INNOVATIONS IN AUDIO • AUDIO ELECTRONICS • THE BEST IN DIY AUDIO

www.audioxpress.com

audioxpress

ADVANCING THE EVOLUTION OF AUDIO TECHNOLOGY

Fresh From the Bench

Benchmark HPA4 Headphone/Line Amplifier

By Gary Galo

Hooke Verse Bluetooth Binaural Recording Headphones

By Luke McCready

You Can DIY!

The Tenacious Bass-8 Powered Subwoofer

By Thomas Zarbo, Jr.

Updated High-Voltage Delay for Tube Amplifiers

By Jan Didden

Audio Scope

A New Column for Audio Experts!

By Michael Steffes

Audio Electronics

Mu Followers and Mu Stages: What's Mu with You Two?

By Christopher Paul

Sound Control

Is Your Neighborhood Acoustically Friendly?

By Richard Honeycutt

R&D Stories

Achieving Frictionless Lower-Power Application Connectivity with Data-Over-Sound

By Dr. Adib Mehrabi

$9.00US $10.00CAN

NOVEMBER 2019

Measurement Mics from AP

Calibrated & TEDS-compliant
1/4” & 1/2” microphones
CCP- & Phantom-powered models

Results you can trust.

Audio precision

www.ap.com

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

November 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

From Spatial Awareness to Room Emulation

Having recently visited the Rocky Mountain Audio Fest (RMAF) show in Denver, CO (read my report online—https://bit.ly/2m8dWfW), I was impressed by how different the systems being demonstrated sounded. Sometimes it had to do with the signal source, others with the exquisite (and expensive) audio chains in use, but I was also particularly impressed by how a simple pair of Active Tower ATC SCM50SE loudspeakers sounded, simply connected to an ATC CDA2 Mk2 CD/DAC Preamp—an experience close to the brand's studio monitors, with which I am probably more familiar. More than once, the room acoustics did not help, in particular with speakers that were simply too large for hotel rooms.

One of the first rooms I visited, and which left me with lasting positive impressions, had the Paradigm Persona 7F loudspeakers and Persona SUBs demonstrated with Anthem STR electronics. As I noted, it was probably the fact that they had measured the room and were using the Anthem Room Correction system (ARC—https://www.anthemarc.com). The positive impression lasted to the point that many other excellent systems after that sounded poor in comparison.

The ARC software and Anthem series of sound processors is an example of an increasing number of solutions available to compensate loudspeakers for room interactions, and its impact in high-end audio should not be underestimated.

Having visited the CEDIA 2019 show, which happened in Denver a few days after RMAF, I could repeat the positive impressions with demonstrations of home cinema solutions, where room calibration tools and DSP processors from Dirac (https://www.dirac.com/room-calibration) and Trinnov Audio (http://www.trinnov-audio.com) were used with great impact, of course, carefully tuned by experts, and in demonstration rooms also featuring reference acoustic treatment. But what fascinates me with those solutions is the possibility that very soon they will be able to perform measurements and processing fully automatically.

As recently proven by the Apple HomePod speaker, this could soon become a reality. The Apple patented "spatial awareness" technology allows the HomePod to automatically adjust the sound to its surroundings. Being equipped with a complete microphone array, used for voice far-field capture, the HomePod is able to sense its location in a room, automatically adjusting the frequency response, directivity, and dispersion of its multiple drivers to compensate for being in a shelf, or siting on a kitchen table, delivering the best possible response for such a compact speaker.

This is something that I believe could be expanded in both consumer products and even professional audio solutions, leveraging—like Apple is doing—powerful modern processors, multiple sensors, microphone arrays, and DSP algorithms.

Of course, many applications will benefit from a combination of spatial awareness, and user adjusted calibration tools, like Sonos explored in its Trueplay speaker calibration tool. As long as we can rely on accurate microphone measuring systems to deliver consistent results, I believe we could soon see the audio industry benefiting from complete self-adjustable processing solutions for speaker calibration and room compensation.

Another exciting area that I believe will very soon contribute to significant improvements in this area is the use of room emulation techniques and virtual acoustics, as companies such as Sonarworks (https://www.sonarworks.com) and Waves (https://bit.ly/2XLRGWw) are exploring in this domain. Very exciting.

New Audio Scope Column!

Please join me in welcoming Michael Steffes to the pages of *audioXpress*. A name that will be familiar to *audioXpress* readers who have followed his many articles, reference application notes and his extensive work with Burr-Brown, Texas Instruments, and other IC vendors. We are certain that his column will be an important contribution for readers with a deep interest in audio electronics and development.

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.	59	KAB Electro-Acoustics	59
All Electronics Corp.	57	Linear Integrated Systems	37
Audio Precision, Inc.	2	Menlo Scientific, Ltd.	41
Avel Lindberg, Inc.	29	NXP Semiconductors	15
Benchmark Media Systems, Inc.	23	Parts Express International, Inc.	68
Celestion	9	Redco Audio	57
Hammond Manufacturing, Ltd.	3	Swans Speakers Kits	29
Hypex Electronics BV	25	Warkwyn Associates	31

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.

Contents

audioXpress

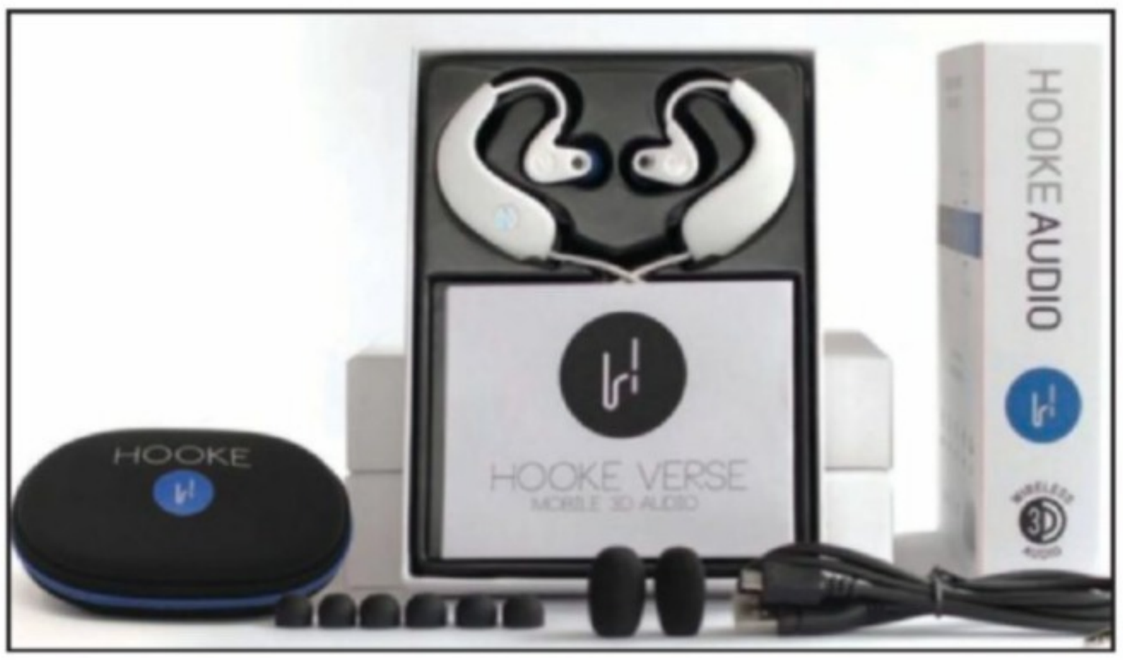

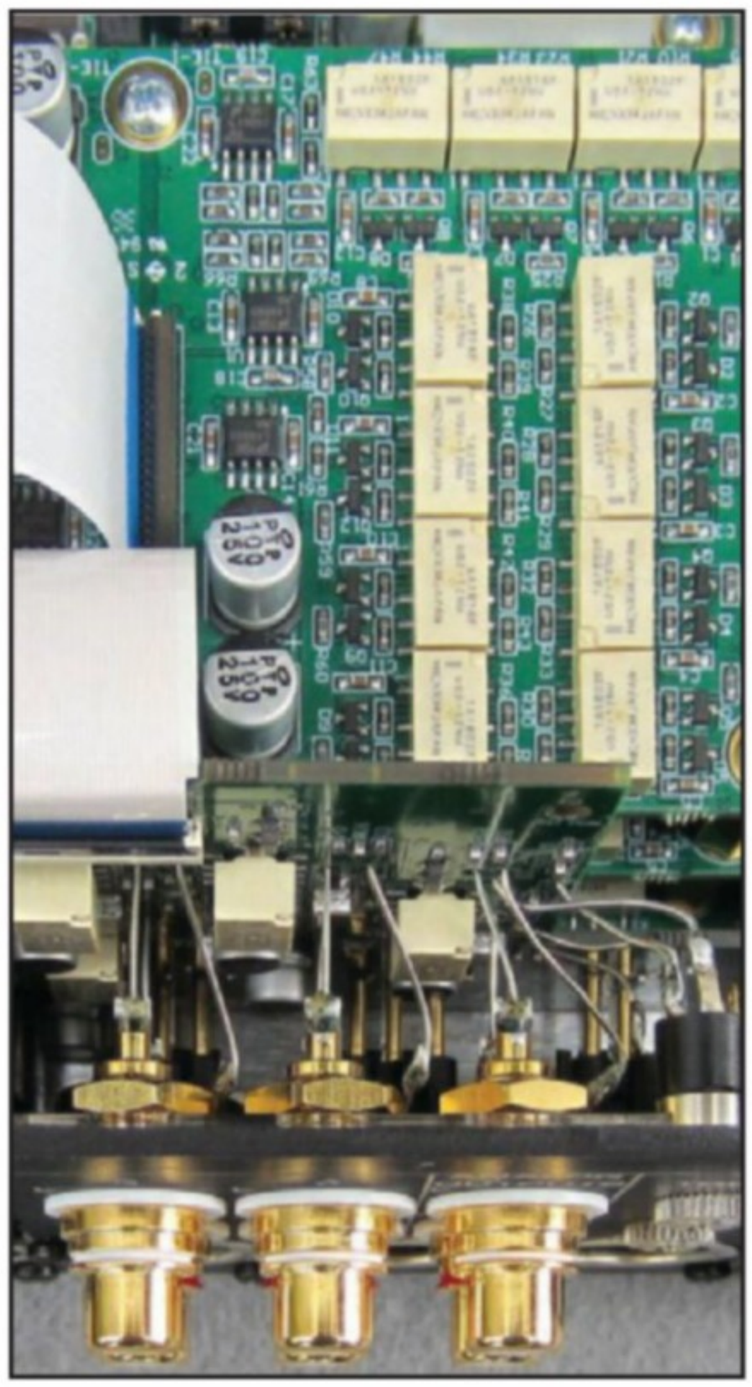

Features

8 Achieving Frictionless Lower-Power Application Connectivity with Data-Over-Sound
By Dr. Adib Mehrabi
This article explores how data-over-sound can be realized in embedded scenarios, and how Chirp's new technology, using the Arm Cortex-M4 and Cortex-M7 processors, can be implemented by third-party engineers in their connected products.

15 The NXP NXH3670 Enables Ideal Wireless Headset Experiences
Sponsored Content by NXP Semiconductors
Inspired by a wish to enhance gaming audio, NXP Semiconductor designed a single-chip, wireless audio streaming solution that delivers high performance and completely wireless audio communication.

16 Benchmark HPA4 Headpone/Line Amplifier
By Gary Galo
The Benchmark HPA4 is a reference headphone/line amplifier designed to deliver amplification of analog sources for audio enthusiasts, mixing and mastering engineers, and music aficionados in homes, recording studios, and mastering suites. Gary Galo explores it in detal.

26 Hooke Verse Bluetooth Binaural Recording Headphones
By Luke McCready
Luke McCready tests the Hooke Verse, part Bluetooth headphones and part wireless binaural recording device, to determine how it sounds and if it truly empowers creators-on-the-go with three-dimensional recording.

38 Two Mu Valve Topologies
By Christopher Paul
In this article, Christopher Paul tells the story of MU circuits, which he describes as being among the more interesting simple hollow-state voltage amplifiers.

50 A High-Voltage Delay for Tube Amplifiers The Sequel
By Jan Didden
In this article, Jan Didden offers a modern update on his original high-voltage delay design, which he built to complement tube amplifiers.

54 The Tenacious Bass-8 Powered Subwoofer
By Thomas Zarbo, Jr.
In last month's *audioXpress*, Thomas Zarbo, Jr. detailed his Bantams microspeaker system. This month, he shares the perfect accompaniment to that system, a subwoofer he calls the Tenacious Bass-8.

