors inside tube amps, please study my all-tube kit, Trans-SE10.[5] In this article, I continue with the Menno-cell solution.

The Audio Section

Figure 1 shows the schematic of the SE18 design. On the left, you will find the inputs, their selector, and a volume control. In the middle, you will find two Menno-cells. The first cell-1 functions as a preamplifier and linearizes cell-2. The combination of cell-1 and cell-2 makes this second cell an ideal VCCS with negligible distortion. The current output of cell-2 drives the power tube. The resistor R** (220 kΩ) creates the essential local Trans-feedback.

The diode D1 prevents the voltage between control-grid and cathode to become positive and protects the power tube at start-up.

The 22 turns trimpot P1 sets the quiescent current through the power tube at 115 mA. You can measure this with a DC voltmeter over R11. Its reading should be 115 mV. The trimming of P1 is a little sensitive; you might have to repeat it several times. Take notice of the "cw" indication in the schematics. Turning P1 to "cw" means more quiescent current.

The Power Supplies

Figure 2 shows the power supply section of the stereo amplifier. The special power transformer can be ordered from my website. The application of two power supplies—Vn = 120 V and B = 360 V—is essential. The power tube only uses the 360 V supply. A quiescent current of 115 mA flows per channel. For stereo, the total current is 230 mA.

The complete power tube section is lifted 120 V above ground with the Vn supply. Normally, this lifting is done with a cathode resistor plus a paralleled capacitor. However, to create 120 V with 115 mA current requires a resistor of about 1 kΩ with a heat dissipation of 14 W (28 W in stereo). That is a lot of heat! A simple Vn supply does this task much better with negligible heat production.

The two Menno-cells use about 5 mA supply-current per channel. So, 10 mA for the complete stereo amplifier. This small current is delivered by the Vn supply. I applied a very effective voltage doubler circuit[6], not often seen in the audio amplifier literature.

Photo 2: Menno van der Veen addresses his students during one of the bi-weekly meetings of his TubeSociety school in the Netherlands.

The circuits with Fp1 and Fp2 are not voltage stabilizers, but rather floating active buffers that react slowly on variations of load or mains. They have a residual ripple voltage below 1 mVrms and an output impedance of approximately 1.5 Ω. If you try to get such good specs with tubes, you will need a very complex tube circuit. Every time I apply my version of these super-C-source-followers, I am amazed by their excellent qualities. Don't forget to place Fp1 directly on the chassis for cooling, because it generates approximately 2.8 W of heat. Solder R23 as close as possible to the Fp1 gate to prevent oscillation. The same for R26 and the gate of Fp2.

The audio section also needs 200 V and 300 V supplies. These last two voltages are created with the Zener diodes Z1 and Z2. Because they power the input and VCCS sections, extra low ripple is essential. The Zeners also function as constant load supplies, which sounds superior in preamp sections.

An extra remark about grounding: The mains-ground is directly connected to the chassis for optimal

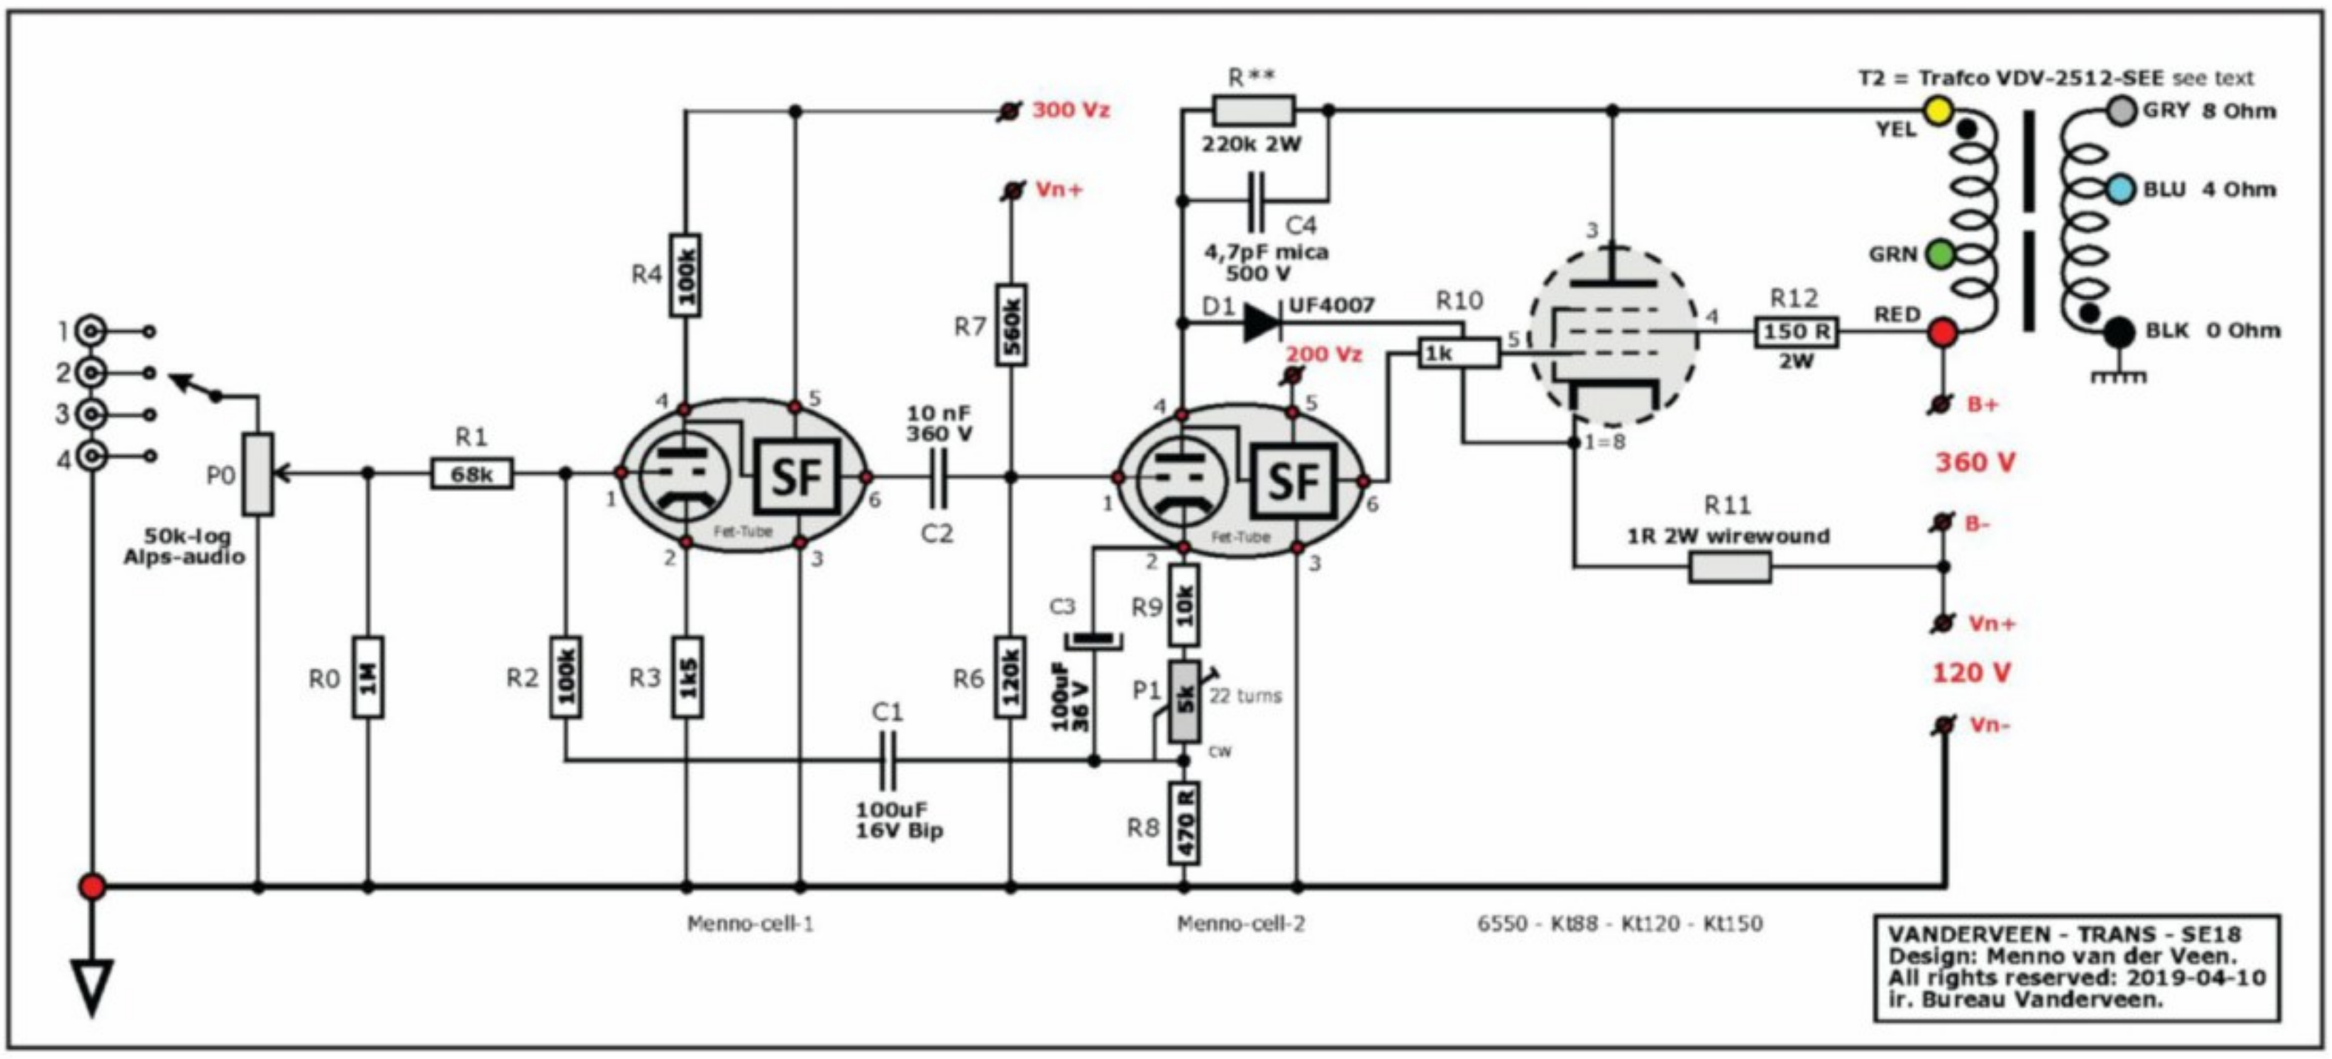

Figure 1: This is the schematic of the audio section of the Vanderveen-Trans-SE18.