Volume 50 – No. 11

November 2019

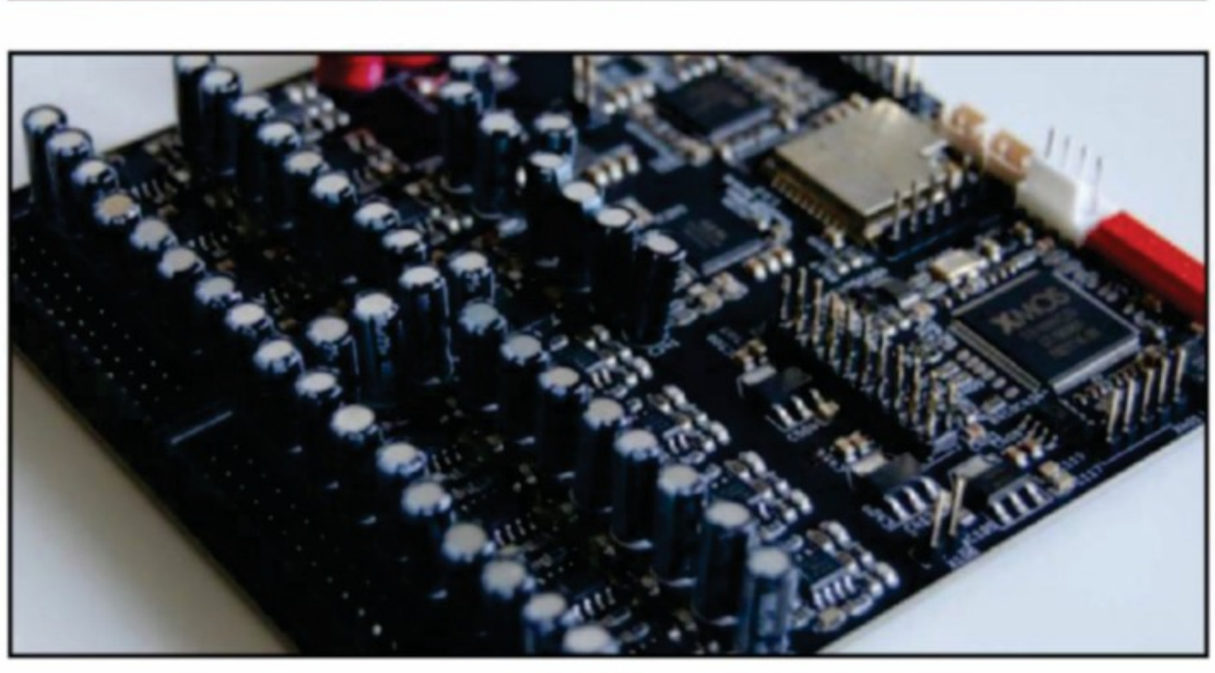

Columns

SOUND CONTROL

32 Is Your Neighborhood Acoustically Friendly?

By Richard Honeycutt

AUDIO SCOPE

36 Can Audio Experts Learn from High-Speed Amplifier Developments?

By Michael Steffes

QUESTIONS & ANSWERS

46 Introducing Aurora Open-Source DSP

An Interview with Raphael Knoop of Auverdion

By Shannon Becker

HOLLOW-STATE ELECTRONICS

62 The Hammond Growl

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

Achieving Frictionless Lower-Power Application Connectivity with Data-Over-Sound

This article explores how data-over-sound can be realized in embedded scenarios, and how Chirp's new technology, using the Arm Cortex-M4 and Cortex-M7 processors, can be implemented by third-party engineers in their IoT products.

By
Dr. Adib Mehrabi
(Head of Research, Chirp)

The rise of IoT devices in the home and workplace has created a world where data and connectivity are becoming increasingly complex. Today, we are seeing smart functionalities and connectivity being added to the plethora of devices already used for entertainment, utilities, and production. The Industrial Internet of Things (IIoT) presents an even more complex connectivity landscape, with tens, if not hundreds of devices, controllers, and sensors sitting side by side in the same factory or workplace. As a team of audio scientists and DSP engineers, Chirp's mission is to simplify machine-to-machine connectivity using sound. Recognizing the growing need for seamless, low-power and low-cost device-to-device communication in the emerging IoT landscape.

In early 2019, Chirp capitalized on the digital signal processing (DSP) capabilities of the Arm Cortex-M4 and Cortex-M7 processors, to engineer a software-defined data-over-sound solution that is robust and reliable without being resource intensive, enabling connectivity and application logic to reside on a single core. This article explores how data-over-sound can be realized in embedded scenarios, and how this technology can be implemented by third-party engineers in their IoT products.

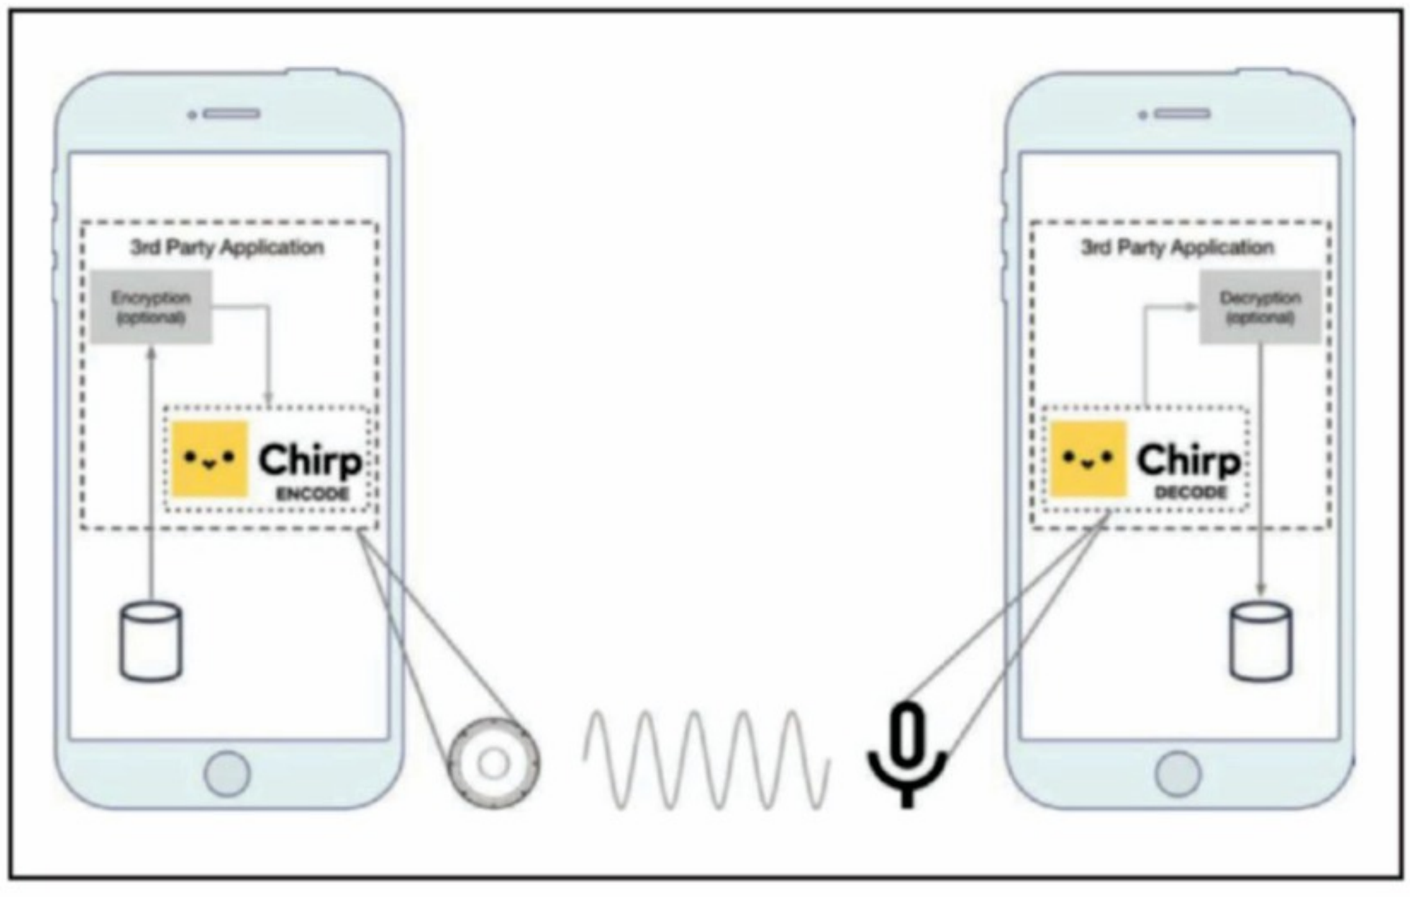

Figure 1: Encoding, transmission, and decoding of data using sound

What is Data-Over-Sound?

The basic idea of data-over-sound is no more complex than a traditional telephone modem. Data is encoded into an acoustic signal, which is then played through a medium (typically the air, although it could equally be a wired telephone line or VoIP stream), and received and demodulated by a "listening" device. The listening device then decodes the acoustic signal and returns the original data. This process is illustrated in **Figure 1**.

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.com

CELESTION

The Missing Piece of the Connectivity Puzzle

There already exists a plethora of connectivity solutions, including extremely short range (NFC and QR); short range, high bandwidth (Bluetooth and Wi-Fi); long range, low power (Sigfox and LoRa); and short range, low power (Zigbee). Each technology has certain advantages that make it more or less suitable for specific applications. Before we delve into the technicalities of Chirp's solution, we'll cover some of the properties of data-over-sound that make it a cost-effective and versatile medium for application connectivity.

As any technology provider in the IoT arena will appreciate, application-specific connectivity needs are not always as simple as choosing a technology based on cost, data rate, and range. Things can get a lot more complicated when one considers the finer details of device support, backward compatibility, user experience, and security. Further detail is introduced by the technical specifications of the protocol, such as the nature of the transfer medium (RF, optical, acoustic), number of required channels, security concerns, or whether two-way and one-to-many communication is needed.

An overview of these considerations is given in **Table 1**. This is by no means an exhaustive list, and each use case will bring its own specific requirements. Nonetheless, data-over-sound has certain advantages that make it particularly attractive for a number of use cases that are suited to smart home, smart workplace, and IoT devices, including:

- Device interoperability: The simple hardware requirement of a speaker and/ or microphone make data-over-sound arguably the most wide-reaching wireless communication technology in terms of device compatibility. Mobile phones, voice controlled devices, and any device with an alarm speaker are able to communicate using data-over-sound. This includes many legacy devices. Data can also be transmitted over media channels, such as radio and TV, and over existing telephone lines.
- Frictionless UX: Data-over-sound requires no pairing or configuration, making data transfer as simple as pressing a button.
- Physically bounded: Because sound waves respect room boundaries, particularly in the near-ultrasonic range commonly used in data-over-sound, transmissions do not pass between neighboring rooms. This means that it can be used for detecting the presence of a device with room-level granularity.
- Low or no bill-of-materials (BOM): Data-over-sound adds nothing to a device BOM if the specification already includes audio I/O, which is becoming increasingly commonplace. If audio I/O is only required for data-over-sound

	Chirp	QR	NFC	Bluetooth	Wi-Fi	Li-Fi	ZigBee 802.15.4	LoRa	Sigfox
Two-way communication	•			•	•	•	•	•	•
One-to-many broadcast	•					•	•		
Non-line-of-sight transmission	•			•	•		•	•	•
Works in RF-restricted environments	•	•	•			•			
Zero setup/pairing/configuragtion	•	•	•						
Available to applications by default	•	•							
Low-power operation	•	•	•	•			•	•	•
Can transmit with sub-$2 electronics	•	•	•	•			•		
Can receive with sub-$2 electronics	•			•			•		
Wireless broadcasts confined to room boundaries	•	•	•						
Transmit over ranges > 10 m	•	•		•	•	•	•	•	•
Can limit the transmission range to < 1 m	•	•	•	•					
Supported by typical mobile devices	•	•	•	•	•				
Supported by low-end mobile devices	•	•		•	•				
Typical usable maximum data rate	1 kb/s	3 kb	424 kb/s	25 mb/s	70 mb/s	1 gb/s	250 kb/s	50 kb/s	100 b/s
Typical maximum range	100 m	10:1	20 cm	100 m	50 m	10 m	100 m	10 km	40 km