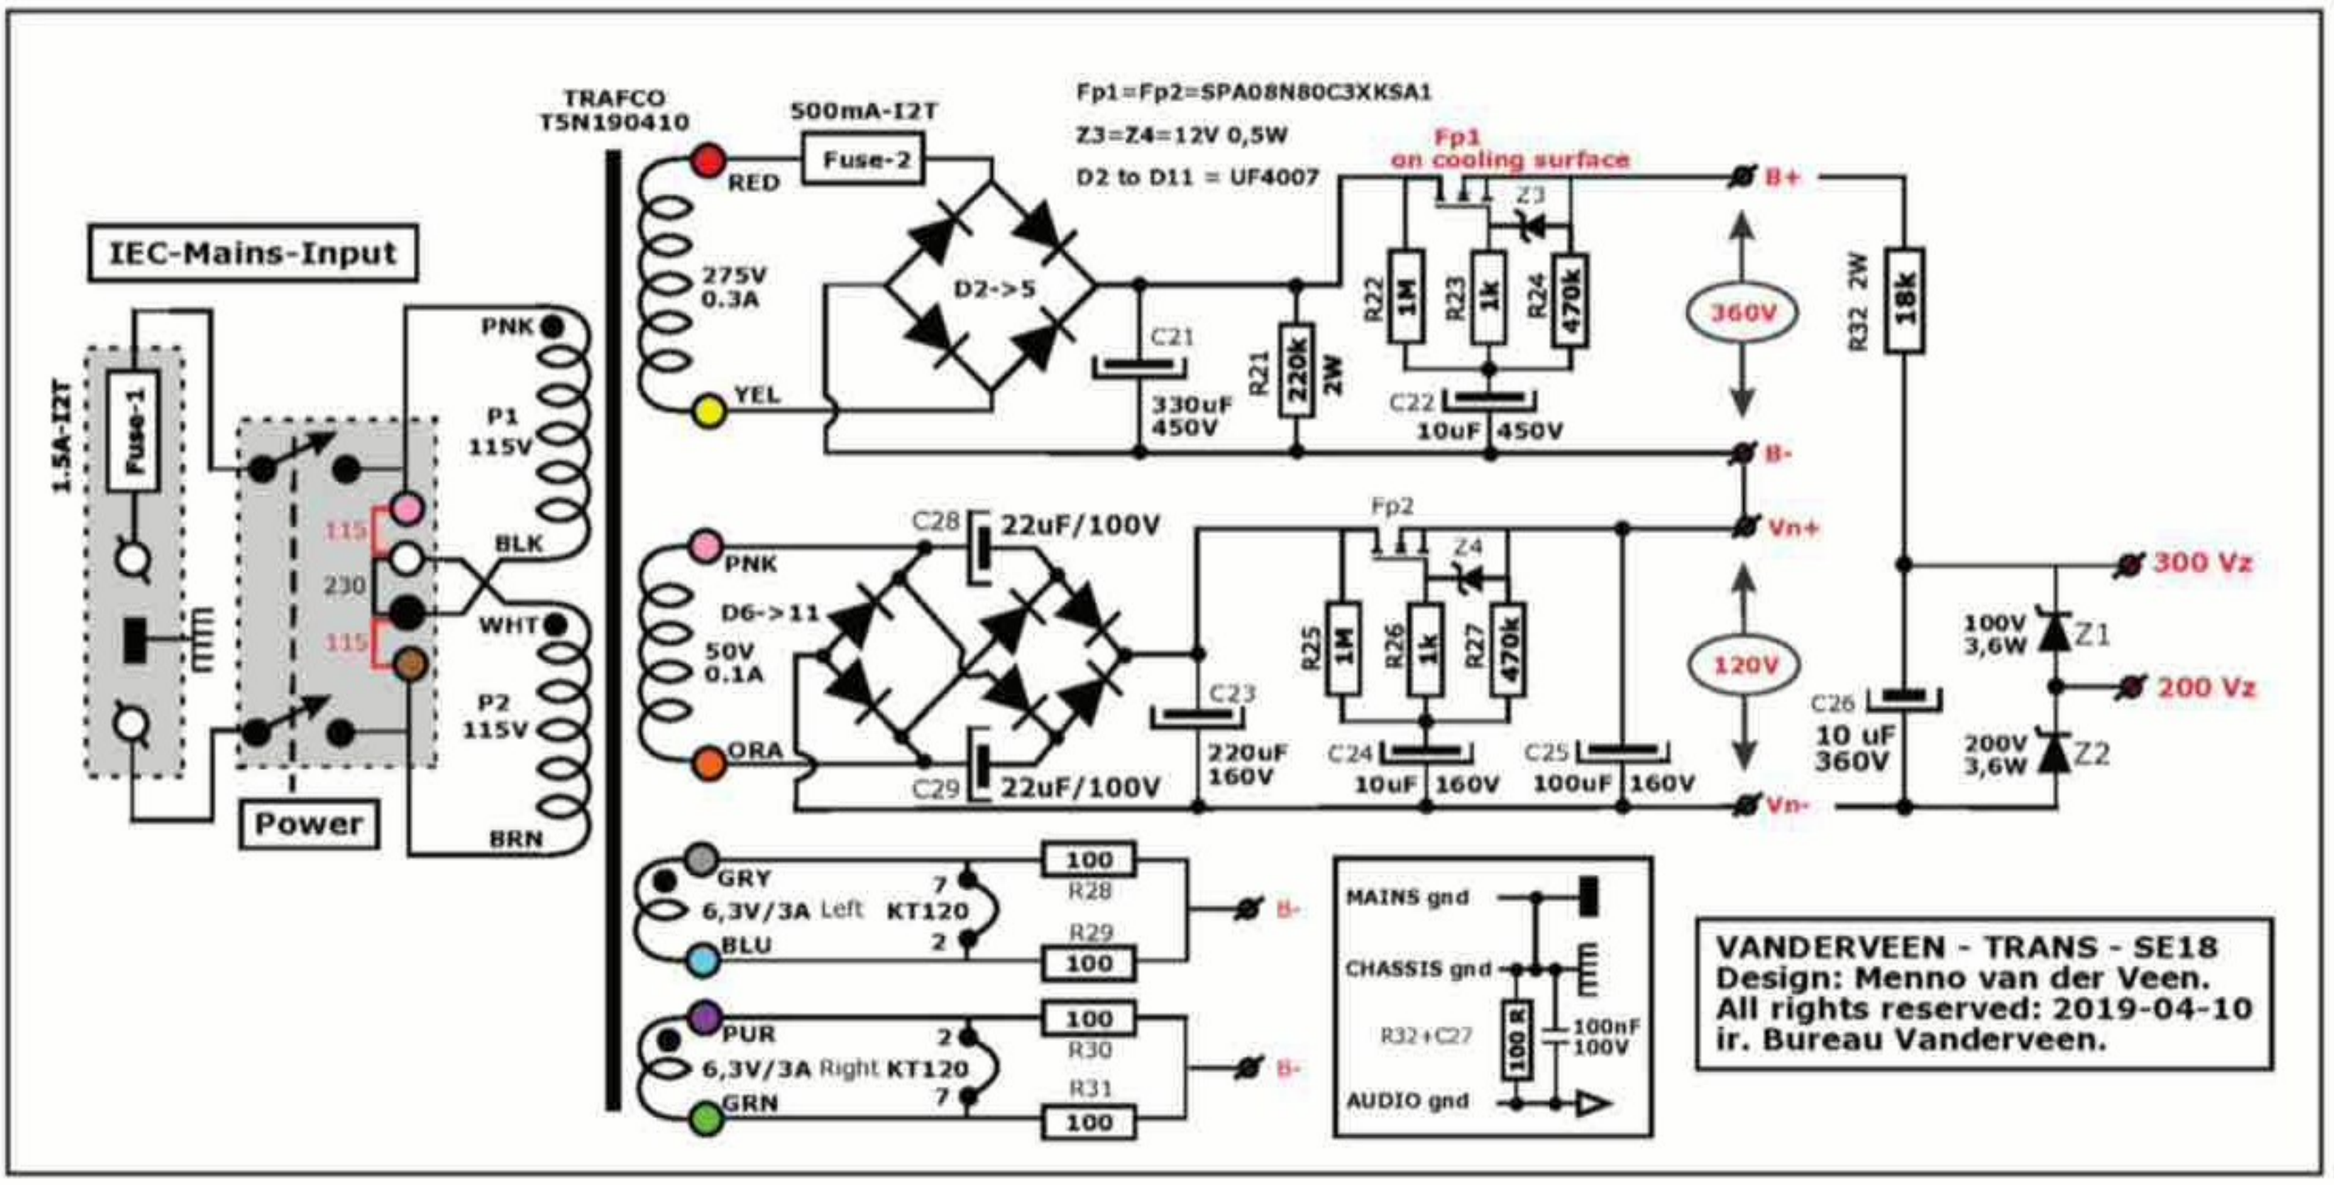

Figure 2: Here is the schematic for the stereo power supply section.

safety. The audio-ground goes to the chassis through R32//C27. This circuitry prevents grounding hum between mains-grounded equipment. Be aware that no connection is made between audio-ground and the chassis at any other spot inside the amplifier. There is only one place for R32//C27: closest to the input sockets!

The Menno-Cell and the Transformers

Figure 3 shows the schematic of the Menno-Cell. The circuit has three FETs (BSP135). The first FET F1 is an amplifier with very large output impedance (almost an ideal current source); F2 is a source follower; and F3 is a current source that linearizes F2. The resistors R1 and R2 at the gates prevent oscillation, while the Zener diodes protect at start-up. The resistor R3 sets the F3 current source at approximately 1.2 mA. My student Erwin Reins designed a clever PCB shown in **Figure 4**, which is available for free Gerbers download from my website.[4]

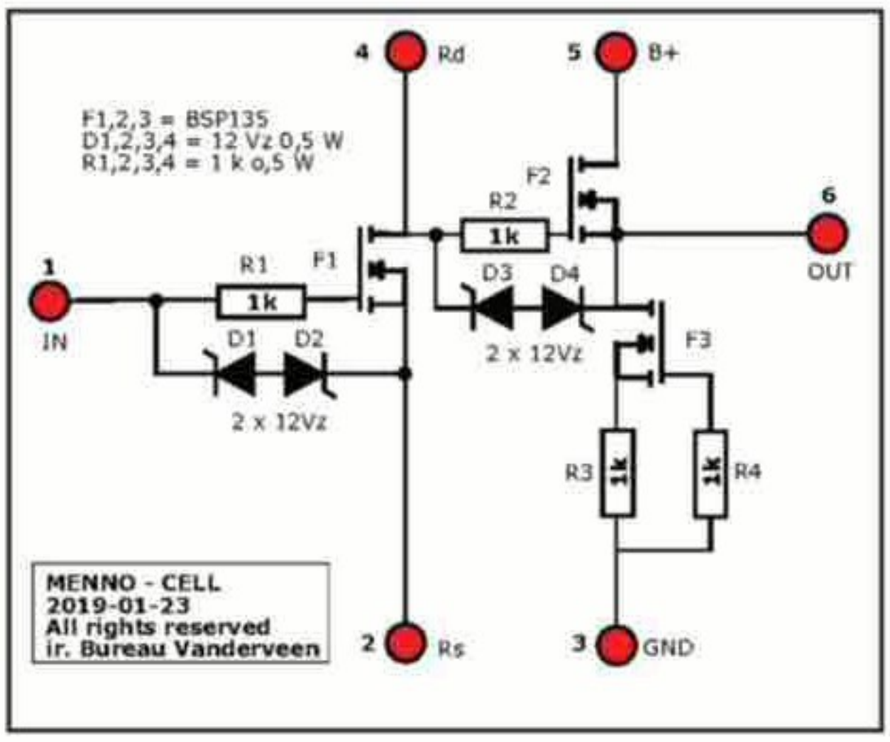

Figure 3: Schematics of the Menno-cell

There are two toroidal output transformers for this Trans-SE18 amp. The 2512-SEE has a maximum power of 12 W at 22 Hz. It easily can handle the power of the amplifier, but not at the lowest bass frequencies. I can measure such bass saturation, but in practice I don't hear it. The other OPT is the 3025-SE, which can handle 17 W at 22 Hz. See my website[7] for both specifications. The best of both is the 3025-SE, with biggest core (and price).

The Construction

If you look at Photo 1, you can see the power transformer is in the middle. The transformers are protected with black Mylar. I spray painted the mounting discs. The left front switch is power and the right one is free for any application. I gave it a function of 10 dB gain-reduction by switching 150 kΩ in series with R1. Such a function is handy when there is a phone call or for non-disturbing listening. The case was also designed by Reins and its dimensions are 35 cm × 29.5 cm × 8 cm.

The volume knob has a built-in two color LED, connected to a small circuit that measures the left and right cathode currents (measured over R11). When the current is below 60 mA per power tube, the LED glows red. This indicates that mains power is ON. When both quiescent currents are above 60 mA, the LED switches to yellow. This indicates that the amplifier is (almost) ready for listening. I leave the decision of this

quiescent current sensor circuit up to you. It is based on a TL072, operating as a comparator. Its supply voltages are taken from a 6.3 VAC filament winding, with rectifying (2 × 1N4007) and buffering (2 × 100 µF/16 V), delivering +8 VDC and -8 VDC to feed the op-amp and the bicolor LED. **Photo 3** (on the right side) shows the realization of this circuit. The inside view of the amp (Photo 3) shows the strip at the bottom, which is the power supply section. Enlarge, and you will find all the details.

In my test amplifier, I did not apply an input selector plus four inputs. One stereo input was enough for all the measurements and listening tests.

I placed the two Menno-cells close to the power tube socket. Please use the shortest connections between the power tube and these cells. If you don't do that, nasty oscillation can occur, resulting in unpredictable quiescent currents.

To measure possible oscillation: connect a 10 MHz (or better) oscilloscope to the output terminals. If you measure 50 mV high-frequency signals, then your amp is heavily oscillating. You can cure it by means of shortest and smartest layout of the interconnections. Enlarge Photo 3 to see how I did it.

For the First Start-Up

So, you've built your version of the Trans-SE18 and checked everything at least a hundred times. So, let's switch on the mains and start enjoying music or producing nasty smoke. Allow me to make some helpful remarks for your first time switch-on.

- Check that P1 is turned to minimum resistance (turned counter-cw)
- Measure the power supply voltages Vn and B without the power tube; see **Figure 5**.
- Measure the voltage at the tube socket point 5. It should be less than 70 VDC, see Figure 5.
- Plug in the power tube, switch-on mains while measuring the voltage over R11 (= 1 Ω). It should read almost zero
- Slowly turn P1 clockwise until the voltage over R11 rises and bring it to 60 mVDC
- Do the same procedure for the other channel
- Now continue to turn P1 clockwise to bring the voltage over R11 to 115 mVDC; do the same for the other channel
- If all of this does not function, you probably have oscillation. Check the layout of wires and change.
- If all functions, check the voltages indicated in Figure 5. Your readings should be close to the given data. The correct reading of 115 mV over R11 is essential.
- Now it is time to enjoy the music.

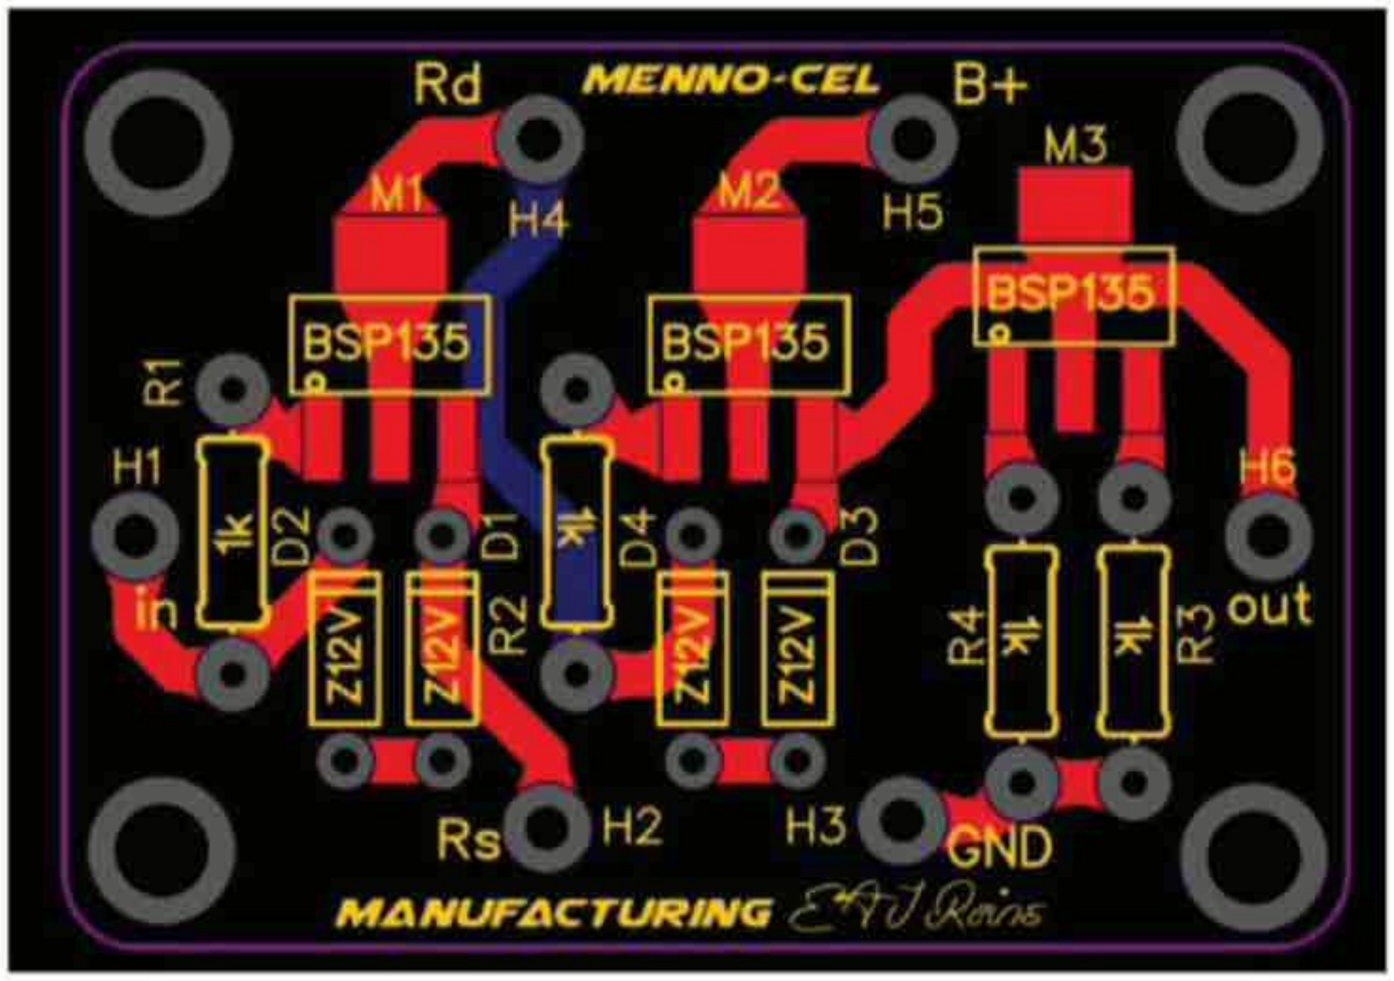

Figure 4: Erwin Reins designed a clever PCB, which is available for a free Gerbers download from my website.