Table 1: Feature comparison for Chirp vs. alternative connectivity technologies

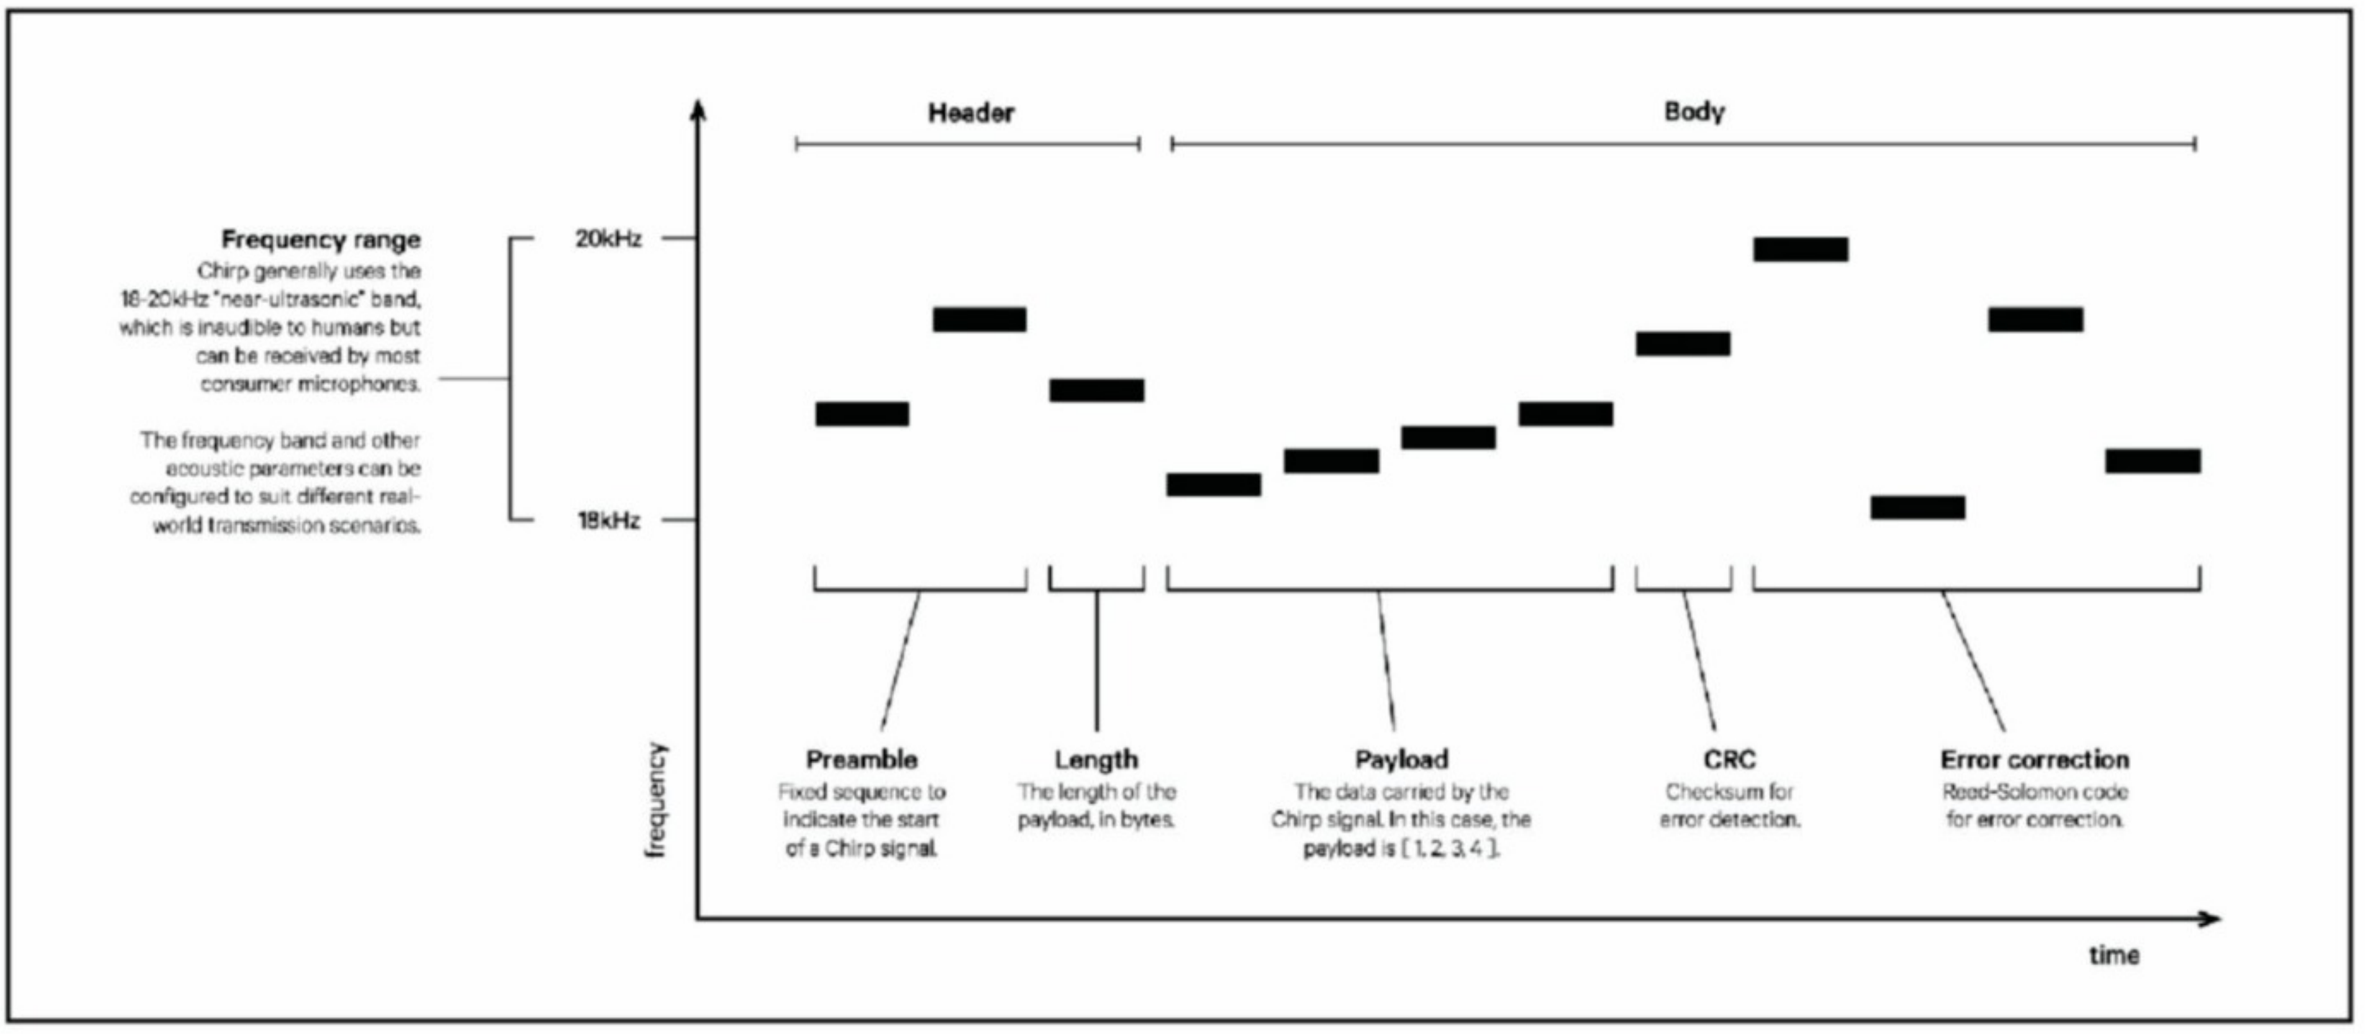

Figure 2: Anatomy of a Chirp signal

functionality, the BOM is less than $2 for sender and receiver components.

- Zero power: The advent of "wake-on-sound" MEMS microphones (e.g., the Vesper VM1010) enables devices to communicate using data-over-sound, while draining virtually no battery power in between communications (< 10 µA).

Simple to Make, Difficult to Master

So how does it work? Chirp uses Frequency Shift Keying (FSK) for its modulation scheme, due to its robustness to the multipath propagation present in real-world acoustics, compared to schemes such as Phase Shift Keying (PSK) or Amplitude Shift Keying (ASK). For spectral efficiency, Chirp uses an M-ary FSK scheme, encoding input symbols as one of M unique frequencies. Each symbol is modulated by an amplitude envelope to prevent discontinuities, with a guard interval between symbols to reduce the impact of reflections and reverberation at the decoding stage.

As with other data transmission protocols, the message can include specific sequences to indicate the start and expected length of a message and to establish timing and synchronization, as well as additional bytes for Reed-Solomon Forward Error Correction (FEC) error detection and CRC error detection, as shown in **Figure 2**. Encoding a Chirp signal in this manner is a lightweight process in terms of memory and CPU, requiring only the generation of error correction symbols and sinusoidal oscillator synthesis.

The encoding process and resultant audio signal is very well understood in the worlds of audio signal processing and communications systems. Given an ideal communication channel (with minimal noise and distortion), the decoding process is as simple as identifying the peak frequencies in the audio stream, and mapping them onto symbols to recover the original message. The robustness of this method using near-ideal channels is probably best demonstrated by the ubiquity of audio FSK systems in the form of the famous telephone modem.

However, in the world of acoustic transmission over the air, things are not so simple. Chirp's technology is typically used in challenging acoustic environments, with potentially high levels of background noise (e.g., from crowded city streets, busy cafes, or noisy domestic environments containing music, television and people talking over each other). In addition to background interference, we know that audio is particularly susceptible to distortion from the acoustic environment in terms of reverberation, bringing with it issues of resonance, phase cancellation, and comb filtering effects.

If these factors alone are not challenging enough, Chirp prides itself on the ubiquity of its technology—potentially supporting any device with a microphone and/or loudspeaker. Chirp's SDKs offer support for most platforms, including iOS, Android, Windows, macOS, Web (JS), and embedded boards running Arm processors, along with a dedicated toolkit for Amazon Alexa.

Each device has unique operating characteristics in the audio signal chain, including the frequency and phase responses of the transducers, ADC, DAC, and physical layout (e.g., casing), not to mention the onboard proprietary DSP found on many modern devices. All of these parts have the potential to

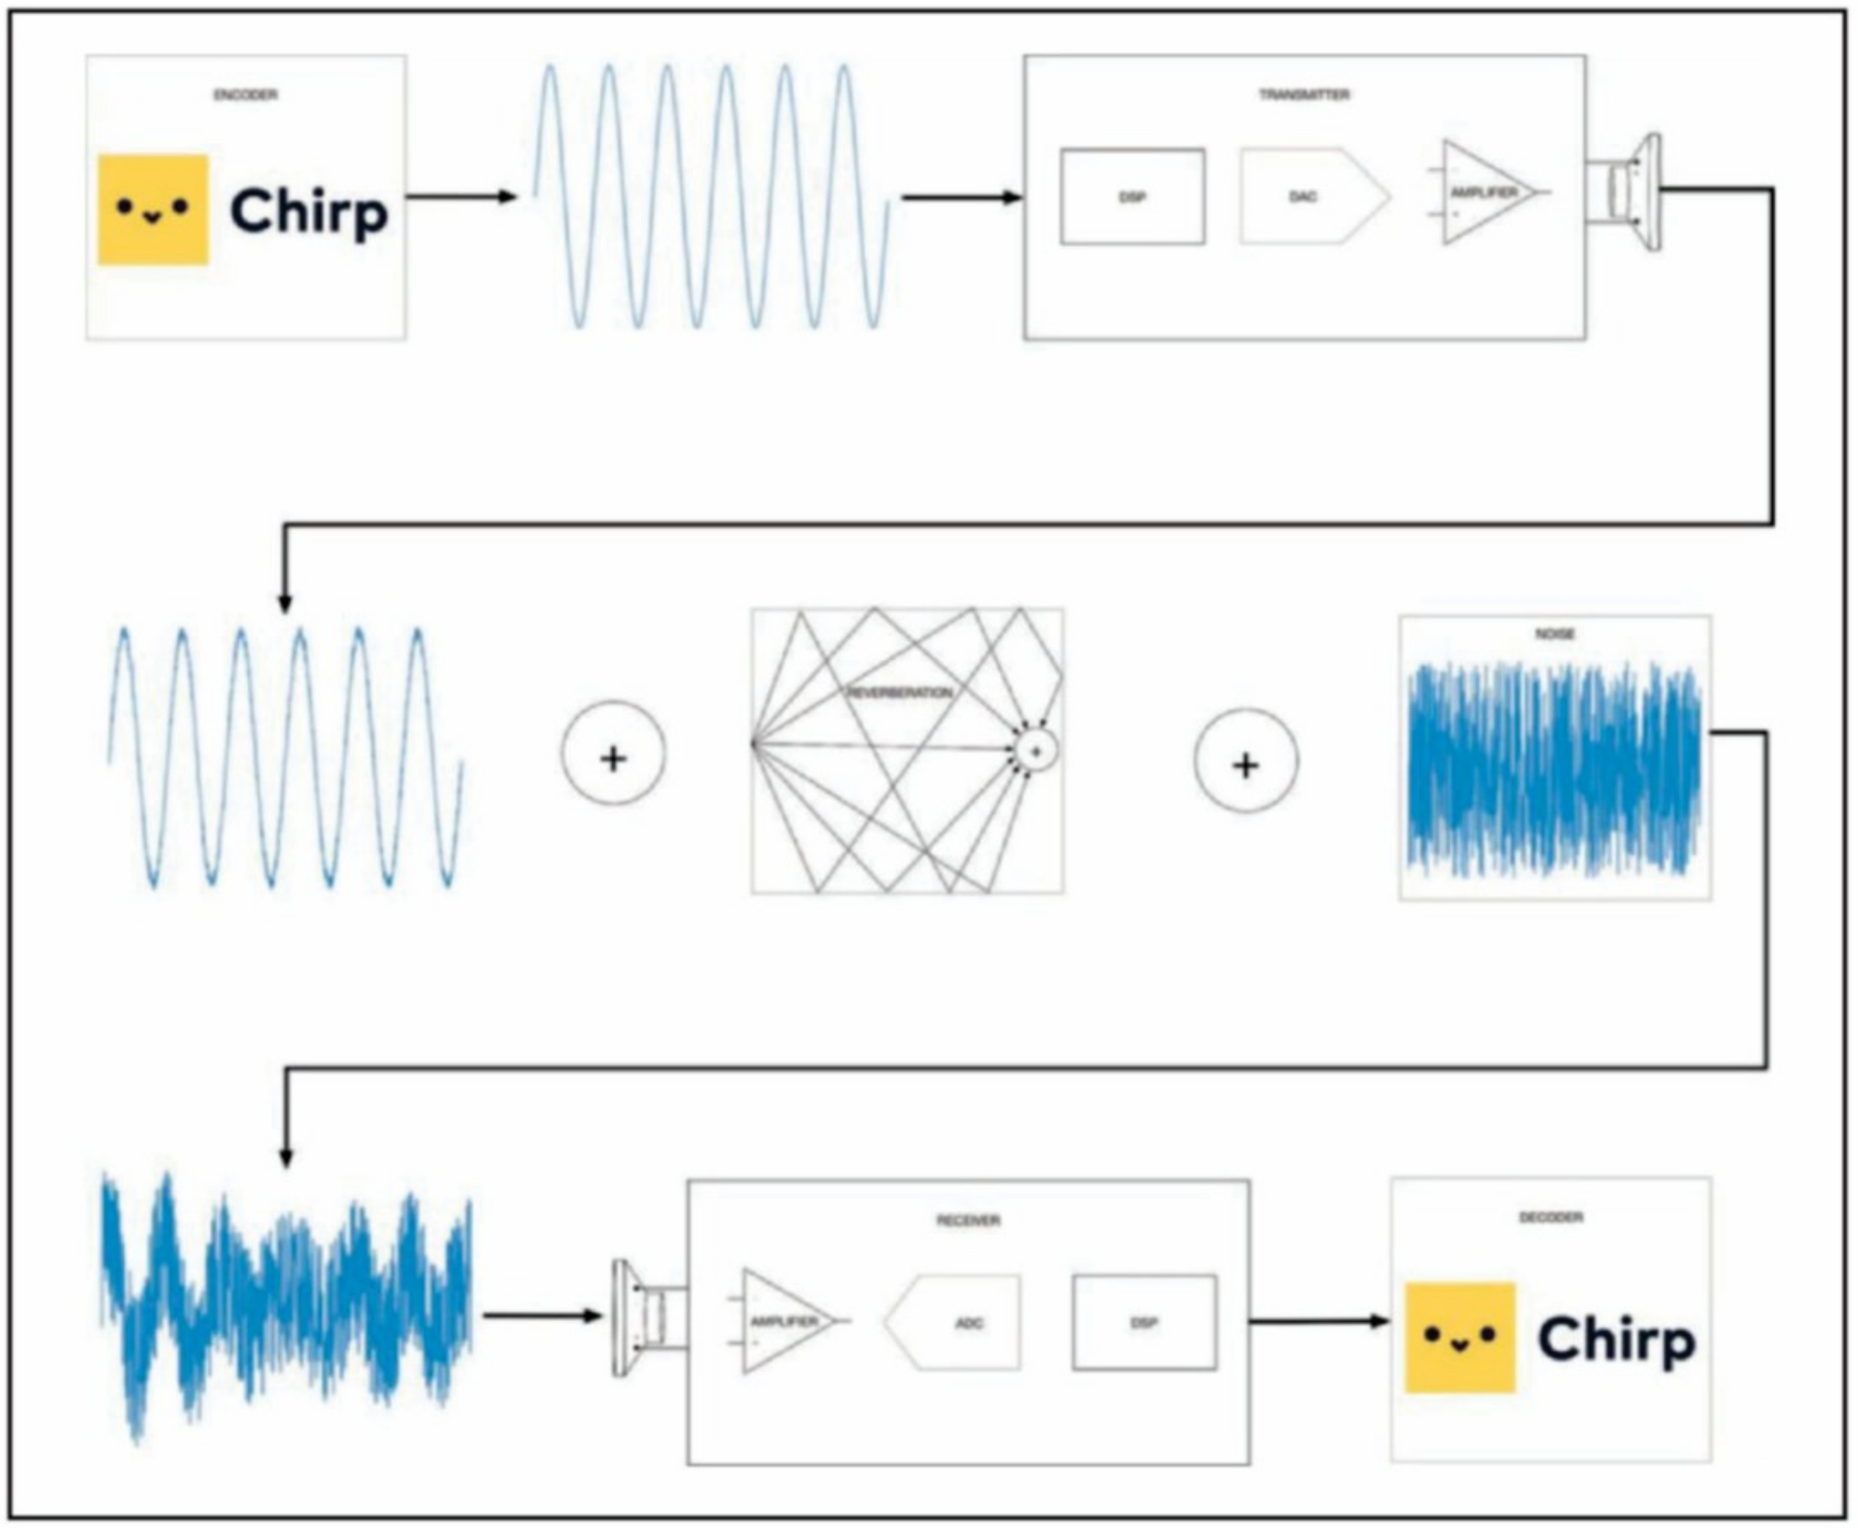

Figure 3: Sources of noise, distortion, and interference in the signal chain

interfere with and distort the signal during its transmission from sender to receiver, as shown in **Figure 3**.

At Chirp, seven years of research focus and the resultant IP has enabled us to build robust audio-based communications systems that perform reliably in a wide range of acoustic environments. Of the existing data-over-sound solutions on the market, the main differentiator for Chirp is this reliability and robustness of the internal decoding technology.

Our proprietary decoding system is able to overcome many of the challenges typically associated with data-over-sound systems. As one might expect, the bulk of Chirp's processing requirements come from the decoder. This is where high-performance DSP power is required, and where DSP-specific features of the onboard processor become critically important.

So how reliable is Chirp? The short answer is that it depends on the system requirements. There is an inherent trade-off between range, data rate, and reliability (see **Figure 4**). All of Chirp's standard profiles are tested to perform at >99% reliability under normal operating conditions. This level of performance can be maintained under particularly challenging acoustic conditions (e.g., with low signal-to-noise ratios and sub-optimal transducers); however, it is ultimately a function of range and data rate. One can expect reliable transmission from 1 cm to 100 m, subject to data rate requirements. This has been demonstrated and stress-tested in acoustically extreme environments, including nuclear power stations with upward of 100 dB(A) of background noise, across ranges exceeding 100 m.

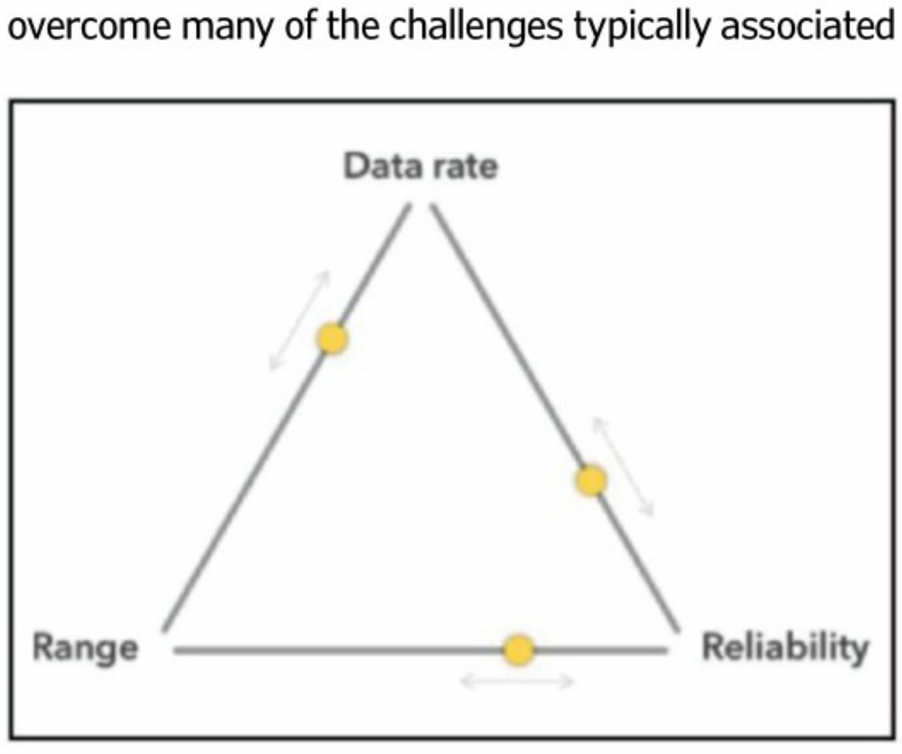

Figure 4: The trade-off between data rate, range, and reliability

Applications and Use Cases

In response to the connectivity requirements discussed in the previous sections, certain

application areas have emerged as the drivers of data-over-sound in IoT contexts, primarily in scenarios in which it significantly reduces friction or human effort (see **Figure 5**). We will now outline four such areas, where data-over-sound is the only data communication solution that meets all of the connectivity requirements.

IoT Device Provisioning: Configuring a new smart device with network credentials and functional configuration remains a disproportionately complex task, particularly for headless devices. Although technologies such as WPS have emerged to address this challenge, these typically require support on the wireless access point or wider infrastructure or, in some cases, physical access to the AP itself. Typically this use case requires a solution that offers seamless and rapid setup, with no pairing/configuring, one-to-many transmission (including non-line-of-sight and >10 m range), and no additional hardware requirements on a router. Chirp offers an approach to provisioning that is offline, locally-bounded, and does not require any infrastructure modifications. Credentials are encoded as audio, optionally layered with cryptography for secure scenarios, and broadcast over-the-air to nearby smart devices. The BOM needed to receive and decode credentials can be as minimal as an Arm Cortex-M4 processor plus a digital MEMS microphone.

Proximity and Presence Detection: This use case requires a one-to-many connectivity solution that respects room boundaries and has universal support for devices. A key property of acoustic signals is that their propagation respects room boundaries, particularly in the near-ultrasonic range. As a consequence, there is a growing uptake for near-ultrasonic acoustic beacons for presence sensing at room-level granularity. Devices can send out sonar-like beacon packets at regular intervals, enabling them to be discovered by nearby peers. Chirp's implementation uses multi-channel encoding to support multiple devices within hearing range, with carrier sensing to detect free channels.

Two-Way Acoustic NFC: Data-over-sound supports near-field communication scenarios while offering a two-way full-duplex mode of communication, thus addressing a critical limitation of NFC. Enabling two-way data exchange means that devices can perform challenge-and-response dialogues—for example, Diffie-Hellman key exchange for secure financial transactions, or securely sending un-spoofable receipts to a merchant or buyer. The low-power requirements and low cost of audio components, combined with the fact that data-over-sound is software-defined thus requires no additional hardware, means that acoustic NFC is a powerful emerging approach to peer-to-peer payments and minimal-cost POS devices.

Telemetry in RF-Restricted Environments: In many sensitive environments, RF-based communications are prohibited due to the risk of sparks or interference with equipment that pre-dates RF regulation. In addition, communications systems in these environments can require non-line-of-sight two-way communication, with up to 100 m range. Chirp overcomes these limitations, enablng the industrial IoT to harness the benefits of wireless communication without limitations of the EM-spectrum.

Data-Over-Sound on Arm Cortex-M Processors

Chirp has recently developed a data-over-sound SDK for the Arm Cortex-M4 and Cortex-M7 based MCUs. The DSP features on these chips allow for high sample rate audio processing without the need for hardware acceleration, enabling the encoding and decoding of ultrasonic chirps. This places the technology in the hands of device manufacturers

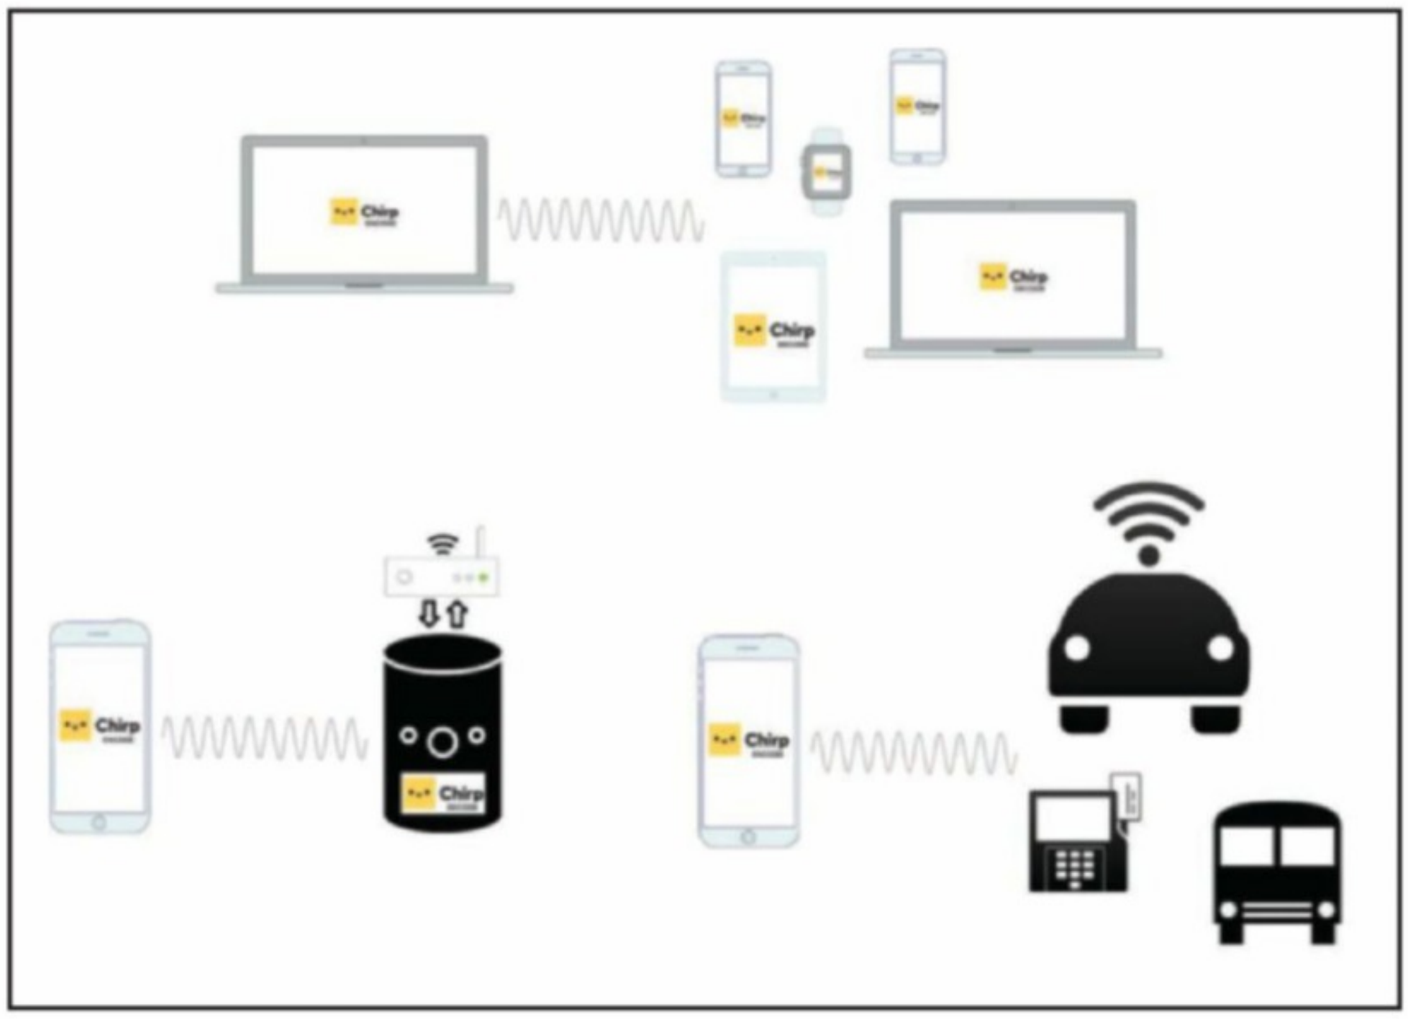

Figure 5: Chirp provides a simple, effective communication system for broadcasting to multiple devices, provisioning smart devices, and conducting wireless transactions.