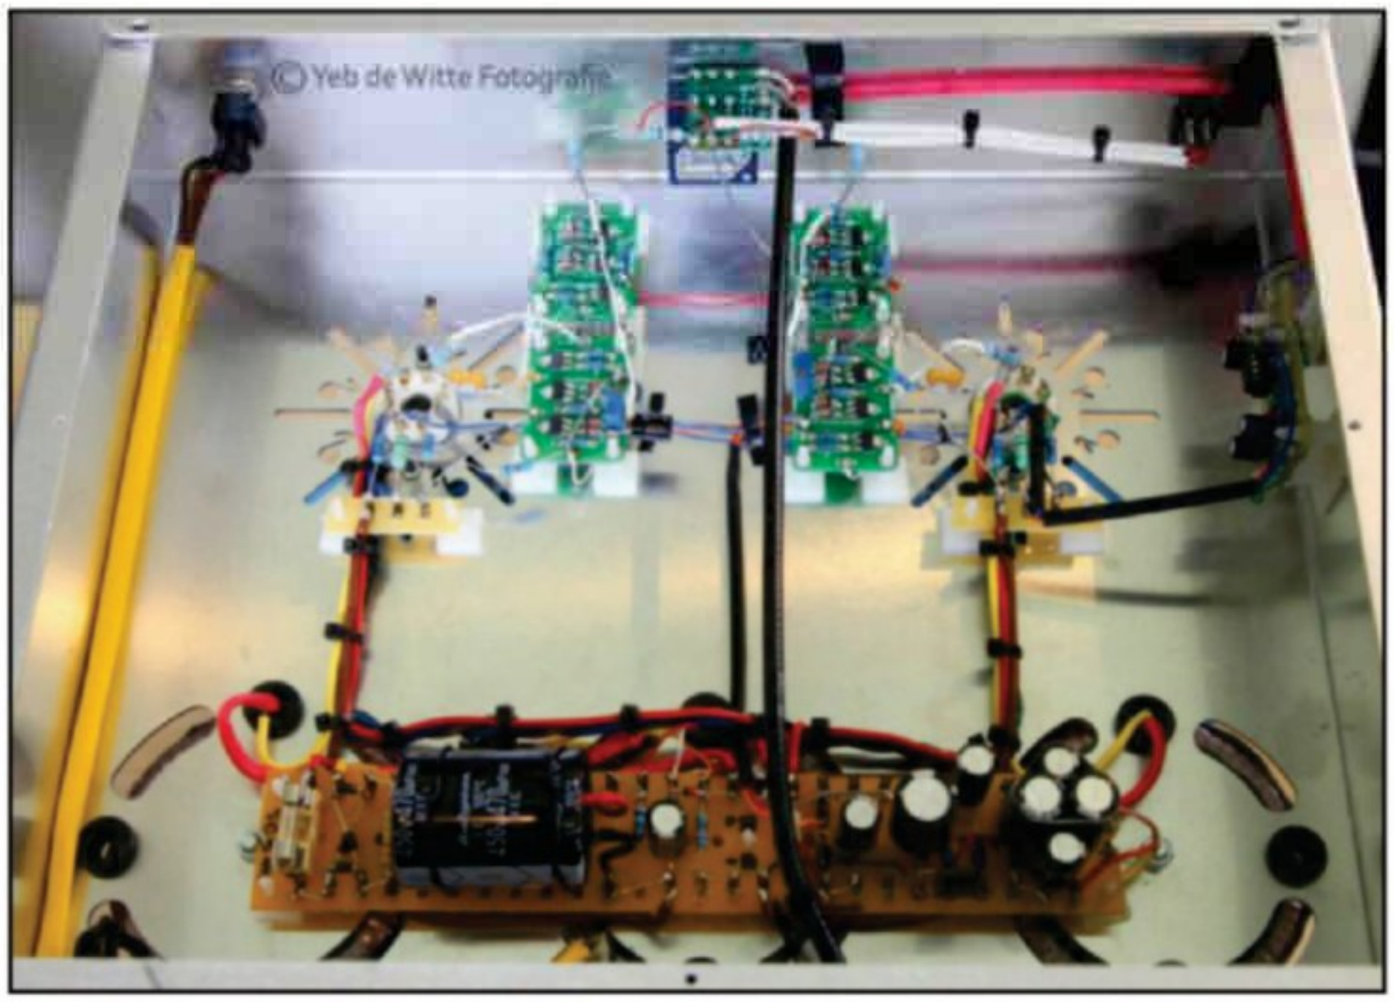

Photo 3: The interior view of the Vanderveen Trans-SE18 amplifier is shown here.

About the Author

Menno van der Veen studied engineering physics at the university. He taught physics at an upper secondary school and teacher trainings college. At the age of 40, he founded his engineering firm, which focues on tube amplifiers and toroidal output transformers. He wrote many articles and books about these subjects. About 12 years ago, he also founded his TubeSociety academy. Tube-amps and toroidal transformer research, consultancy, design, teaching, and playing the guitar are the main activities of his present work.

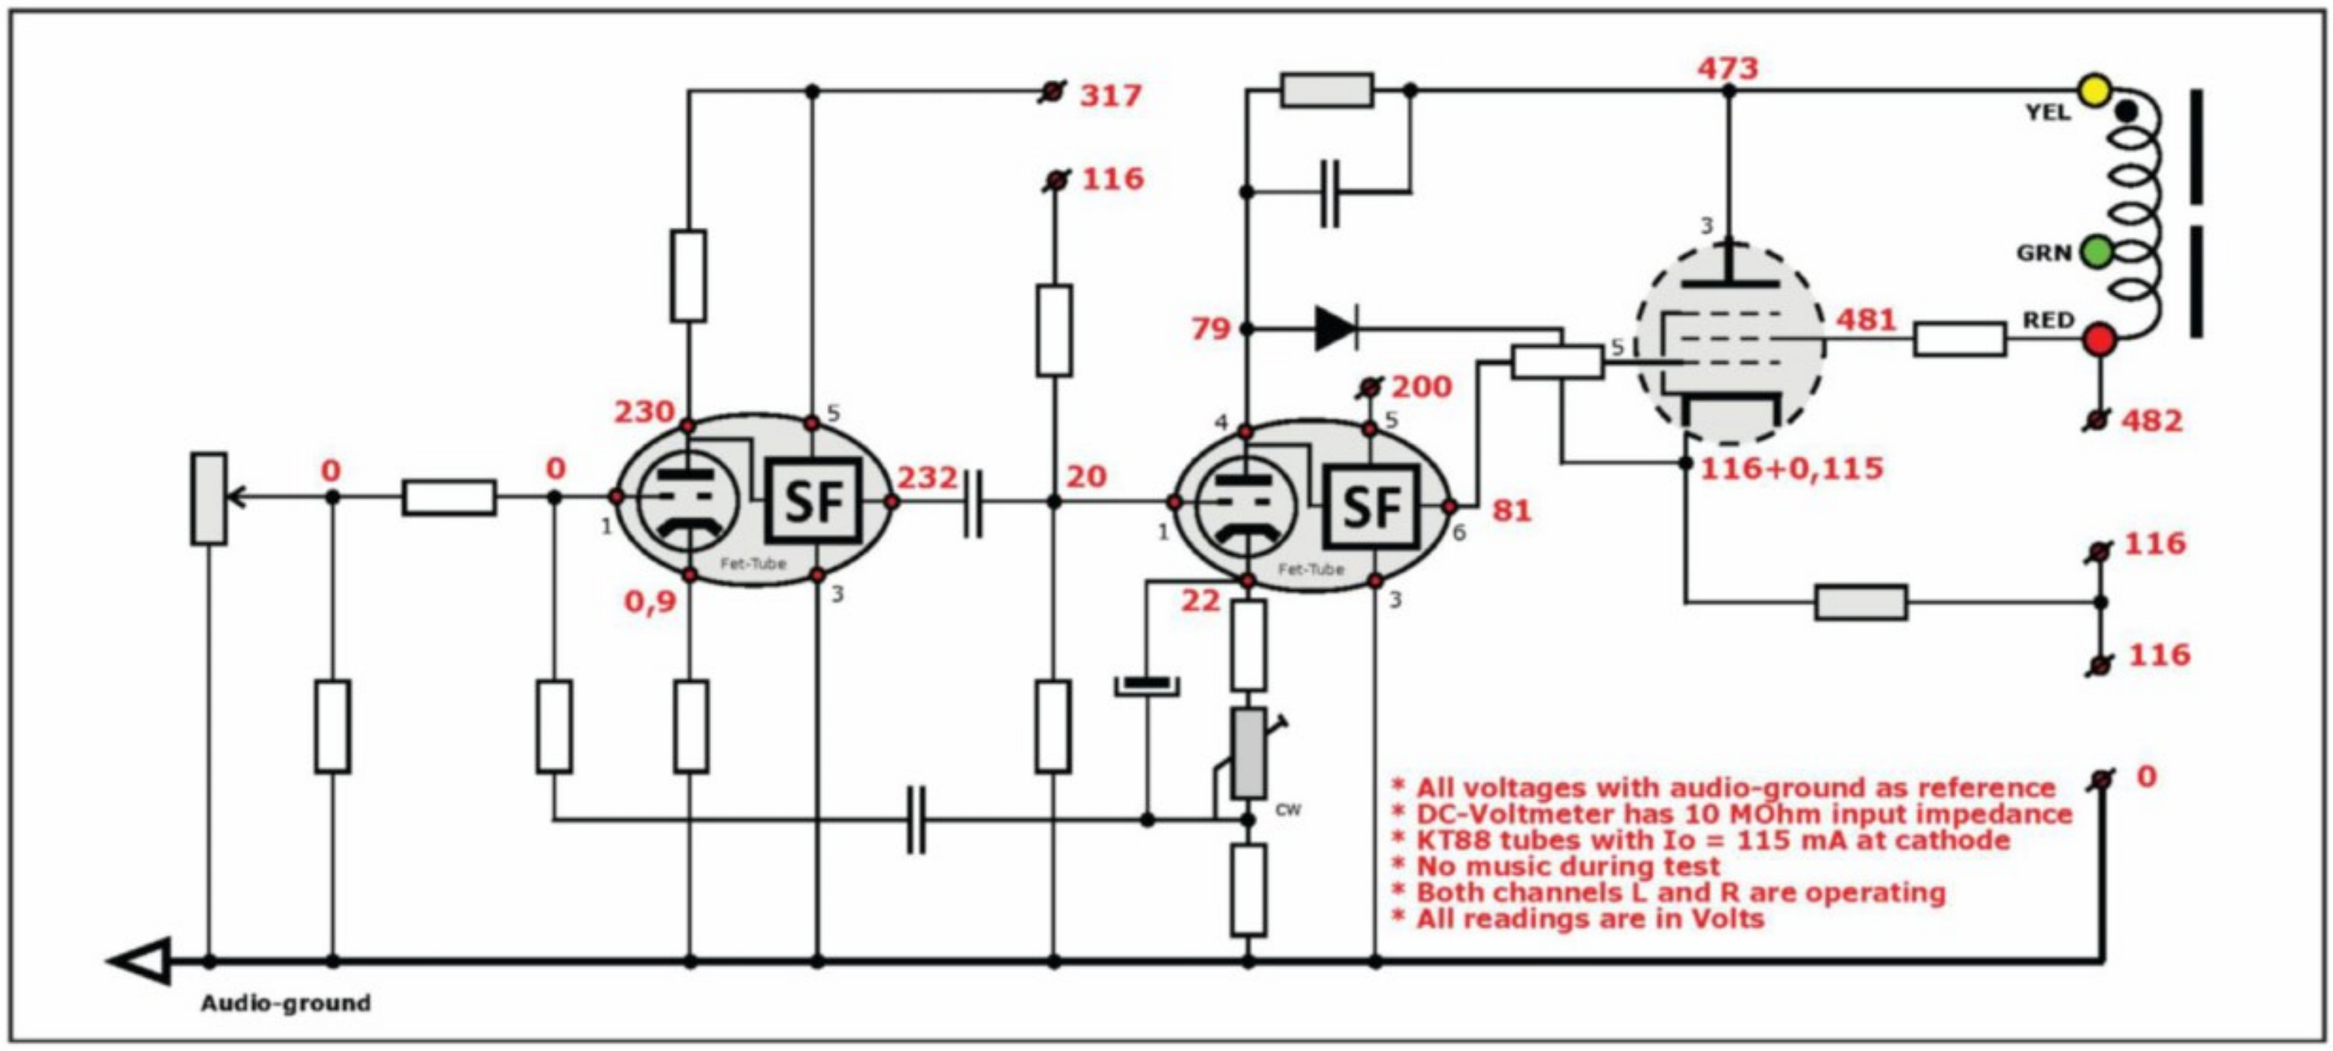

Figure 5: Your voltages should correspond to all the voltages shown here inside the amplifier.

I measured important DC voltages inside the amplifier, enabling you to check correct operation. See Figure 5. All measurements are with a digital voltmeter with audio-ground as reference.

After switching ON, the amplifier takes about 20 seconds to heat up. If you would listen to music during start-up, you clearly will notice that the amp moves towards its optimal operating point.

Measurements

Figure 6 shows the frequency characteristic from 5 Hz to 40 kHz. The -3 dB frequencies are at 7 Hz and 43 kHz, without any higher resonances.

Figure 7 shows the harmonic distortions. From 20 Hz to 1 kHz, you see a clear decline in distortion of all harmonics. However, the third (H3) and higher harmonics decline more than the second (H2) harmonic. See my book for more explanation about this typical Trans-effect, which mainly is caused by local Trans-feedback in combination with the primary inductance Lp of the OPT.

Above 1 kHz, H2 inclines again a little. The Trans-corrected power tube starts to interact with the interwinding capacitance Cip of the OPT. Around 1 kHz, we get the lowest total harmonic distortion (THD) of 0.04 %, where H2 is not dominant. Especially this last result removes the "curtain" before the soundstage, as explained and discussed in my previous *audioXpress* articles (see Resources).