About the Author

Dr Adib Mehrabi is an accomplished research scientist with a background spanning audio engineering, acoustics, DSP, machine learning, and speech processing. He holds a BSc (Hons) in Audio Technology from UWE, Bristol and a PhD in Computer Science from Queen Mary University, London. Adib currently heads up the research efforts at Chirp, where his work focuses on building DSP and AI-based systems for audio processing.

and developers who wish to take advantage of a frictionless, low-cost solution for device provisioning, zero-power monitoring, or RF-free telemetry in general.

The suite of Chirp C SDKs are compact but fully featured, allowing developers to use Chirp's technology even on low-power, memory-constrained devices. The Chirp Arm SDK is optimized for the Cortex-M processors, with support for both the Cortex-M4 and Cortex-M7. It exists as a static library (approximately 740 KiB in size) with a straightforward API for sending and receiving data via Chirp signals.

The SDK implements the physical and data link layers of the data-over-sound OSI stack. Chirp is software-defined, meaning that any additional layers can be easily implemented on top of the SDK if features such as encryption, application-specific protocols, or end-to-end transmission/QoS are required.

Chirp can be incorporated into the software stack for existing projects by integrating the library's object (.a) and header (.h) files. API calls are available for the following functions:

- Initialization and termination, to allocate and free the main structures of the SDK
- Audio processing interaction, in which arrays of audio samples are exchanged between the audio I/O drivers and the Chirp SDK
- Payload handling, including creation and transmission of data
- Parameter getters and setters for properties such as software volume, SDK states, and internal parameters (e.g., audio sample rate)
- Configuration of user profile and audio transmission properties, for application specific communication protocols
- Debugging helpers, for converting payloads and error codes into human-readable strings

The protocol definitions are handled via a cloud service and associated with specific users/ clients. Protocols are often updated and in some cases users will require bespoke protocols for specific applications. This provides a simple way to update the protocol without having to change any of the code implementations.

Decoding messages requires the processing of relatively large amounts of audio data in a short space of time. Using the CMSIS open-source DSP library, the Chirp Arm SDK maximizes the performance of the Cortex-M processors by using fully optimized instruction sets. In particular, the SIMD instructions allow the Chirp decoder to work on large audio buffers in real time without incurring CPU overruns.

In order to monitor and manage the memory usage, Chirp has developed the Chirp Memory Manager, an open-source library for tracking memory usage in C. This provides a single entry point for allocating and freeing dynamic memory. Using the memory manager, it is possible to track the dynamic allocations and query the current allocated memory at any time in the life cycle of the program.

If there is sufficient headroom in the memory, additional statistics can be calculated and stored during the lifecycle of the program, logging information regarding every allocation and free (relative time, file, function, type allocated, amount of the allocation).

Using the CMSIS library and Chirp Memory Manager, the memory and CPU requirements of the Chirp Arm SDK have been reduced to levels that not only perform within the hardware limitations, but leave ample headroom for additional processes and sophisticated application logic to run alongside Chirp.

The Future of IoT with Data-Over-Sound

By expanding the suite of Chirp SDKs to include the Arm Cortex-M4 and Cortex-M7 processors, Chirp has substantially lowered the barrier to incorporating data-over-sound into key IoT applications.

One of the major threats to the mass adoption of IoT technology is the age-old conflict between performance and price. We are entering an age where people expect even the cheapest smart devices to be able to process real-time audio and perform intelligent digital signal processing (DSP) on the edge device. The Cortex-M series of processors addresses this need by providing an extremely cost-effective solution that is capable of high-performance real-time audio DSP.

This solution presents the fundamental concepts and benefits of data-over-sound connectivity, and explores key application areas within the IoT, including provisioning smart devices and facilitating secure near-field communication in low-cost, low-power scenarios.

Moreover, the status of data-over-sound as a pairing-free, one-to-many medium means it addresses some of the pain points of longstanding alternatives such as Bluetooth and Wi-Fi, presenting an attractive general purpose solution for frictionless data transmission.

To learn more about more about Chirp's SDK for Arm's Cortex-M processor series, go to https://blog.chirp.io/arm-whitepaper-2.

The NXP NXH3670 Enables Ideal Wireless Headset Experiences

By
NXP Semiconductors

Whether it's for gamers trying to chat with friends during an online match or fitness gurus in need of an upbeat playlist during their morning workouts, the benefits of completely wireless headphones are obvious. By eliminating the mangled wires and connecting the audio of consumer devices directly to the ear, you not only enhance the comfort and convenience of the headphone experience but equip the user with a much-desired level of freedom.

Yet in developing wireless audio, there are a multitude of concerns. Battery life, latency, and effectively maintaining audio connection are all factors that may make or break the user experience. Audio drop off, be it in a remote conference call or in the middle of a videogame deathmatch, could be disastrous. Therefore, the question of wireless audio is not a question of why, but of how. How exactly can it be pulled off?

NXP's NXH3670 has stepped up to answer that questions. Inspired by a wish to enhance gaming audio, this single-chip, wireless audio streaming solution delivers high performance, completely wireless audio communication. The NXH3670 Gaming and Communications solution provides robust, high-quality, and ultra-low latency audio and voice streaming with its optimized architecture, embedded DSP, and hardware support.

Flawless Audio

In order to avoid audio synchronization issues, a gaming headset requires extremely low latency. In the midst of a high-stakes virtual deathmatch, a split-second delay may give the upper hand to your rival. With an ultra-low latency of less than 20ms, the NXH3670 ensures that your audio is streaming to you in real time without any disruptions. The NXH3670 is also equipped with a robust interference avoidance scheme, which ensures a seamless connection even in the presence of Bluetooth and Wi-Fi interference.

Industry-Leading Battery Lifetime

While other wireless headsets may falter with high power consumption—constantly needing a recharge—the NXH3670 enables an ultra-long battery life. This is made possible by NXP's efficient streaming protocol, which reduces the power consumption to 8.5 mW, the lowest active energy consumption in the market. This allows you to enjoy our seamless, continuous audio for hours on end.

Comfortable, Lightweight Design

The ultra-low power consumption of NxH3670 allows for comfortable and stylish headset designs which not only equip the user with the utmost audio quality, but with a sleek, lightweight, and incredibly comfortable look and feel—making it all the better to wear.

The NXH3670 keeps the joy of wireless audio alive, and does so with unparalleled comfort, audio latency, and battery life. While primarily designed for headsets, the NXH3670 also can enable wireless audio streaming for household and office devices like soundbars and wireless speakers. If you are interested in learning more about the NXP's NXH3670, and how it can empower your designs, please visit www.nxp.com/NXH3670 .

Benchmark HPA4 Headphone/Line Amplifier

In this article, Gary Galo reviews the Benchmark HPA4, a reference headphone/line amplifier designed to deliver amplification of analog sources for audio enthusiasts, mixing and mastering engineers, and music aficionados in homes, recording studios, and mastering suites. Resulting from Benchmark's collaboration with THX, the HPA4 features a THX-888 headphone amplifier for an output signal with high voltage, current, and and power complemented by low distortion, noise, and crosstalk.

Photo 1: This is Benchmark's HPA4 Headphone Amplifier/Line Amplifier. This new product combines the LA4 Line Amplifier with a state-of-the-art headphone amp.

By
Gary Galo
(United States)

The HPA4 is the latest product from Benchmark Media Systems, an upstate New York audio manufacturer. Benchmark calls it a headphone amplifier, but it's really two products in one, combining the LA4 Line Amplifier introduced in 2018, with a headphone amplifier (see **Photo 1**). It is also based on the same THX Achromatic Audio Amplifier technology used in Benchmark's acclaimed AHB2 power amplifier. The signal path in the HPA4 is entirely analog, with four stereo, line-level inputs—two fully balanced XLR and two unbalanced RCA (see **Photo 2**). Outputs include stereo balanced XLR, a summed-mono balanced XLR, and stereo unbalanced RCA. The HPA4 is designed for full, professional operating levels using the XLR I/O, up to +28 dB, or 19.5 VRMS, in and out.

Overview

Nearly all of the control functions are accessed through the front-panel touchscreen (see **Photo 3**). The only visible control is a large volume knob on the left side. Features include separate volume and balance adjustments for the headphone and line outputs and individual gain adjustment for each input, up to ±10 dB. The gain adjustments are handy for matching playback levels between various source components. The Settings menu allows customizing various functions, including input identification and screen brightness. There are also a variety of powering options, including 12 V triggering. Benchmark's clear, well-written manual explains everything in detail, and is available as a download on its website.

Benchmark Media Systems, Inc.
203 East Hampton Place, Suite 2
Syracuse, NY 13206
800-262-4675
www.benchmarkmedia.com
Price: $3,099 US
Optional Remote Control: $100 US

The HPA4 is designed to be used with the same optional remote control as Benchmark's digital-to-analog converters (DACs). Benchmark has a firmware upgrade for its DAC2 and DAC3 converters that will enable the DACs to be put in a Compatibility Mode for use with the HPA4. The new firmware is V. 2.0 and consists of an 8-pin DIP module that plugs into an empty socket on the DAC's main PC board. The firmware upgrade costs $59 US and can be found on the Support/Parts page of the website. The Compatibility Mode avoids duplicating DAC and HPA4 functions. Input Selection and Volume now control the HPA4. The old firmware can be re-activated by moving a jumper on the PC board.

Relay Volume Control

Benchmark's website offers a good overview of the innovative technology used in the HPA4, and a look "under the hood" reveals a clean and sophisticated execution (see **Photo 4**). The line amplification and volume control circuitry are on the upper, stacked PC boards, toward the rear panel. Each board contains five Texas Instruments LME49860 dual op-amps, the same low-distortion, low-noise, high-output current device used in the DAC2 and DAC3. The balanced inputs are terminated differentially, to achieve the highest-possible common-mode rejection. Benchmark uses a double-differential system, which is essentially differential in and differential out, so internally the amplification circuitry is also fully-balanced.

Volume control is accomplished with a unique system of precision relays with gold-plated contacts, and low-impedance resistor networks made with 0.1% tolerance precision resistors. To keep the noise levels as low as possible, Benchmark uses thin-film resistors with low temperature coefficients and conservative power ratings. The relay/resistor networks are actively buffered—the buffering provides high input impedances and low output impedances. The volume adjustments are accomplished with a hybrid of attenuators inside the feedback loops of the active devices, and passive inter-stage attenuation.

There are four independent, fully-balanced, 256-step volume controls—two for the line outputs and two for the headphone outputs. None of this would be possible without the digital control board, which is the bottom rear PC board partially visible in Photo 4. The heart of this board is a Xilinx Spartan 6-series Field-Programmable Gate Array (FPGA) chip. The volume control knob operates an optical rotary encoder that turns continuously in either direction. The encoder includes an acceleration feature, which makes it easy to make large changes in volume very quickly, while allowing precise adjustments when the knob is rotated more slowly. When rotated slowly, each detent results in a 0.5 dB change, so the 256 steps cover a 128 dB range. The FPGA chip also provides precisely timed, make-before-break relay connections, eliminating the pops, clicks, and zipper noises often associated with relay switching, no matter how fast the volume control is rotated.

It's no secret that potentiometers in the signal path can add small amounts of noise and distortion, and with current measurement technology (e.g., the Audio Precision AP2722 system used by Benchmark) that noise and distortion can often be measured. When you're dealing with equipment capable

Photo 2: The rear panel of the HPA4 includes two stereo balanced XLR inputs and two stereo unbalanced RCA inputs. The fully-balanced outputs include stereo and summed-mono XLR connections, along with stereo unbalanced RCA connectors. The 12 V trigger connections can be used as either inputs or outputs, and the power connector accepts standard IEC power cords.

Photo 3: The attractive touchscreen display controls most of the functions of the HPA4. Channel balance can be separately adjusted for the headphone and line output groups. The gear symbol in the lower right selects the Settings menu. Volume can be individually trimmed for each of the four inputs, and the user can give each input a custom name.

of noise and distortion in the neighborhood of -130 dB, everything matters, and with a signal path as transparent as a DAC3/HPA4/AHB2 system, signal degradations masked by less capable equipment now become audible. Adding a balance potentiometer will make things even worse.

Even the best stereo potentiometers rarely have channel-to-channel tracking better than 1 dB across their rotation. Integrated circuit (IC) volume controls can be precise, but none are sonically transparent enough for perfectionist applications. The accuracy issue can be addressed with stepped attenuators and precision resistors, but I haven't seen any beyond 64 positions, making fine adjustments impossible. No stepped attenuator can offer more than a fraction of Benchmark's 256 positions, in 0.5 dB increments.

The relay system makes it possible to adjust channel balance and the gain of individual inputs, without any additional controls or circuitry. The FPGA chip simply switches the volume attenuator relays as needed to achieve these adjustments. Input switching and output muting are also done with relays, all controlled by the FPGA (see **Photo 5**). Benchmark has an application note on its website offering further details on the relay system, and how it's been implemented with the LME49860 op-amps. The complete app note provides additional technical information and is well worth reading.

I measured the review sample using my HP/Agilent 34401A's decibel function and found the channel-to-channel tracking to be within 0.02 dB over a 0 dB to -40 dB range (I didn't see the need to continue measuring below -40 dB). This certainly lives up to Benchmark's claim of "near perfect" tracking. I also found the accuracy of the 0.5 dB increments to be within 0.02 dB over the same range. This is extraordinary accuracy!

Photo 4: This is an inside view of the HPA4. The stacked left and right line amplifier circuit boards are visible toward the rear panel. Each line amp board contains active circuitry based on LME49860 op-amps and volume switching using gold-contact relays and 0.1% precision resistors. The THX-888 headphone amplifier board is on the bottom, toward the front. The digital control board is partially hidden by the analog boards, and the switching-mode power supply module is on the left.

THX AAA Headphone Amplifier

THX's Achromatic Audio Amplifier topology uses feed-forward error correction to reduce distortion by 20 dB to 40 dB. The crossover distortion produced by conventional Class AB amplifiers is virtually eliminated. With feed-forward error correction, amplifier errors are extracted from the output audio signal, inverted in polarity and fed to the output to cancel the errors. A small, error-correction amplifier is needed to buffer the error signal, and any amplifier employing error correction can be no better than the error-correction amplifier itself. This is the technology used in the AHB2 power amplifier, and THX has recently applied this concept to a line of headphone amplifiers.

THX lists seven AAA headphone amplifiers on its website, covering low-cost, mobile applications to high-end, high-power audiophile-grade designs. Benchmark wanted a headphone amplifier with the same level of performance as the AHB2 power amp, so nothing short of the top-of-the-line would suffice. That top-of-the-line amplifier is the THX AAA888. The 888 is not a module, but rather a circuit topology that manufacturers can incorporate into their own PCB layouts. Benchmark actually participated in the final development of the 888 amplifier, which is a dual-mono design using a pair of Texas Instruments (TI) OPA564 power IC op-amps, one per channel.

The HPA4 includes both a 1/4" phone and an XLR4 connector for the headphone outputs. John Siau, Benchmark's VP and Director of Engineering, notes that the main advantage of the XLR4 connector is lower contact resistance than a 1/4" phone tip-ring-sleeve (TRS) interface. If a TRS connector must be used, it should have clean, gold-plated contacts. The outputs of each OPA564 drive the positive terminals of the headphone connectors.

Driving headphone voice coils differentially has

gained some popularity in certain audio circles, but Benchmark's view is that this has no performance advantage because a transducer can't tell the difference between balanced and single-ended drive. The quality of the ground connection is crucial, and the XLR4 connector eliminates the shared ground connection between the two channels. Benchmark's favorite headphone has been the Sennheiser HD-650, which you can purchase on Benchmark's website. These headphones are supplied with a removable cord, terminated in a 1/4" TRS phone plug. But, Sennheiser also makes an upgrade cable, the CH 650 S, which has a 4-pin XLR connector.

The OPA564 comes in a 20-pin PowerPAD package and is capable of delivering 1.5 A of output current. PowerPAD is a package developed by TI that provides efficient thermal transfer to a solder pad on the PCB. An application report is available on the TI website. Benchmark rates the headphone amplifier at 6 W into 16 Ω, so this really is a small power amplifier rather than one of the high-current op-amps more typically used in headphone amplifiers.

The headphone amp in the HPA4 has an output impedance of near-zero ohms. This is because there are no series resistors at the outputs of the OPA564 amplifiers, which would seriously degrade damping factor. Series resistors also change the headphone's frequency response because the impedance of a headphone, like nearly any transducer, changes with frequency. The THX AAA technology keeps the distortion levels extremely low even when driving difficult, low-impedance loads. The headphone amp includes protection against DC offset, clipping, short-circuits, as well as excessive temperature and output current. Any faults will trigger the output muting relays, protecting your headphones and your ears.

TI also specifies a full-power bandwidth of 1.5 MHz and a slew rate of 40 V/µSec. The 888 also uses a pair of OPA1612A dual op-amps as the error-correction amplifiers. The OPA1612A's rated THD of 0.000015% makes it ideal for this purpose. By itself, TI specifies THD for the OPA564 as 0.003%, or -90 dB. THX's AAA feed-forward error correction reduces the headphone amplifier THD to 0.00006%, or -125 dB. For those who are used to thinking of THD in percent, the formula for conversion to decibels is: dB = 20Log(%/100).

There's an online calculator for conversion in either direction on the Sengpiel Audio website. Benchmark's Application Note "Interpreting THD Measurements" is also worth reading.

Benchmark is a firm believer in switching-mode power supplies because they eliminate the hum fields caused by conventional, 60 Hz power transformers. Switching supplies are necessary to achieve signal-to-noise ratios (SNR) that realize the full potential of 24-bit DACs. Switching-mode power supplies (SMPS) have a bad reputation in some audio circles but, like any other design concept, execution is everything.

Benchmark achieves superior audio performance by using a power supply that operates at a switching frequency of 800 kHz—much higher than many switching-mode supplies. The HPA4 uses the same power supply module that Benchmark has used successfully in the DAC2 and DAC3 converters. Equally important is power supply filtering and regulation. For Benchmark, these are long-solved problems, as the superior performance of its DACs shows. The lack of radiated hum fields from Benchmark's products means that the HPA4, DAC, and AHB2 power amp can be stacked one on top of another with no degradation in noise performance.

The construction quality of Benchmark products has always been very high, and the HPA4 lives up to their reputation. This 8-pound chunk exudes quality as soon as you lift it out of the box. The front and side panels are anodized, machined aluminum, and the rear panel and bottom plate are 0.125" (3.18 mm) anodized aluminum stock. The front panel is 0.25" (6.35 mm) thick, matching the appearance of their DACs and power amp; the side panels are 0.84" thick (21.3 mm), which contribute to the robust feel, and also provide effective heat dissipation.

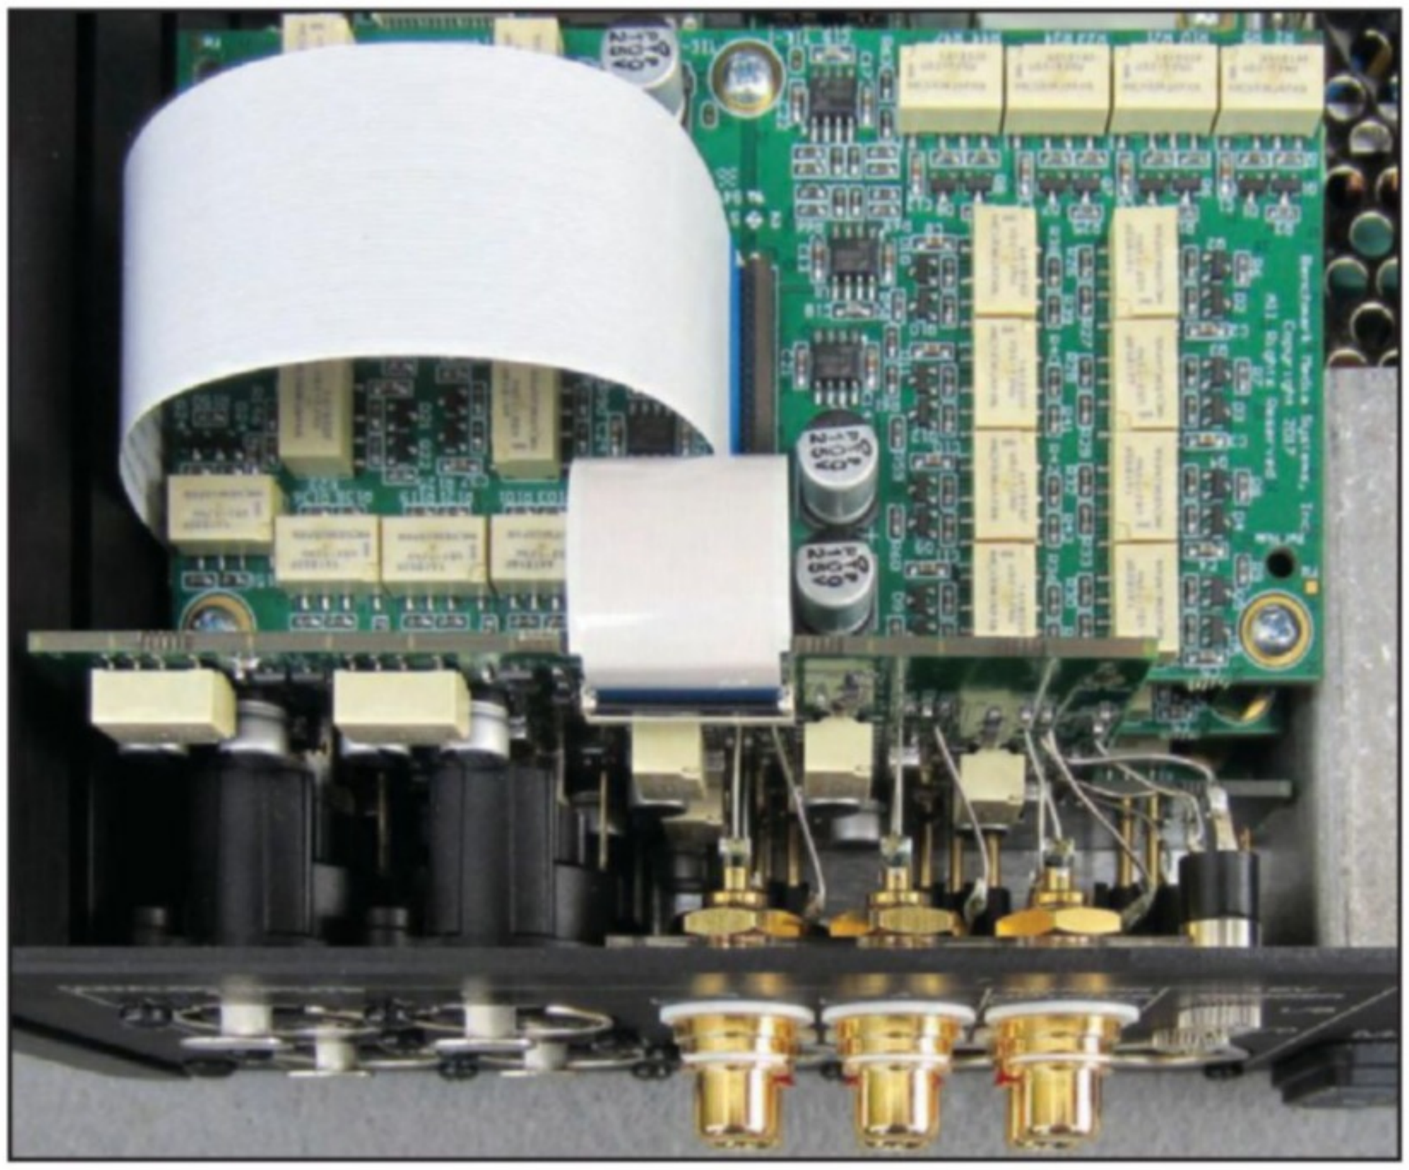

Photo 5: This is another inside view of the HPA4, showing the vertical input/output board and one of the analog boards. All input switching is done with gold-contact relays. The use of relays for volume control and input switching is more precise than conventional methods, and eliminates wiring between the controls and PC boards. This helps ensure the lowest-possible noise and distortion.

Specifications for Benchmark's HPA4

THX-888 Headphone Output
THD: < -125 dB (0.00006%)
SNR: > 131 dB, unweighted, 20-20 kHz
SNR: > 135 dB, A-weighted
Frequency Response: - 0.003 dB at 10 Hz, -0.001 dB at 20 kHz
- 3 dB Bandwidth exceeds 0.1 Hz to 500 kHz
Output Impedance: near 0 Ω
Output Noise: < 2.45 µV at Unity Gain, 20-20 kHz
Maximum Output Power: 6 W into 16 Ω
Maximum Output Current: 1.5 A
Maximum Output Voltage: 11.5 Vrms into 300 Ω
Crosstalk: < -133 dB at 1 kHz, -115 dB at 10 kHz (XLR4)

Balanced Line Outputs
THD: < -125 dB (0.00006%)
SNR: > 135 dB, unweighted, 20-20 kHz
SNR: > 137 dB, A-weighted
Frequency Response: - 0.003 dB at 10 Hz, -0.001 dB at 20 kHz
- 3 dB Bandwidth exceeds 0.1 Hz to 500 kHz
Output Impedance: 60 Ω
Output Noise: < 1.9 µV at Unity Gain, 20-20 kHz
Maximum Input and Output Voltage: 20 Vrms (+28 dBu)
Crosstalk: < -133 dB at 1 kHz, -116 dB at 10 kHz

Dimensions
8.65" W x 3.88" H x 8.33" D - Including Feet
8.65" W x 3.47" H x 8.33" D - Excluding Feet

Weight
8.0 lbs., 12 lbs. shipping

Power Consumption
18 W Typical Input Power
35 W Maximum Input power
0.33 W Standby Power (when off)
Internal Universal Power Supply

Table 1: Manufacturer's specifications for the HPA4

About the Author

Gary Galo retired in 2014 after 38 years as Audio Engineer at The Crane School of Music, SUNY at Potsdam, NY. He now works as a volunteer in the Crane Recording Archive doing preservation, restoration, and digital transfer of vintage Crane recordings. He is also a Crane alumnus, having received a BM in Music Education in 1973 and an MA in Music History and Literature in 1974. Gary is a widely published author with more than 300 articles and reviews on both musical and technical subjects, in over a dozen publications. Gary has been writing for *audioXpress* and its predecessors since the early 1980s. He has been an active member of the Association for Recorded Sound Collections (ARSC) since 1989, and a frequent recording and book reviewer for the *ARSC Journal*. He has given numerous presentations at ARSC annual conferences, many of which have been published in the *ARSC Journal*. He was the Sound Recording Review Editor of the *ARSC Journal* from 1995-2012, and co-chair of the ARSC Technical committee from 1996-2014. Gary has also published numerous book reviews in *Notes: Quarterly Journal of the Music Library Association*, written for the *Newsletter of the Wilhelm Furtwängler Society of America*, *Toccata: Journal of the Leopold Stokowski Society*, and he is the author of the "Loudspeaker" entry in *The Encyclopedia of Recorded Sound* in the US. He has also written several articles for *Linear Audio*. He is a member of the Audio Engineering Society, the Boston Audio Society, and the Société Wilhelm Furtwängler.