The last measurement (**Figure 8**) shows the output impedance of the amplifier. At the highest frequencies a rise in output impedance is visible, caused by the leakage inductance Lsp of the output transformer. Below say 2 kHz the output impedance is 0.67 Ω, which is a typical result of the Trans principle.

The maximum output power of the amp is 18 W, measured at the onset of soft symmetrical clipping. The voltage amplification is 27 dBV/V = 22.4 x. With reference to 18 W in a 4 Ω load, the input sensitivity of the amplifier is 380 mV. The input impedance is 29 kΩ, equal to the 50 kΩ of the volume control in parallel with R1 = 68 kΩ and R0 = 1 MΩ.

In summary: It all looks sound and understandable. This amplifier follows the typical Trans-rules. What is special are the low harmonic distortion with H2 equal or below H3 in the mid-range of hearing. Also the good speaker damping is striking. But how does it sound?

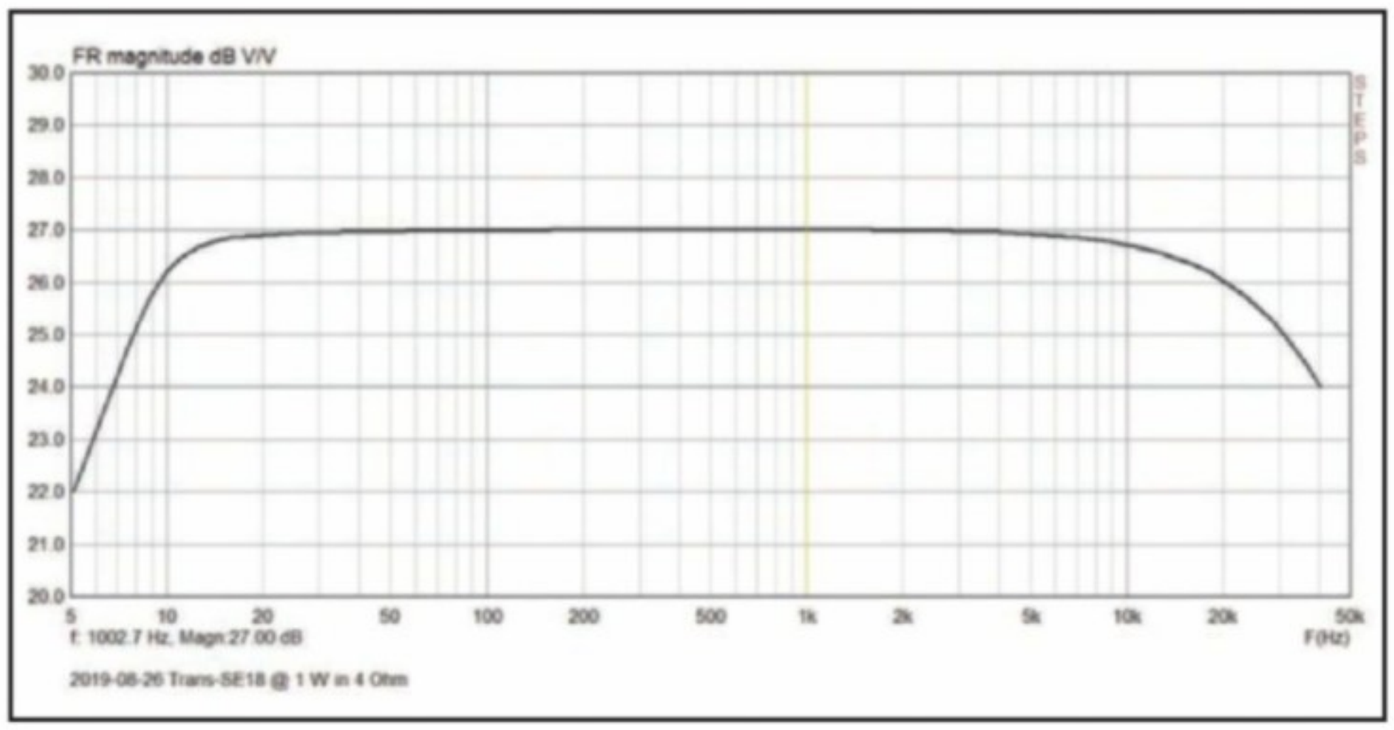

Figure 6: Frequency characteristics shown at 1 W in 4 Ω

Subjective Evaluation of the Trans-SE18

The amplifiers sound open with many details in the soundstage, from left to right and in depth and height and front stages. I know that the lowest H2 causes this detailed character. The bass reproduction is quick, as a result of the low output impedance, which creates good speaker damping. Also, the Trans-technique minimizes the magnetic transient distortions in the core of the OPT. Music consists of transients and Trans transfers those much cleaner. All these facts explain the good detail and space reproduction.

Next, the ease and fluency of listening. These are totally comparable with the standard 300B-SE. However, what is really striking is that the Trans-SE18 shows so many more details, without disturbing the joy of listening. After many years of tube amps design and evaluation, my preference clearly shifts to such details. I am not the only one reporting such observations; see others opinions. [8,9] Last, I would like to remark that the 18 W output power gives you impressive dynamics, especially with my 100 dB/Wm horn speakers.[10]

Summary

With Trans and pentodes and modern semiconductors and toroidal transformers, you can make an amplifier with remarkably low distortion, clean transient reproduction, and pleasant dynamics. Its sound character is very close to the famous 300B-SE amplifier. However, the bonus with Trans are the clean details. General opinion says declining H2-H3-H4-H5- sounds the best. My research says for optimal details ask for lowest THD with minimal H2 and largely reduced magnetic transient distortions.

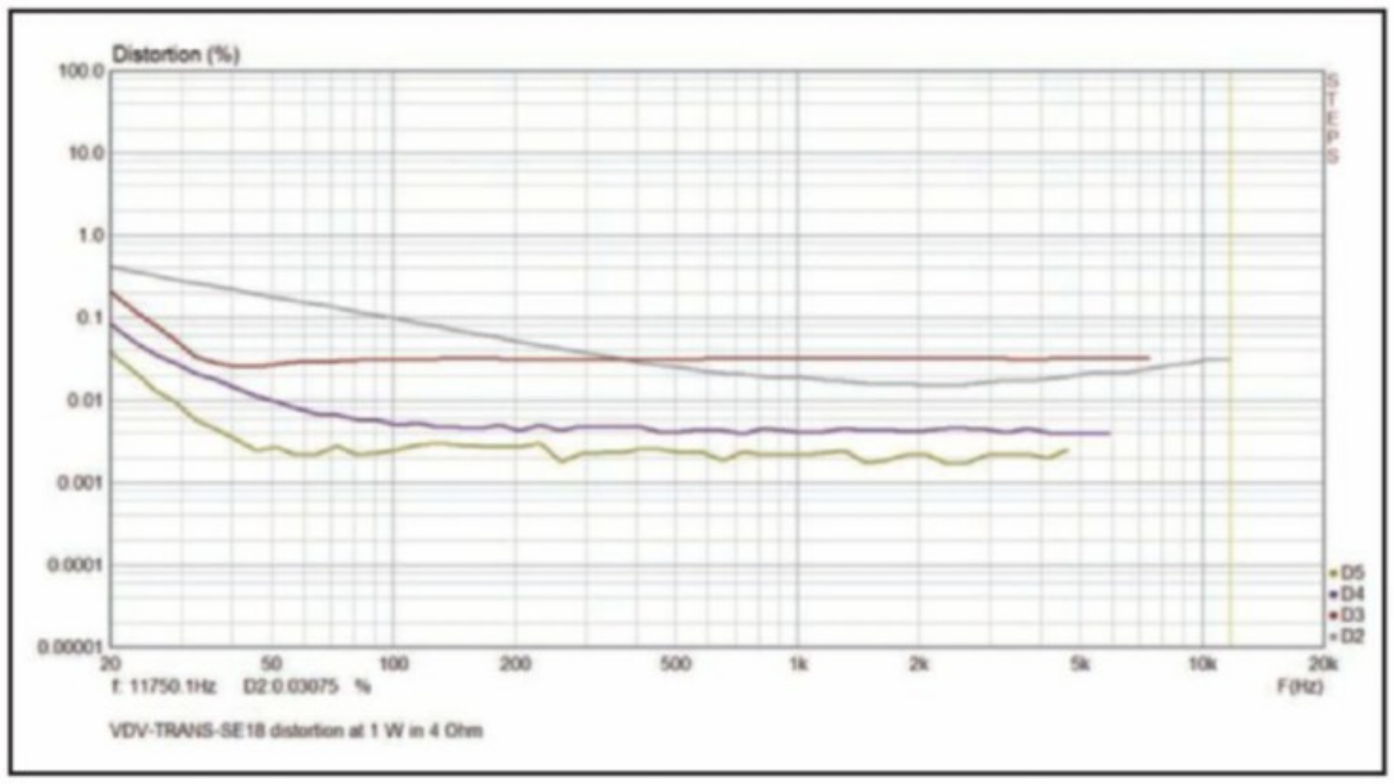

Figure 7: The harmonic distortions shown at 1 W in a 4 Ω load.

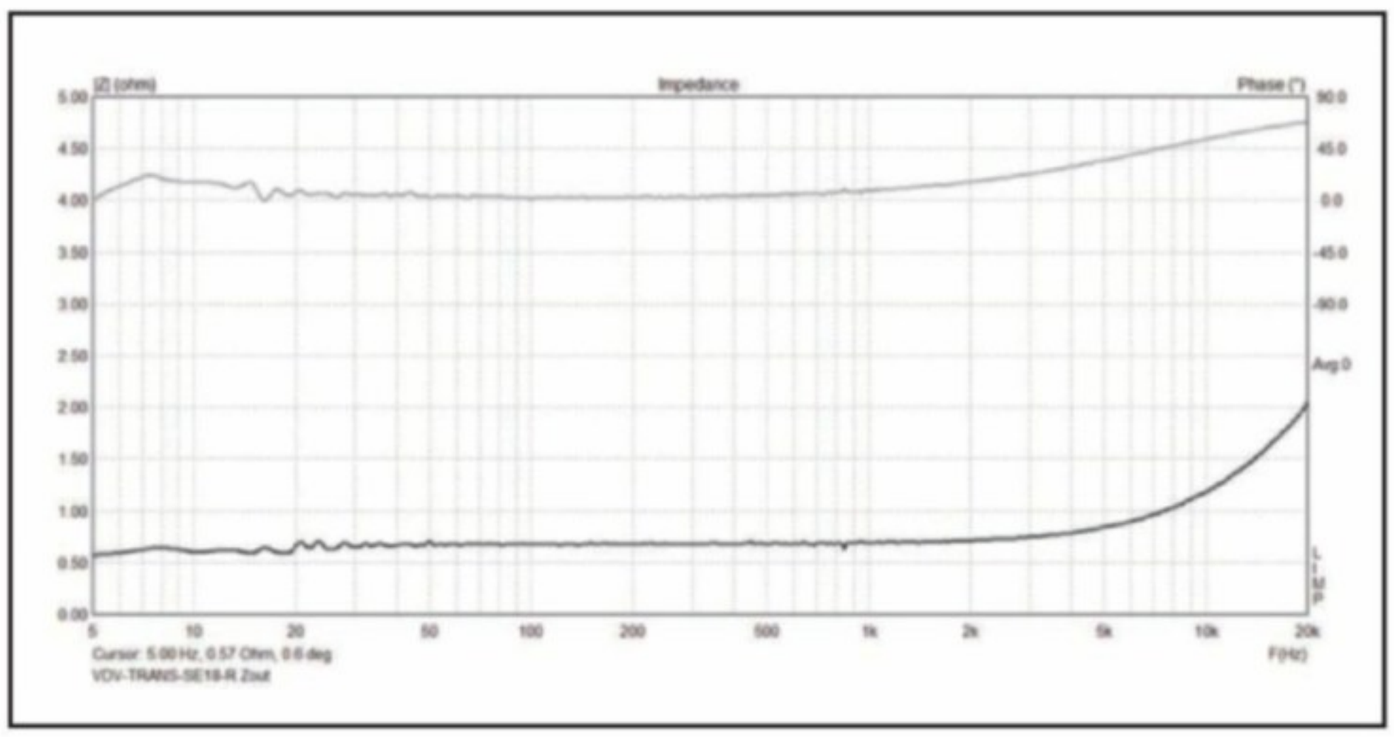

Figure 8: The upper curve is the phase, see the right hand scale. The lowest curve is the output impedance, see the left hand scale.

Resources

[1] M. van der Veen, *Vanderveen-Trans-Tube-Amplifiers*-e-book, www.elektor.com

[2] M. van der Veen, "TubeSociety 2A3-300B-SE Amplifier Project (Part 1)," *audioXpress*, April 2018.

[3] M. van der Veen, "TubeSociety 2A3-300B-SE Amplifier Project (Part 2)," *audioXpress*, May 2018.

[4] Vanderveen ir. Bureau and TubeSociety, www.mennovanderveen.nl | TubeSociety | project-2018

[5] Vanderveen ir. Bureau and TubeSociety, www.mennovanderveen.nl | Products | Kits | Trans-SE10

[6] Eidgenössische Technische Hochschule Zürich; *Elektroniker*, July 1989, pp.83..84

[7] Vanderveen ir. Bureau and TubeSociety, www.mennovanderveen.nl | Products | Transformers | SE-transformers

[8] P. Millet, "A Single-Ended 'E-Linear' Power Amplifier," *audioXpress*, April 2005, http://www.audioxpress.com/article/a-single-ended-e-linear-power-amplifier

[9] F. Blöhbaum, "Multiplied Transconductance Amplifiers (Parts 1 and 2), *Linear Audio* 6 and *Linear Audio* 8, www.linearaudio.net.

[10] Quintafila Niagara MSR: www.facebook.com/quintafila.it

Plug-In Modifications for Hollow-State Guitar Amplifiers?