Performance

Benchmark lists an array of impressive specifications for the HPA4 (see **Table 1**). The line amplifier and headphone amp have -3 dB points of 0.1 Hz to 500 kHz, and THD of -125 dB (0.00006 %). The ultra-wide bandwidth ensures virtually perfect differential phase response well beyond the audible spectrum. Benchmark's frequency-response graphs show a ruler-flat frequency response for the line amplifier from 10 Hz to 200 kHz (the frequency limits of the Audio Precision AP2722). With the headphone amplifier driving 300 Ω and 16 Ω, the response is down 0.8 dB and 1.1 dB, respectively, at 200 kHz. The phase response of the headphone amplifier with a 300 Ω load is within 0.1°, 10 Hz to 200 kHz. The unweighted line amplifier SNR of -135 dB, 20 Hz to 20 kHz, actually exceeds that of the DAC3 HGC and the AHB2; SNR at the headphone output is -131 dB.

The instruction manual has 18 performance graphs done on the Audio Precision system, showing the outstanding measured performance of the HPA4. I've picked a few to illustrate the unit's extremely low noise and distortion (the full manual is available as a download on Benchmark's website).

The graphs in **Figure 1** show Fast Fourier Transform (FFT) plots of the HPA4's idle channel noise, balanced in to balanced out (see Figure 1a) and balanced in to headphone out (see Figure 1b). The left and right scales are calibrated relative to +28 and +22 dBu, respectively (+22 dBu is the level required to drive the headphone amp and the AHB2 power amp to their full, rated output). The graphs show that the headphone amplifier retains the exceptional noise performance of the line amplifier, with idle noise at or below -160 dB over most of the audible spectrum.

Figure 2 shows AC line-related hum at the headphone output. The graph shows hum to be at least -152 dB relative to the headphone amp's maximum output of +22 dBu. This graph makes a good case for a properly designed SMPS, which virtually eliminates any trace of power line hum.

The graphs shown in **Figure 3** are FFT plots showing total harmonic distortion (THD) of a 1 kHz tone, at an output level of +22 dBu. The upper graph (a) is balanced in to balanced out. The third harmonic measures better than -124 dB (0.00006%) and the fifth harmonic measures better than -141 dB (0.000009%). The lower graph (b) is balanced in to headphone out, with the headphone output driving its full rated 6 W into a 16 Ω load. The second harmonic measures better than -127.5 dB (0.00004 %) and the third harmonic measures better than -126.5 dB (0.00005 %). The manual also contains 10 kHz THD comparisons of the line and headphone outputs—

they're essentially identical, with second and third harmonics measuring better than -124 dB (0.00006%) at both outputs. A 16 Ω load at an output level of 6 W is a worst-case situation, not likely to be encountered in any real-world application, by anyone who values their hearing (!). The fact that the headphone output shows essentially no degradation of the signal fed to it, under extreme output and loading conditions, is a testament to the effectiveness of the THX AAA error correction. Benchmark's data shows the measured performance of the HPA4 to be state-of-the-art, and beyond reproach.

Listening

I'm using the HPA4 as a line amplifier connected between my DAC3 HGC and pair of AHB2 power amps. I have one AHB2 per channel, in stereo mode, passively bi-amping my Audio Concepts Sapphire III/Sub 1 loudspeaker systems. The AHB2 power amps are set to Benchmark's recommended +22 dBu, which is their lowest sensitivity position (toggle switches all the way down). The output attenuators on the DAC3 are set to 0 dB, and for most of my listening, I've set the DAC3 in Compatibility Mode.

My OPPO Digital UDP-205 digital player's HDMI 2 output feeds a KanexPro HDMI De-Embedder which, in turn, is connected to one of the DAC3's S/PDIF inputs. I also connect the OPPO Digital player's balanced analog outputs directly to the XLR2 inputs on the HPA4, to play SACDs in native DSD mode. All analog interconnections are balanced, using D.H. Labs Air Matrix cables and Neutrik gold-plated XLR connectors.

With the DAC3 HGC connected directly to my AHB2 power amps, my system has performed superbly, with excellent detail and sound-stage

Figure 1: These graphs show FFT plots of the HPA4's idle channel noise, balanced in to balanced out (a) and balanced in to headphone out (b). Green is the left channel, red is the right channel, and yellow is the summed mono output in the top graph (a) only. The left and right scales are calibrated relative to +28 and +22 dBu, respectively. The graphs show that the headphone amplifier retains the exceptional noise performance of the line amplifier, with idle noise at or below -160 dB over most of the audible spectrum. All of Benchmark's measurement graphs are done with the Audio Precision AP2722. (Image courtesy of Benchmark Media Systems)

reproduction, and a sonic presentation that provides hours of fatigue-free listening. Given the superb measured performance of the HPA4, I expected the HPA4 line amplifier to be sonically transparent, so it would not degrade the output of the DAC3 HGC, In other words, its presence in the signal path would not be detectable. What I heard was not at all what I expected. What I've heard, consistently, is an across-the-board improvement over the DAC3 driving the power amps directly. On my system, several things are obvious: greater inner detail and resolution; larger and more precise soundstage; noticeably improved dynamics, low-frequency extension and control. I wasn't prepared for this!

In my memorial tribute to C. Victor Campos (*audioXpress*, March 2018) I mentioned the large quantity of 15-IPS copies of master tapes he had acquired from major record companies, for his program "Adventures in Sound" that aired on WGBH-FM in the 1970s. I did high-resolution digital transfers of several of his favorite tapes for him between 2010 and 2016—some of the original recordings are revered by audiophiles.

My digital transfers of these superb tapes have been staples of my own audio listening, for equipment evaluation, and sheer musical and sonic pleasure. I have played them dozens of times and should, by now, intimately know their sonic qualities. I was wrong. With the HPA4 in my system, I hear inner details and subtle soundstage cues that I had not previously experienced on these recordings. Some of these recordings have low-frequency extension and weight that I didn't know were there. The dynamic range also seems wider. In pieces that begin softly and end with a full orchestra, I've found that I need to start at a lower subjective volume than before, in order to keep levels comfortable at the end. I've noticed these improvements in all of my reference recordings, whether CDs or higher-resolution sources.

How can this be? I posed that question to John Siau. He provided a detailed explanation of how inserting the HPA4 in the signal path actually results in lower noise than with the DAC3 HGC feeding the AHB2 directly. Running the DAC3 HGC in Compatibility Mode allows the DAC to operate at maximum SNR. Running the DAC3 HGC directly into the AHB2 gives the following SNR results:

Full system output: System SNR = 125 dB
-6 dB output: System SNR = 119 dB
-8 dB output: System SNR = 117 dB

Inserting the HPA4 line amp in the signal path results in improved SNR at the same system output levels:

Full system output: System SNR = 127 dB
-6 dB output: System SNR = ~125 dB
-8 dB output: System SNR = 125 dB

But, I suspect there is more going on here than just an improvement in SNR. Siau confirmed that with the DAC3 HGC in Compatibility Mode, the 32-bit digital volume control in the ES9028PRO DAC chip is out of the picture. As good as that digital volume control is, system performance may not be as transparent as it is with the analog relay volume control in the HPA4. In its two-page cut

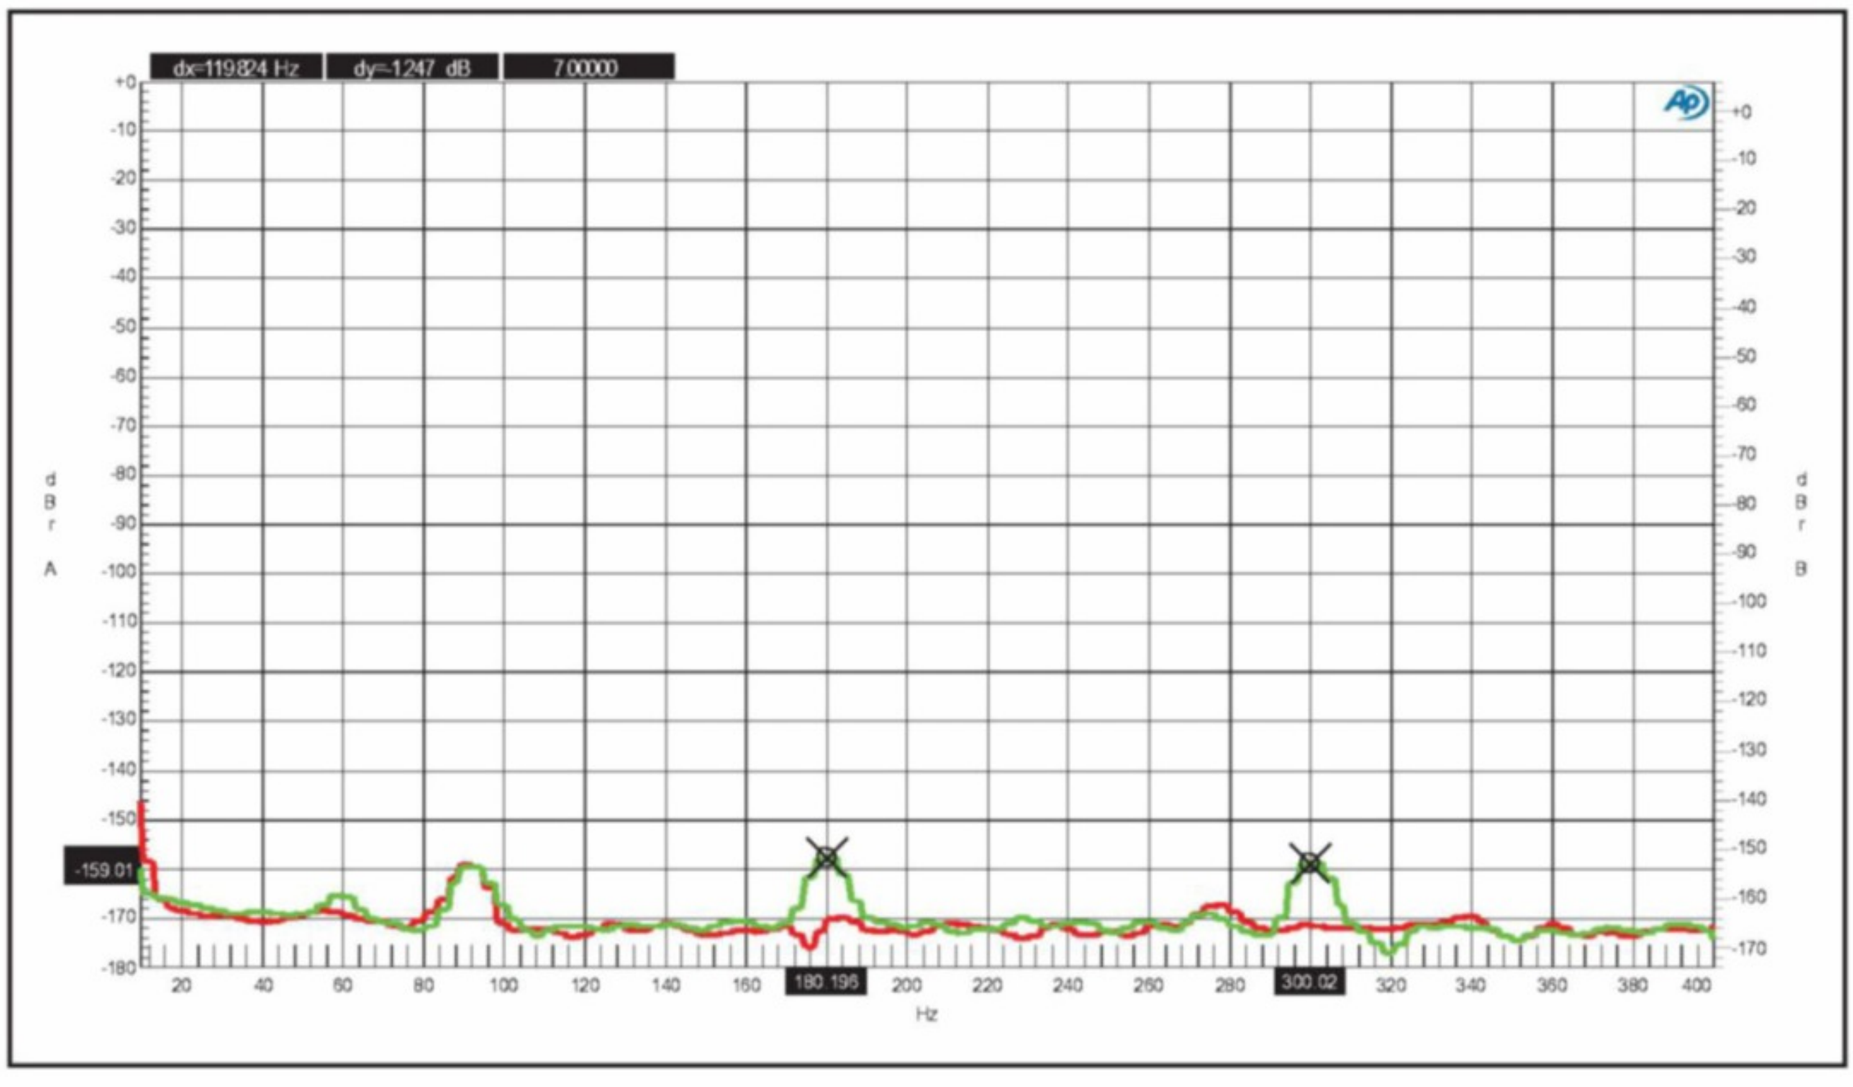

Figure 2: This graph shows AC line-related hum at the headphone output. The graph shows hum to be at least -152 dB relative to the headphone amp's maximum output of +22 dBu. The switching-mode power supply, which doesn't use a conventional 60 Hz power transformer, virtually eliminates any trace of power line hum. (Image courtesy of Benchmark Media Systems)

sheet describing the HPA4, Benchmark notes that the fully-analog relay gain control is "more resolving" than the 32-bit digital gain control in the DAC3.