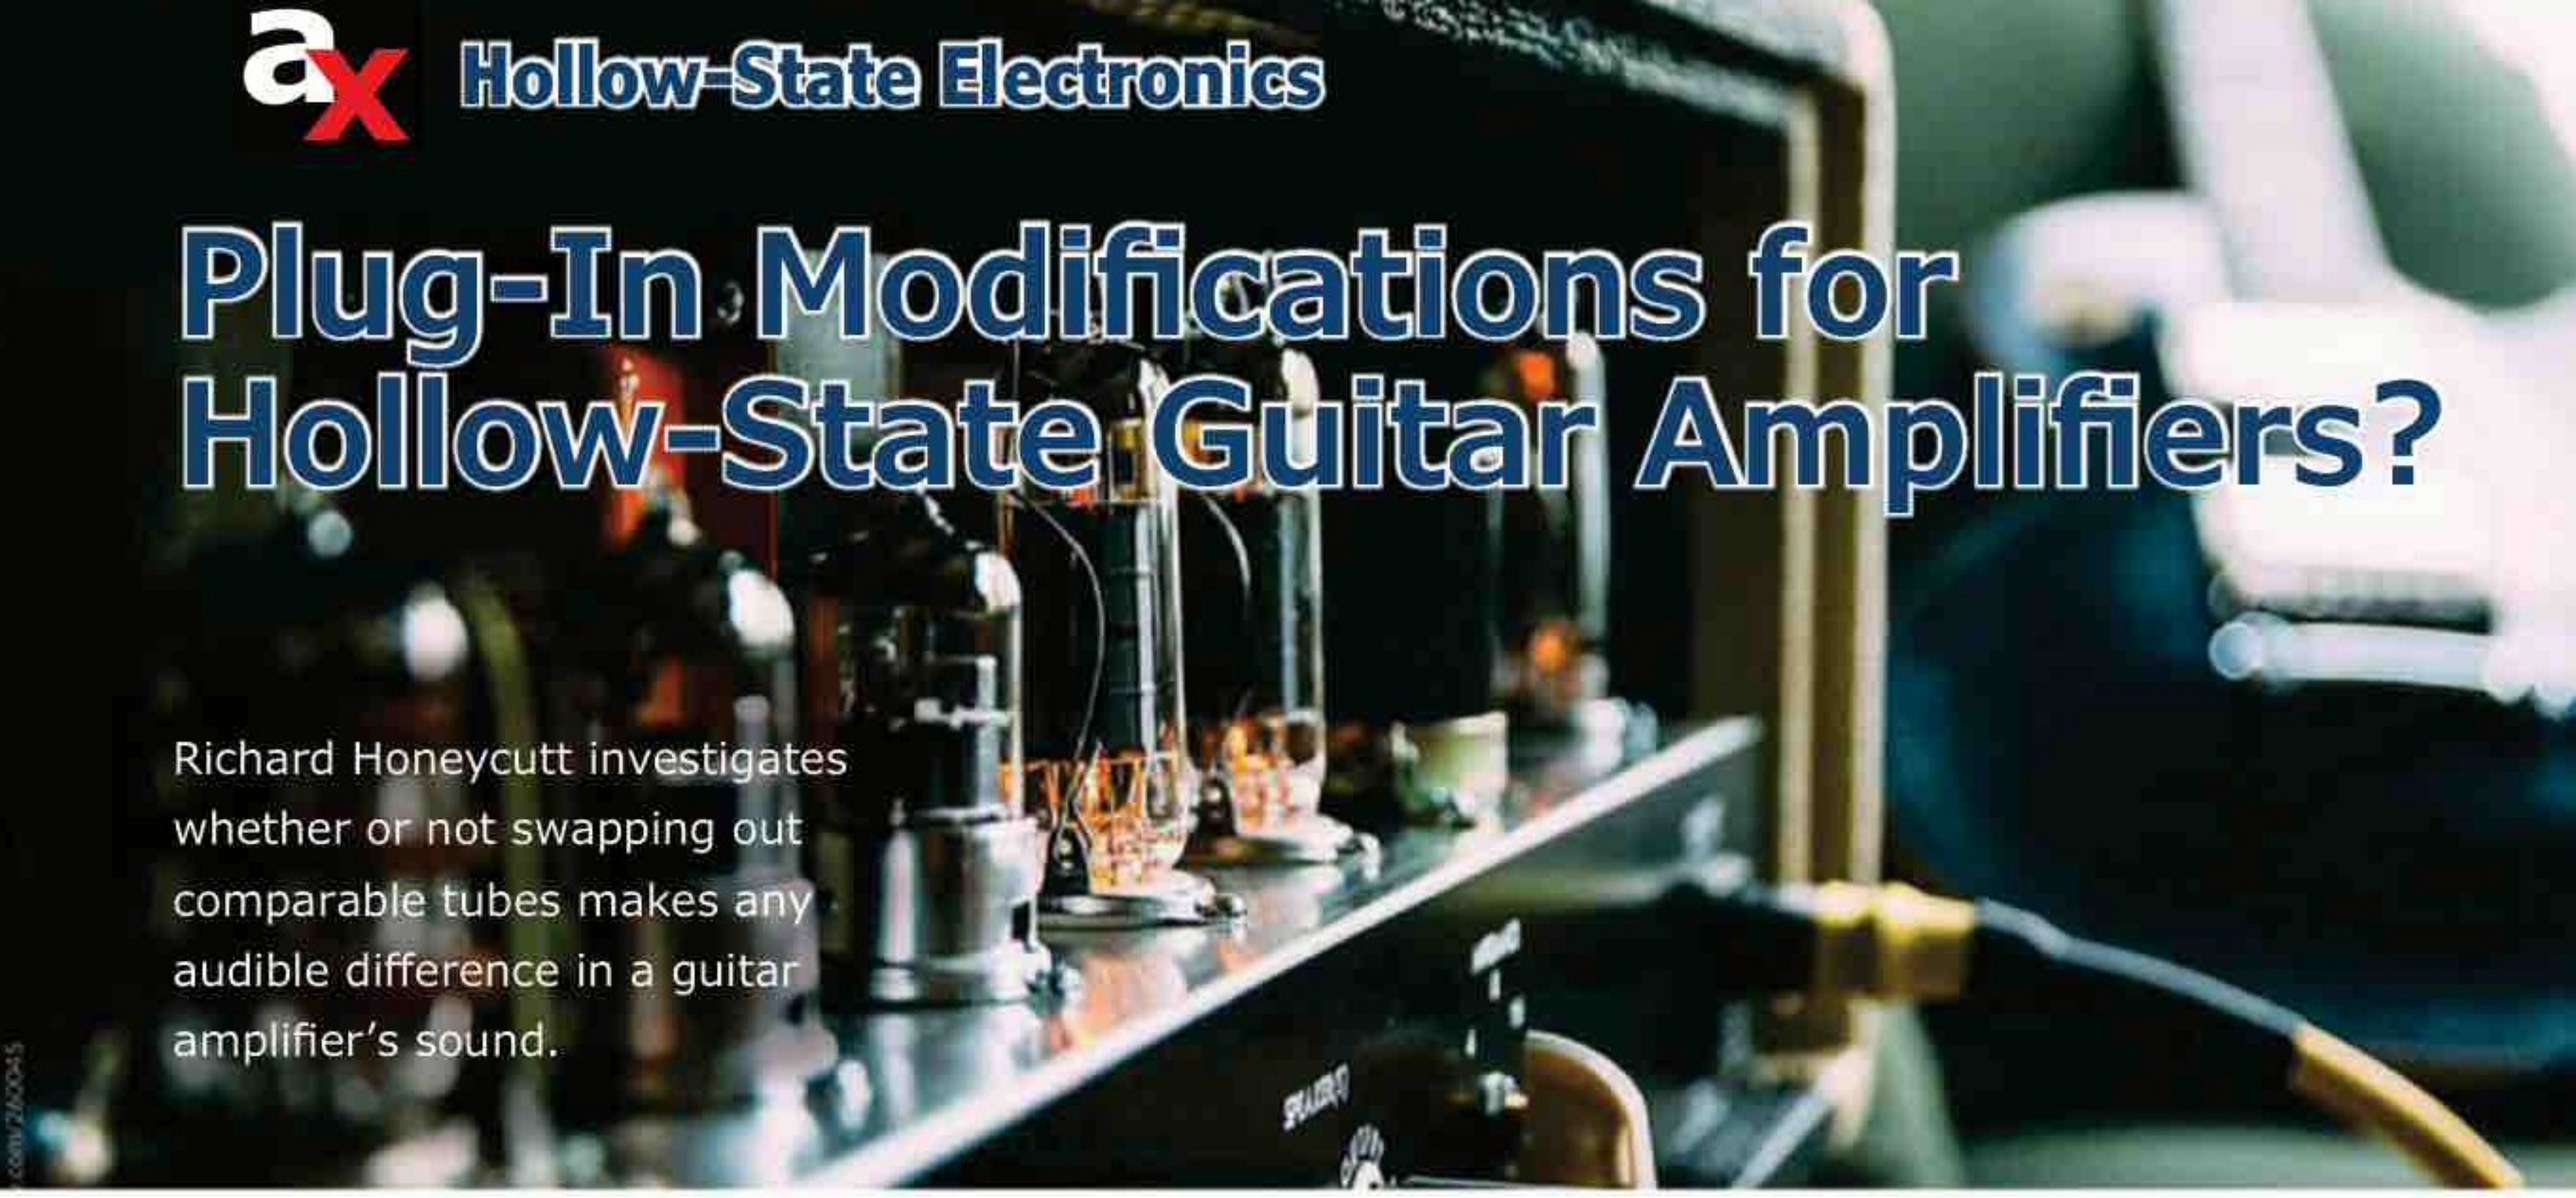

Richard Honeycutt investigates whether or not swapping out comparable tubes makes any audible difference in a guitar amplifier's sound.

By
Richard Honeycutt
(United States)

Back in the good old days of hollow-state electronics—before 1960—whenever your radio or TV quit working, the first step was to look at all the tubes to see if they were illuminated. If not, you'd go to the nearest drug store, hardware store, or repair shop and buy a replacement for any tubes that stayed dark. In the 1940s and 1950s, you might have pulled all the tubes and taken them to one of those businesses to be checked on their self-service tube tester.

We did these things not because tubes were the most likely components to fail (capacitors had that distinction), but it was something that almost anyone could do, ignorance of electronics notwithstanding. It was a lot like the story about a guy who was looking under a street lamp for a lost coin. When a policeman offered to help, he asked the man if he had lost the coin in that area, and he said, "no, but the light's better here!"

Guitarists' feelings about their amps run the spectrum from "I love my amp" to "What a sorry piece of junk; I've just got to replace this." Somewhere in the middle is "My amp's pretty good, but I'm sure I can find a way to modify it to make it better." Among this group, the first try is usually to see if switching some of the tubes will help. The reasoning is much the same for fixing a radio. It helps that many hollow-state guitar amps have a similar tube complement: 12AX7s and maybe 12AT7s for the preamp, vibrato oscillator, reverb driver, and phase splitter; and 6L6GCs for the push-pull output pair(s). Plug-in replacements for these tubes, having slightly different characteristics, abound.

A few years ago, several companies sold tubes that claimed to be different varieties of the 12AX7 and the 6L6GC: there were "hotter" and "quieter" 12AX7s, and 6L6GCs that could be subbed for the beefier 6550 or KT88. In my experience, there was little difference among the "different" 12AX7s, but the beefy 6L6GCs could indeed be used instead of 6550s, though with a shorter life span.

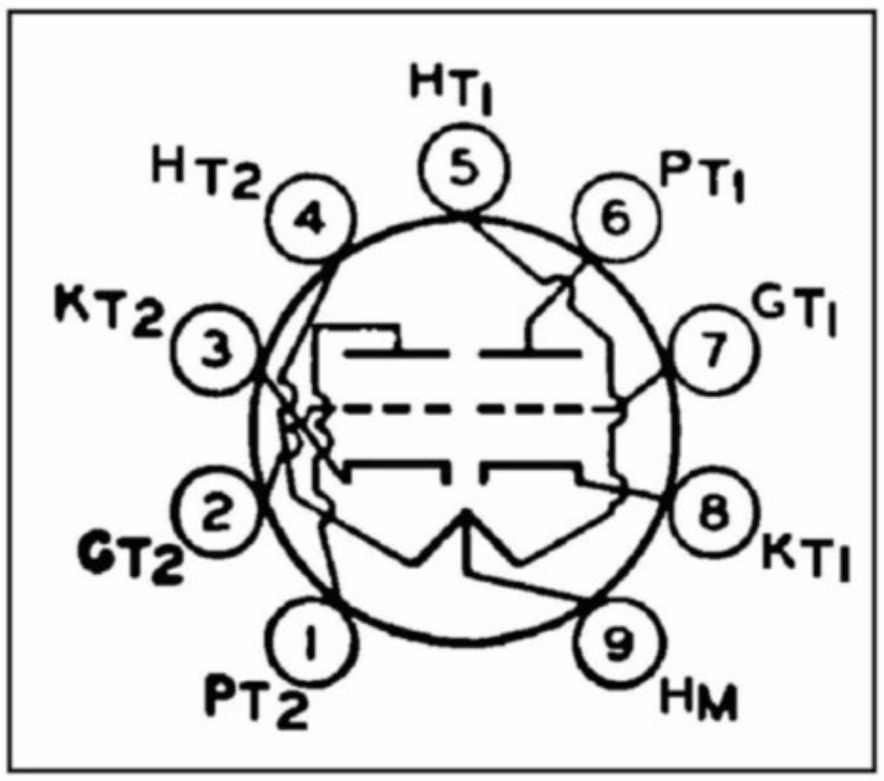

Figure 1: The 9A base is used by the 12A_7 class tubes.

Dual Triode Differences

Let's look at some real, significant differences among the 12AX7 family of dual triodes. The most common of these dual triodes is the 12AX7, which has a μ of 100. This tube is also known by its European model number: ECC83. The 7025 tube is a 12AX7 with a controlled noise and hum characteristic. The 12AT7 (ECC81) is similar to the

12AX7 in that it has a high μ. In this case, though, "high" is 60, not 100. The 12AT7 (ECC81) has a maximum plate dissipation of 2.5 W, as compared to 1.2 W for the 12AX7. This makes it useful for applications in which the tube must supply more current to the load, such as drivers for spring reverb units.

The 12AU7 was the standard dual-triode amplifier in electronic measuring instruments such as vacuum tube voltmeters and oscilloscopes. It

Figure 2: Here is a schematic of a Fender Twin Reverb amplifier.

Figure 3: This is an LTspice model of a Twin Reverb amplifier.

has a medium μ of about 20. The 12AY7 also has a medium μ, but in its case, "medium" is 40. The 12AY7 is not commonly used today.

All of the 12A_7 tubes use a 9A base (see **Figure 1**), and thus are direct plug-in replacements for each other, connection-wise.

Probably the most-copied hollow-state guitar amplifier is the Fender Twin Reverb. Its schematic is shown in **Figure 2**. Many manufacturers use the same basic circuit: two input channels, "Normal" and "Effects." The effects channel of the Twin provides tremolo (erroneously called "vibrato" on the schematic) and reverb.

On many amps of newer design, the effects channel has an extra gain stage to permit introduction of distortion via a channel-switching footswitch. The input tubes, and the ones following the tone controls, are 12AX7s or 7025s. The phase splitter is a differential amplifier (sometimes called a "long-tailed pair"), and it uses a 12AT7. We'll not dwell on the tremolo and reverb circuits in this article. The output stage is two paralleled push-pull pairs of 6L6GCs.

In order to investigate the effects of swapping the 12AX7s and 12AT7, I built an LTSpice computer model of the Twin Reverb, shown in **Figure 3**. Since the focus of the investigation was the effect of swapping certain tubes, I simplified the model by omitting the tone controls and effects circuitry. Because I used a 1-kHz test signal, and the tone/volume controls at settings 5,5,5,5 attenuate the signal by about 22 dB, I inserted resistors R9 and R11 to make the levels work out correctly. I replaced the power supply with five DC voltage sources. In order to account for the droop in B+ that occurs at high output levels, I added R20, a 33 Ω resistor, in series with the 460 V high B+ supply. Not having complete details on the output transformer, I made an LTSpice subcircuit using coupled conductors having the correct primary and secondary inductances to provide the proper turns ratio.

Testing the Tubes

I made the first test run using the original tube types: 12AX7s for preamps and a 12AT7 for the phase splitter. **Figure 4** shows the output wave with the tone and volume controls halfway up, and an input signal of 15 mV_p. As you can see, under the test conditions, we get some clipping at 23 V_p output (33 W into 8 Ω). **Figure 5** shows a Fast Fourier Transform (FFT) of the output wave, indicating

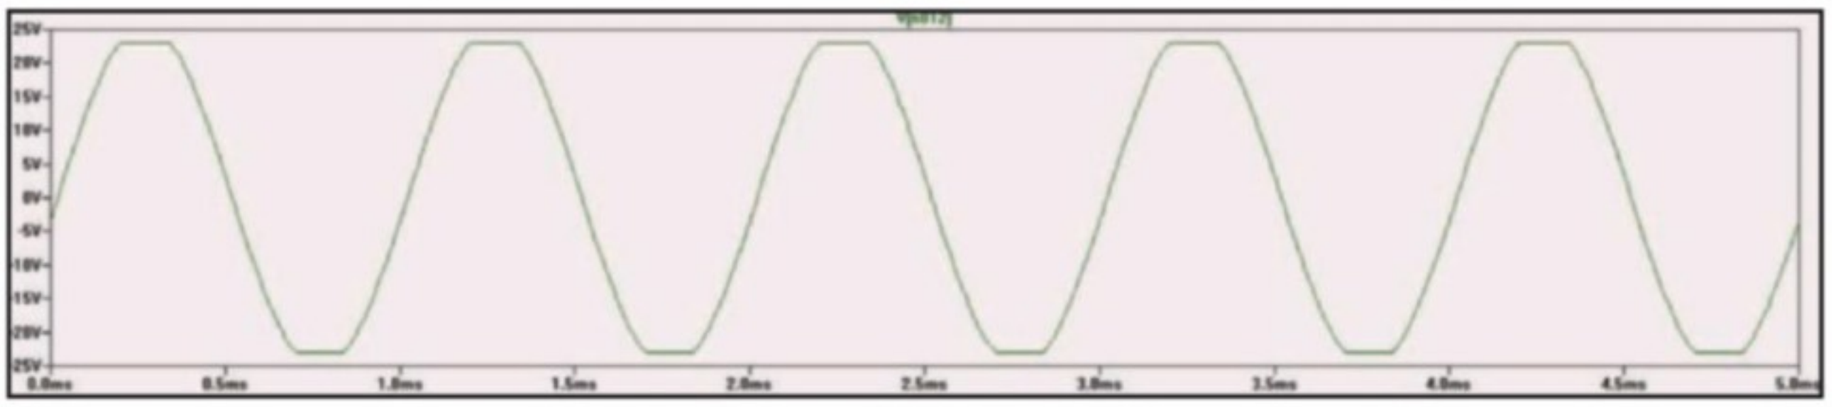

Figure 4: The output wave is shown with the original tube types.

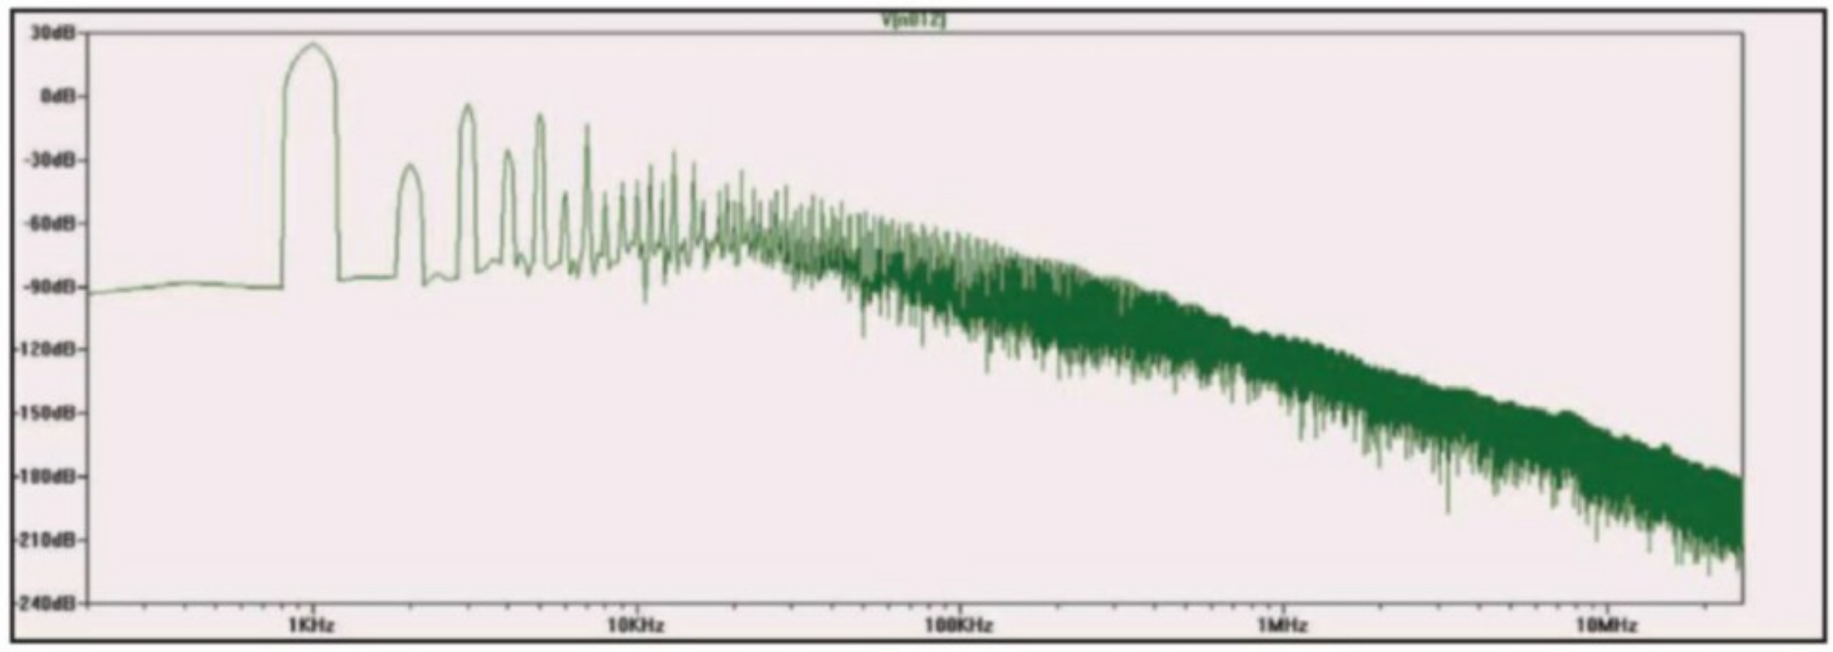

Figure 5: Here is the Fast Fourier Transform (FFT) of an output wave using the original tube types.

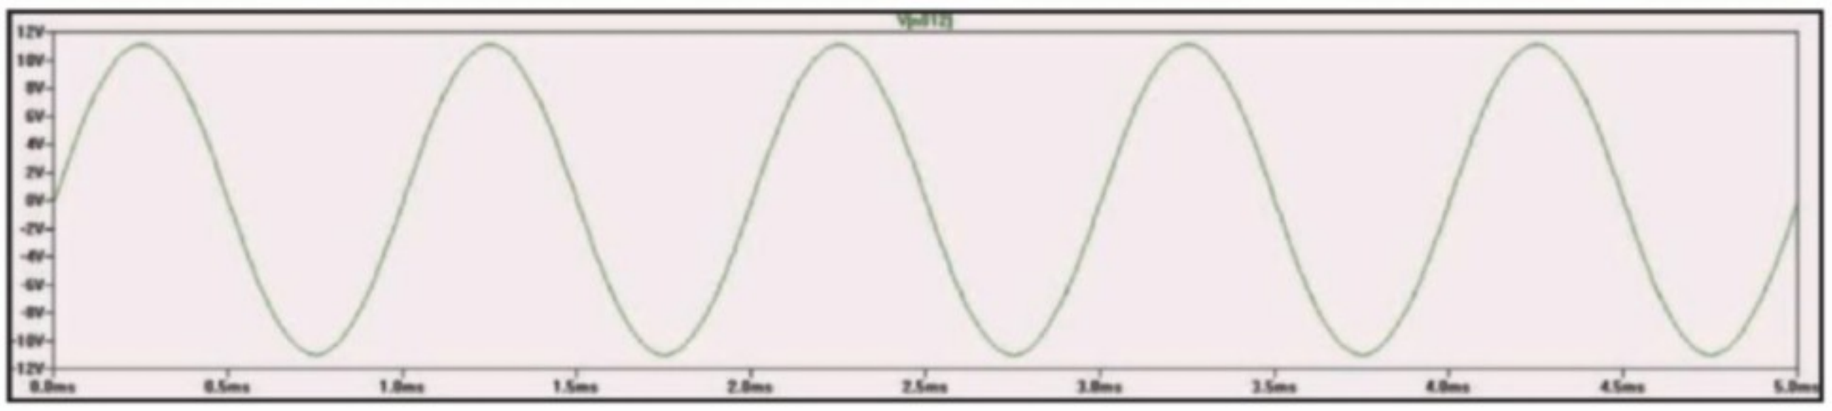

Figure 6: The output wave is shown using the 12AT7 as the input tube.

predominantly odd-order harmonics (about 4.5%), with lower levels of even harmonics (about 3%).

Substituting a 12AT7 for the input tube gives the output shown in **Figure 6**. With the lower gain of the 12AT7, we get a peak output of 11 V, and no visible clipping. The FFT (see **Figure 7**) shows that there are still significant harmonics present: about 0.28% odd harmonics and 0.22% even harmonics.

If we use a 12AU7 for the input tube, the output wave changes (see **Figure 8**). Since the 12AU7 has an even lower gain than the 12AT7, the output voltage is only 1.4 V_p. Harmonic distortion at this low level is only about 0.1%. Under normal circumstances, the volume control setting would be increased if the 12AT7 or the 12AU7 were used, so the output level variations would be cancelled out by this readjustment. If we readjust the values of R9 and R11 to simulate a volume control readjustment that would bring the output with the 12AT7 up to the level of the stock Twin Reverb with a 15 mV input, the output waveform and its FFT are almost indistinguishable from those of the stock configuration. This is not surprising, since the phase splitter stage is the first to clip in this amp, so

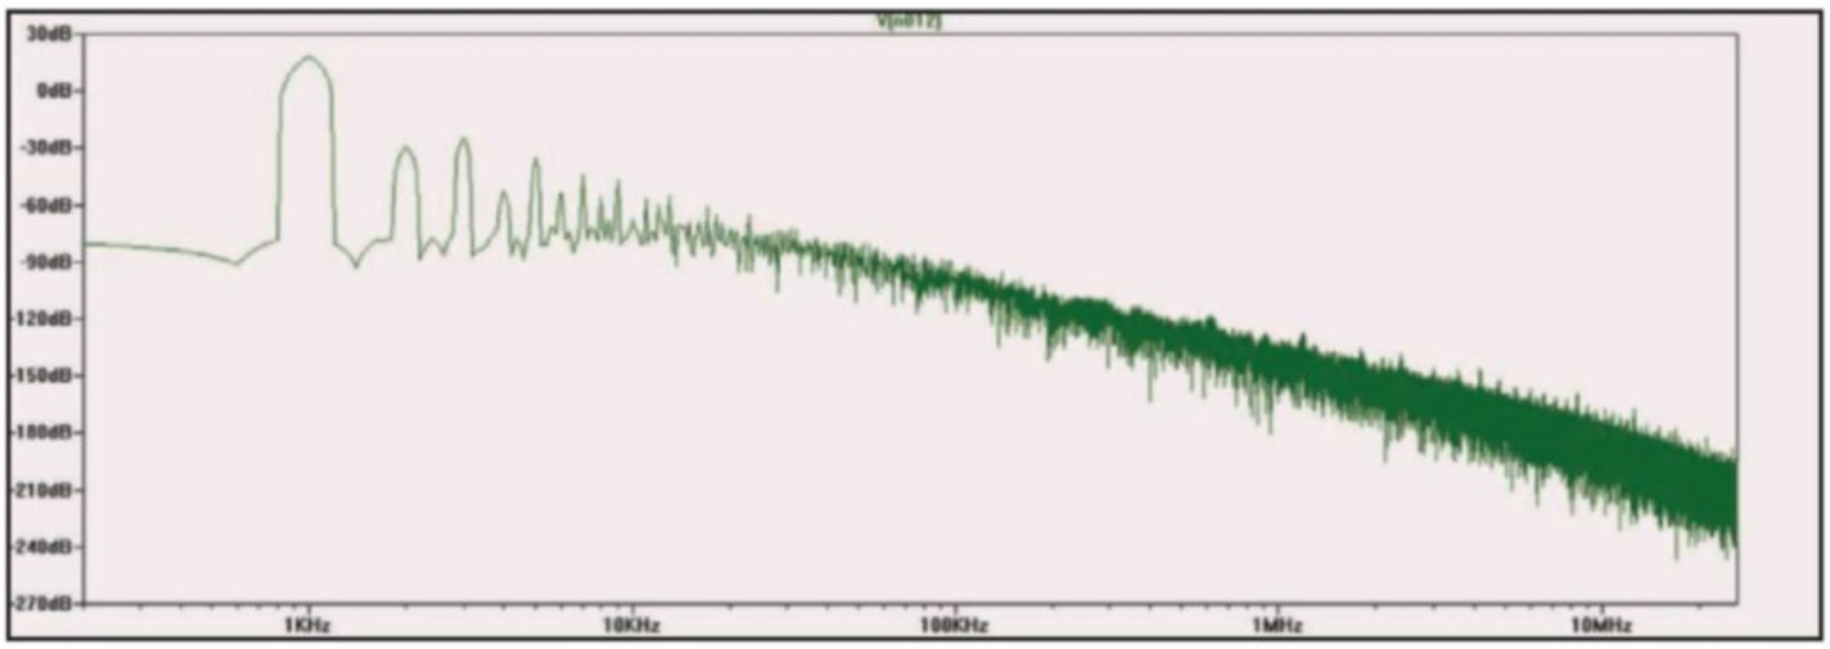

Figure 7: Here is the FFT of an output wave using the 12AT7 input tube.

Celebrating 40 Years!

ACOustics begins with **ACO**™

Measurement Microphones
Polarized and Electret

Community/Industrial
Noise Monitors/Alarms
Indoors and Out

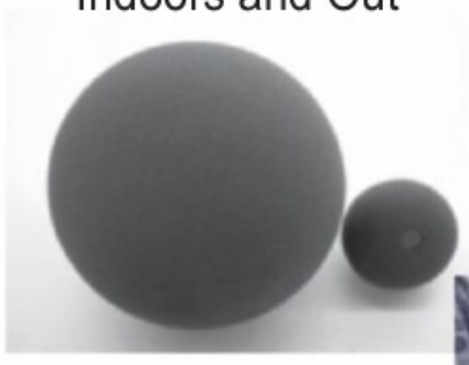

Windscreens
3 to18 inch & Custom
Weather/UV Resistant

Microphone Systems
Standalone
IEPE Powered
Phantom Power
Custom

ACO Pacific, Inc
www.acopacific.com sales@acopacific.com
Tel: +1-650-595-8588

Ready, willing and
AVEL

offering an extensive range of ready-to-go
toroidal transformers
to please the ear, but won't take you for a ride.

Avel Lindberg Inc.
47 South End Plaza, New Milford, CT 06776
p: 860.355.4711 / f: 860.354.8597
sales@avellindberg.com • www.avellindberg.com

changing the preamp tubes would not affect the distortion signature noticeably. The same effect would be expected if a 12AU7 were used for the preamp tube.

Only if the input were very hot—either from a guitar with extremely high-output pickups, or from a booster amp being used—would changing the input tube significantly affect the distortion sound. In that situation, the volume control would have to be turned very low and the first stage would provide most of the distortion. And while we're discussing this, let's note that guitarists who claim that a certain tube provides better bass or more brilliance are self-deceived, since nothing in either preamp, phase-splitter, or power-amp tubes affects frequency response within the range handled by a guitar amp.

So, if the phase splitter is where clipping shows up first, what happens if we swap that tube? **Figure 9** shows the output wave you get if you leave the 12AX7 preamp tube in a Twin Reverb and swap another 12AX7 for the phase splitter instead of the stock 12AT7. Not surprisingly, the added gain of the 12AX7 only makes the clipping a bit more prominent than it is for the stock tubes (as shown in Figure 4). The FFT shown in **Figure 10** indicates very nearly the same distortion signature as for the stock condition.

Figure 11 shows the result of subbing a 12AU7 for the phase-splitter: the output voltage is lower due to the lower μ of the 12AU7. Again, the difference in output level would probably be made up in a performance by the volume control being turned higher, nullifying the effect of swapping tubes.

So what can we say about those dedicated amp-tweakers who are sure they have improved their sound by swapping tubes? Two comments:

- The placebo effect is very real, and people can really hear a difference that is not objective, in the sense that others may not be able to hear it.
- Our aural memory is very short. In fact, the time required to turn an amp off, swap a tube, turn it back on, and let the new tube warm up is longer than our aural memory. If we can't truly remember what the amp sounded like with the old tube, placebo effect has a better chance of deceiving us.

The Output Tubes

So far, we've only discussed the small-signal tubes. But what about the output tubes? Since

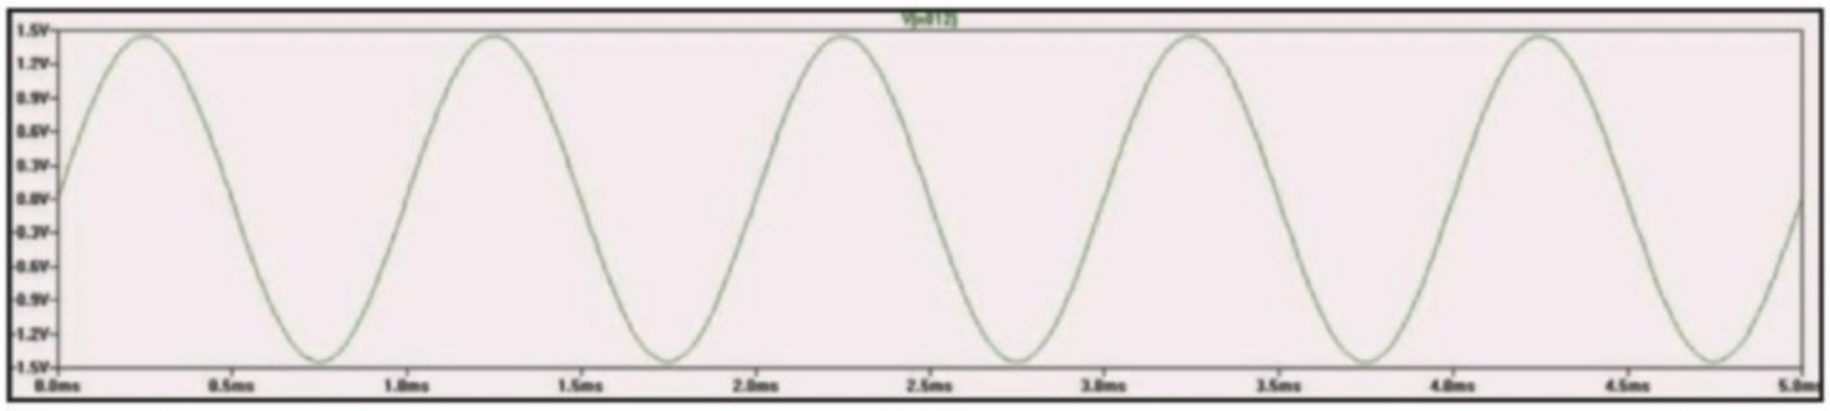

Figure 8: This is the output wave using the 12AU7 input tube.

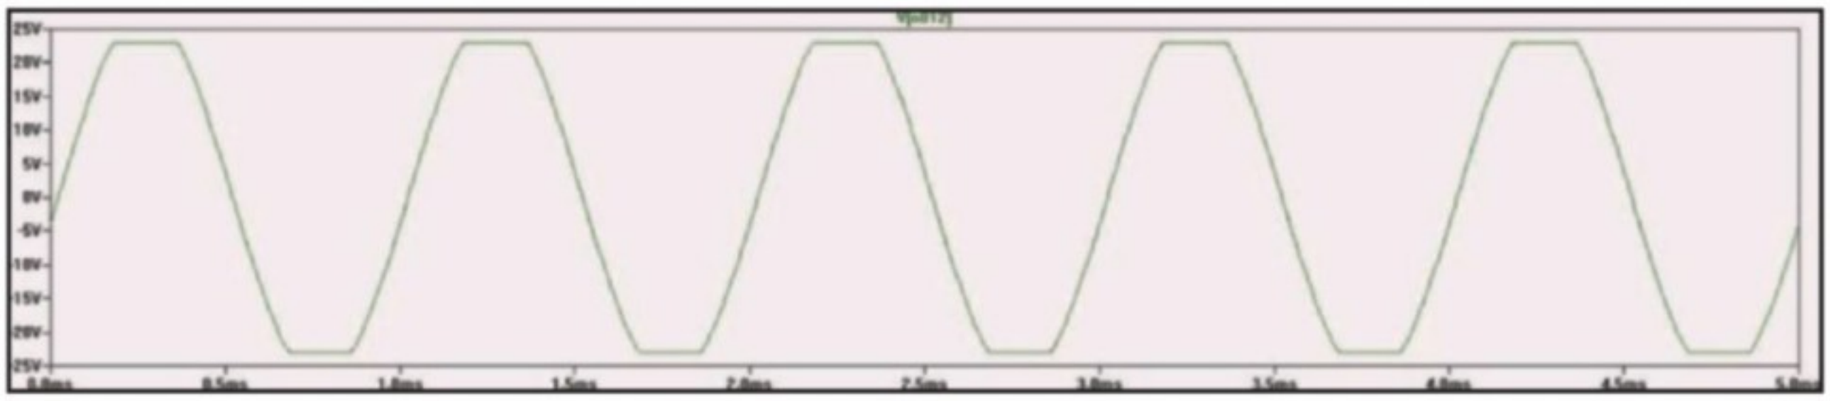

Figure 9: The output wave is shown with the 12AX7 preamp and 12AX7 phase splitter.

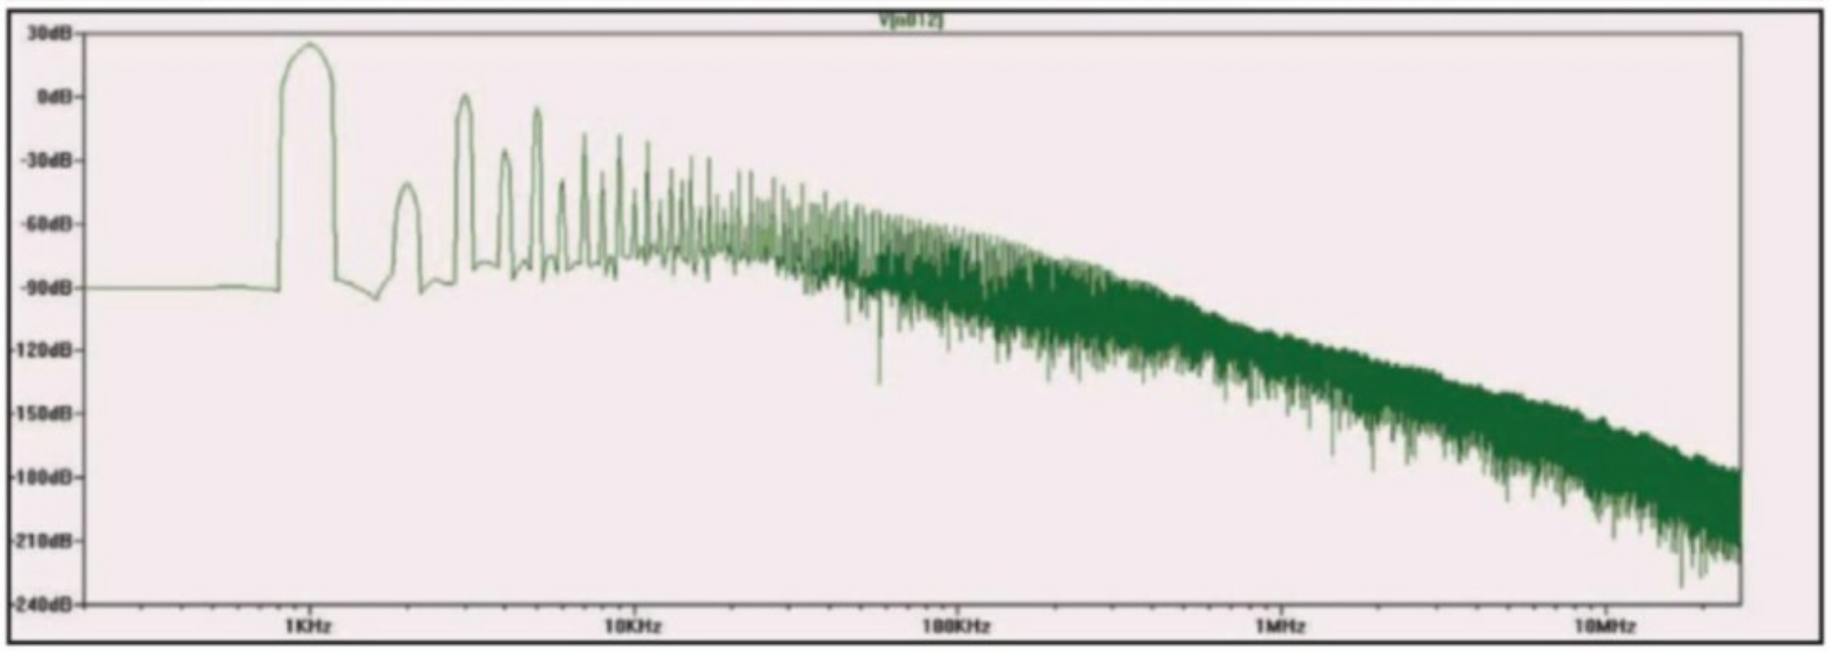

Figure 10: Here is the FFT of the output wave that was shown in Figure 9.

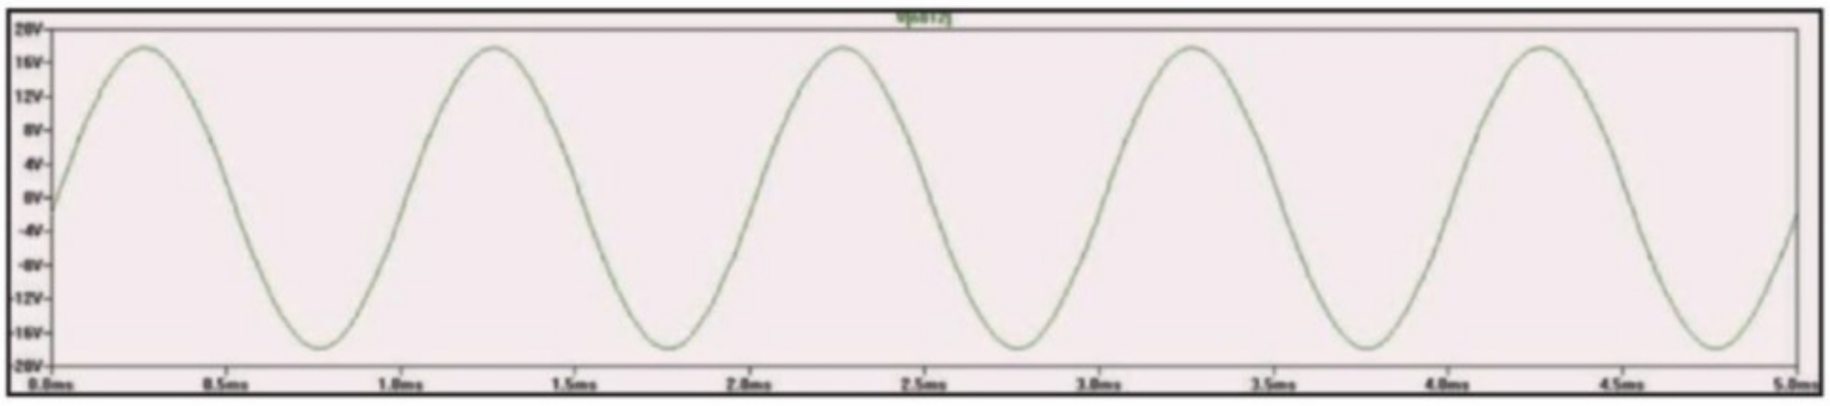

Figure 11: This is the output wave with 12AU7 used as a phase splitter.

different output tubes may interact differently with the output transformer, and I do not have access to a really accurate model of the Twin's output transformer, I didn't investigate swapping output tubes. However, anyone who decides to try swapping other tubes for the stock ones in his/her amp needs to be aware of several cautions:

Never swap a 6V6GC for a 6L6GC, since the 6V6 has a maximum plate voltage of 350 V, and circuits designed for 6L6s often have plate supplies approaching 500 V. The same applies to a 5881.

Never swap an EL34 for a 6L6, because their basing is different (8ET versus 7AC), and you may let out the magic smoke in your amp by making this swap. Amps don't work without the magic smoke! Same goes for swapping a 7027 for a 6L6, except the 7027 uses an 8HY base.

You're usually safe swapping a 6550 for a 6L6, as the only likely issue is that the 6550 draws 1 A of heater current, compared to 0.9 A drawn by the 6L6. Your power transformer may run slightly hotter, but probably not enough to seriously shorten its life.

If you stumble across an old metal 6L6 (as opposed to the 6L6GC), don't be tempted to use it instead of a 6L6GC. The 6L6 has substantially lower maximum plate voltage and current ratings than the 6L6GC, and may fail catastrophically, perhaps taking the output transformer with it.

Will a 6550 sound different from a 6L6GC? To borrow a phrase from Star Trek's Mr. Spock, "Insufficient data, Captain!" ax

Industry Calendar

Here are a few places where you might find a copy of *audioXpress* and possibly meet one of our authors and staff members.

January 16-19, 2020
The NAMM Show 2020
Anaheim Convention Center, Anaheim, CA
www.namm.org/thenammshow/2019/
attend

The 2020 NAMM Show returns to Anaheim, CA, January 16–19—a little earlier than usual. The annual global gathering of the musical instruments, pro audio, and entertainment technology industries will be on display across the NAMM Campus at the Anaheim Convention Center (ACC) and the excitement for this 2020 edition couldn't be greater.

The National Association of Music Merchants (NAMM) Show, the global crossroads of music products, pro audio and sound, and live event technology industries will offer visitors a dynamic way to cover all aspects of music: from products to pro sound and live production. This will be a consolidation of the reinforced show that positively surprised the industry in 2018, and will offer an overview of new products from 7,000+ brands, combined with special live music events, celebrity appearances, dozens of educational sessions tied to today's trends, and the world's largest gathering of top music and audio technology innovators, artists, and professionals.

Expanding on the concept introduced in 2018 with the opening of the new AAC North halls, the NAMM Show is now creating new categories for better identification of activities and exhibits during the event. Labeled as Music, Sound, and Live, visitors will be able to follow their interests as they visit the different halls, conferences, and outside event areas.

In the Music areas, visitors will be able to discover a wide variety of the latest instruments, accessories, and solutions from leading global brands. In the Sound areas, visitors will be able to meet the industry's leading audio experts, and experience new product trends in studio, recording, and pro audio technologies. Finally, the Live label identifies where visitors can find entertainment technology at The NAMM Show.

Since 2018, the NAMM show has gained the full scope of an AES West Coast convention with the exhibits in the new AAC North Level 1 hall, plus the AES@NAMM sessions at the Anaheim Hilton Hotel, while the new Level 2 hall looks effectively like a condensed version of a Prolight+Sound show (the sound part of it...). All this, while the traditional Hall A gains a new life with more space for exhibitors in the home recording and stage product segments. The traditional halls at the NAMM show also benefit largely from wider walkways, and more dispersed noise clusters.

The Technical Excellence and Creativity (TEC) Awards also take place during The NAMM Show on Saturday, January 18.

And once again, taking place concurrent with The NAMM Show 2020, the Audio Engineering Society (AES) Academy will provide a series of hands-on training sessions and technical presentations from expert presenters. And top sound manufacturers will return to present a showcase of line array and loudspeaker technologies at the second annual Loudspeaker System Showcase inside the ACC arena, giving participants the opportunity to demonstrate both their portable systems in a real-world setting to a mix of industry professionals.

February 11-14, 2020
Integrated Systems Europe (ISE) 2020
Amsterdam RAI, The Netherlands
https://www.iseurope.org
www.avixa.org | www.cedia.net
www.iseurope.org

Integrated Systems Europe

ISE 2020 is set to be the most compelling Integrated Systems Europe show ever. Taking place at the RAI Amsterdam from February 11-14, ISE is the world's leading event for the professional systems integration channel and end-users of AV technologies and solutions.

Launched in 2004, ISE 2020 will be the final edition held in Amsterdam—and is confidently expected to be the biggest in the show's history with more attendees and exhibitors than ever before. One of the factors enabling this growth will be the construction of a permanent extension to Hall 5. ISE 2020 is expected to draw nearly 1,400 exhibitors and more than 85,000 registered attendees to its Amsterdam RAI location. The event is a joint venture of the Audiovisual and Integrated Experience Association (AVIXA) and the Custom Electronic Design and Installation Association (CEDIA).

In response to the growing demand for exhibition space at the RAI Amsterdam, ISE is also spreading beyond the RAI by hosting a range of new and established conferences at the nearby five-star Hotel Okura.

Once again, the Audio Forum@ISE will investigate the concept of sound environments and settings. It will explore the relationship between quality sound propagation, architectural and environmental acoustics, the listening experience, acoustic wellbeing and related psychological implications. Audio Forum@ISE is produced by long-standing ISE media partner Connessioni, in collaboration with the Audio Engineering Society (AES).

The ISE exhibition floor is sectioned into different technology zones, giving attendees a concentration of relevant exhibitors in designated areas. This makes navigating around the sizeable show floor, with almost 1,400 exhibitors, more straightforward. The seven technology zones are: Audio & Live Event Technology, Digital Signage & DooH, Education Technology, Residential Solutions, Smart Building, Unified Communications, and the Innovation Zone.

The audio chain, from microphones and headphones to loudspeaker systems and networking, is an ever-prominent force at the show. As, too, are the visual ingredients necessary for making live performances truly unforgettable. Shows are getting bigger and more extravagant as the focus shifts from recorded performances to live.

February 24-27, 2020
GSMA Mobile World Congress
Fira Gran Via, Barcelona, Spain
www.mwcbarcelona.com

GSMA MWC Barcelona
24-27 February 2020

MWC Barcelona will take place February 24-27, with events also staged at Fira Gran Via, Fira Montjuïc, and La Farga L'Hospitalet. Under the theme of "Limitless Intelligent Connectivity," MWC Barcelona will outline the convergence of artificial intelligence (AI), Big Data, and connected devices, all powered by 5G. CEOs from companies including Airbus, Nokia, Novartis, NTT Docomo, Orange, Telefónica, and Viacom will be speaking in the 2020 MWC Keynote Program.

As in previous editions, MWC Barcelona will be a crucial event for many technology and solution companies in the wearable, hearables, smart home, and portable/personal audio space—some of the strongest trends in this event. From semiconductors used in voice, DSP, and wireless audio to MEMS microphones and speakers and all the latest mobile platform technologies, including active noise cancelling and voice detection and recognition, the Barcelona event combines all key suppliers and vendors in a few convenient meeting/demo room areas within two halls, with some of the largest companies located at themed areas with their own booths.

MWC Barcelona is also an important opportunity for companies and professionals in the Media and Entertainment sectors, allowing a vision of the evolution of key technologies and platforms, such as 5G networks, or getting a specific mobile perspective on the solutions for Content Creation, Gaming, and Immersive Experiences.

MWC Barcelona will bring together leading players from across the mobile ecosystem, as well as adjacent industry sectors such as automotive and consumer electronics, highlighting the latest technologies, products and services. More than 2,400 companies will participate at the event and there will be more than 40 country and regional pavilions. As one of the key destinations at MWC, the GSMA Innovation City will return to highlight how mobile products and services are transforming how we live, work and play.

Flexibility Redefined

The power & versatility of APx500 software paired with 3rd party audio interfaces

APx500 Flex Audio Analyzer

Audio precision

www.ap.com

Accurately Measure Speaker Parameters ***Faster Than Ever!***

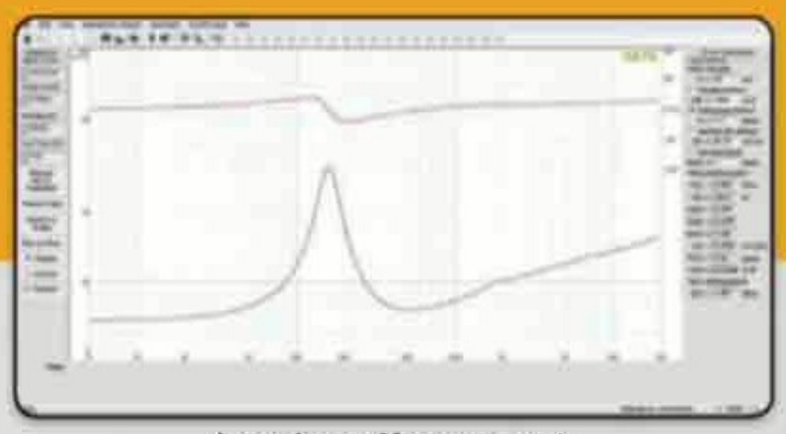

Impedance Measurement

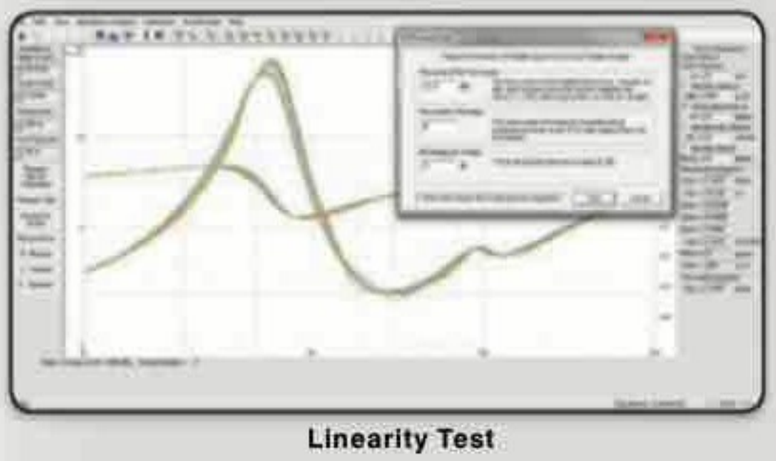

Linearity Test

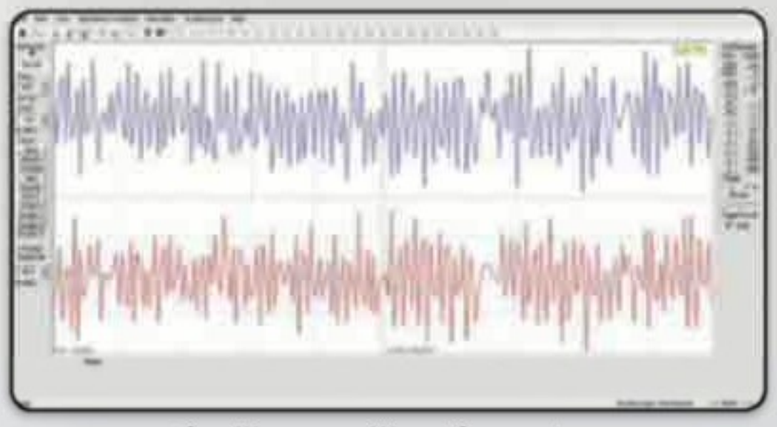

Oscilloscope Wave Generator

DATS V3

DAYTON AUDIO TEST SYSTEM

DAYTONAUDIO

Dayton Audio has taken the most complete and easy-to-use audio test system and improved it in almost every way when creating the DATS V3 Computer Based Audio Component Test System.

Expanding on the accuracy, capability, and reliability of the DATS V2, the DATS V3 sets a new standard for quick and accurate audio component measurements.

DATS (Dayton Audio Test System) V3 is the essential tool in any speaker designer's or audio engineer's toolbox, providing accurate, precise measurement data for any audio transducer or loudspeaker system. Save yourself time, money, and trial-and-error frustration by using DATS V3 to measure loudspeaker drivers, crossover networks, and complete audio systems!

DATS V3 is a unique, multipurpose test system that works with your PC and full-featured test software to measure the impedance response of a speaker driver and derive a complete set of Thiele/Small parameters that describe its operation. It's as easy as connecting a driver and starting the test!

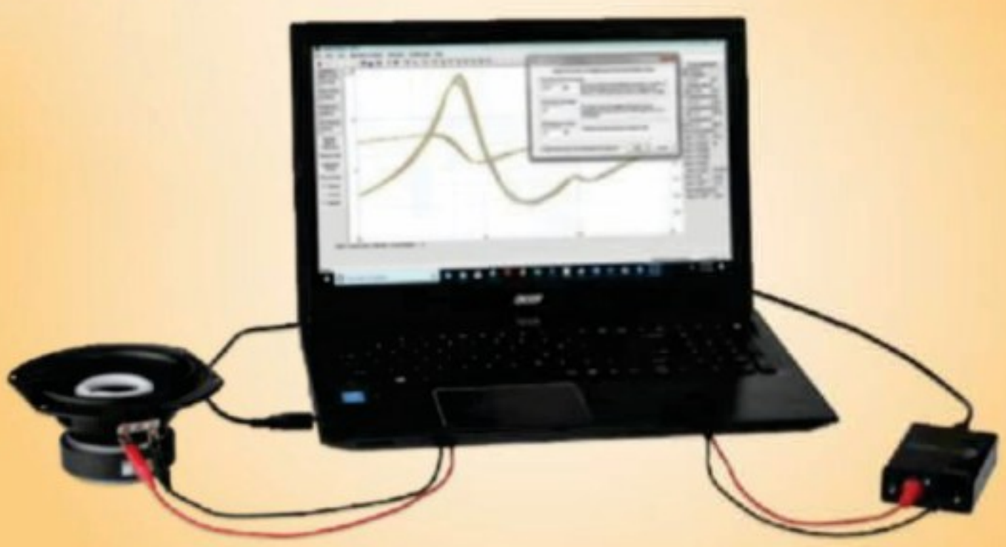

Improvements include:

- Twice the output capability
- Tighter tolerances for improved accuracy
- Double the data points for more precise measurements
- NEW DATS Linearity Test – easily compare multiple drive levels
- Rugged extruded aluminum housing
- Built-in precision calibration resistor
- Detachable test leads and USB cable
- On or under-desk mounting brackets

NOW AVAILABLE!

Custom OEM inquiries welcome.

SAVE $6
off $50 order!

Code: AXMQ3
Expires 12/31/2019

Not valid with any other offer, on prior purchases or backorders.
Not valid on Commercial Sales accounts. No cash value.
Valid for one time use.

Distributed by:

PARTS EXPRESS
AUDIO COMPONENTS & SOLUTIONS

parts-express.com
1-800-338-0531