This prompted me to conduct an experiment. In Compatibility Mode, the DAC3 HGC's volume control is positioned at 2:00, and the calibrated output level is 3.5 dB below maximum output. Since the ES9028PRO's digital volume control is bypassed, that 3.5 dB reduction in level is accomplished entirely by the DAC3's analog volume control. What would happen if the DAC3 were taken out of Compatibility Mode, and run with the volume control at full rotation?

I set the XLR1 input level trim on the HPA4 to -3.5 dB, to compensate for the 3.5 dB increase in level from the DAC. So, the system playback level remained unchanged. What I heard was a subtle increase in ambience retrieval, which put the entire soundstage just a bit further back, behind the front plane of my loudspeakers, than before. In addition to directly driving the AHB2s, I've also had my own custom-built preamp in the line between the DAC3 HGC and the power amps. My preamp has the transformer-coupled balanced outputs I described in the February 2019 issue of *audioXpress*. With this arrangement, I normally had the DAC3's volume control at 2:00, to bring the volume control on my preamp closer to the center of its rotation. I tried this same experiment with my own preamp in the line and could not really hear any difference between the DAC3 at full volume or at 2:00. It's now obvious that my preamp was masking a difference that that the HPA4 easily reveals.

I shared my reactions with Siau and he noted one caveat: Running the DAC's volume control at maximum reduces the headroom available for inter-sample overs (see Benchmark's App Note on intersample overs, and my review of the DAC2 DX). He noted that the 3.5 dB reduction with the output level calibrated is to allow sufficient headroom for the reproduction of intersample peaks. Running the DAC's volume at maximum reduces the available headroom.

Reources

Benchmark Application Notes, Benchmark Media Systems, Inc., https://benchmarkmedia.com/blogs/application_notes.

G. Galo, "A Transformer-Coupled Output for Solid-State Preamps," *audioXpress*, February 2019.

——, "OPPO Digital UDP-205," audioXpess, November 2017, www.audioxpress.com/article/oppo-digital-udp-205-review-a-4k-ultra-hd-audiophile-blu-ray-disc-player

——, "Benchmark DAC3 HGC," *audioXpress*, July 2017, www.audioxpress.com/article/fresh-from-the-bench-benchmark-dac3-hgc-stereo-d-a-converter

——, "KanexPro HDMI De-Embedder" *audioXpress*, July 2016.

——, "Benchmark DAC2 DX," *audioXpress* website, July 2015, www.audioxpress.com/article/Fresh-from-the-Bench-Benchmark-DAC2-DX-Stereo-D-A-Converter

——, "Benchmark AHB2," audioXpress, April 2015, www.audioxpress.com/article/Fresh-from-the-Bench-Benchmark-AHB2-Stereo-Power-Amplifier

Sources

THX AAA Headphone Amplifiers
THX, Ltd. | www.thx.com/mobile/aaa

Texas Instruments PowerPAD IC Package
Texas Instruments | www.ti.com/general/docs/lit/getliterature.tsp?baseLiteratureNumber=slma002&fileType=pdf

Converting THD decibel to Percent and Vice Versa
Sengpiel Audio | www.sengpielaudio.com/calculator-thd.htm

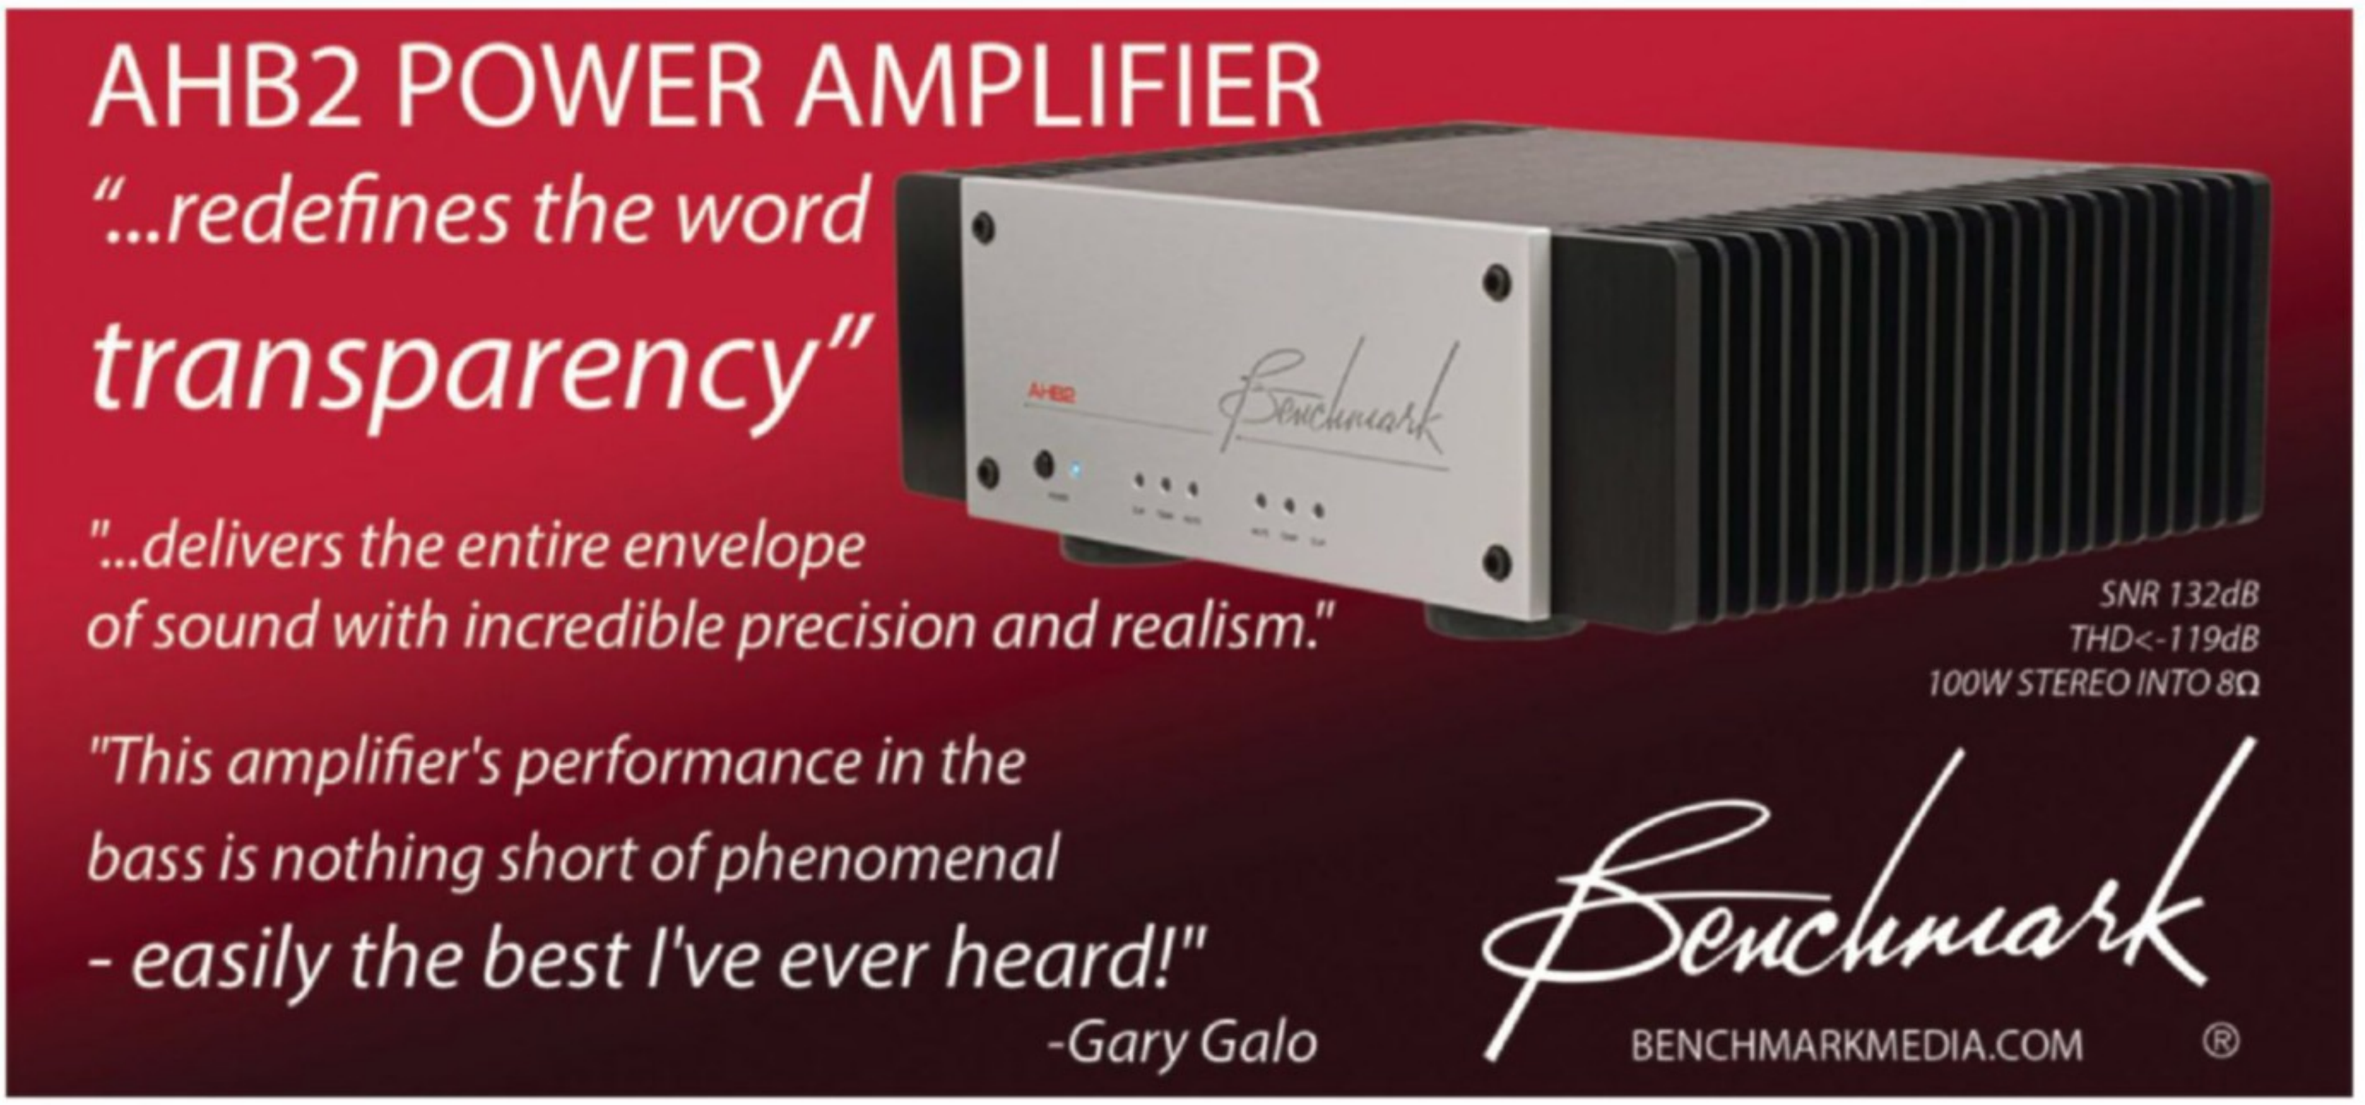

I listen to classical music almost exclusively. Classical music tends to have wide dynamic range, with much of the program material in most pieces well below maximum level. On a lot of popular music recordings, dynamic range is much more limited. I've viewed many such recordings on my digital editor and observed long stretches—and sometimes entire songs—where the volume is at or close to maximum level. This will make intersample overs a potentially greater problem, since they can occur much more frequently. I suggest trying my experiment with the music you enjoy to see what your ears tell you. As I explained in my DAC2 DX review, intersample overs are only a problem with sample rates of 48 kHz or lower. At higher sample rates, they're not an issue.

I evaluated the headphone amplifier with my AKG K-240 headphones, which have a nominal impedance of 62 Ω. The HPA4 drives these headphones with ease, with plenty of headroom, and seemingly no sonic character of its own. What you hear from the line amplifier is faithfully passed to your headphones, with the same level of detail and resolution. The THX AAA error correction really does provide the same benefits for headphones that the AHB2 does for loudspeakers. Compared to the already-excellent headphone amp in the DAC3 HGC, the HPA4 is simply cleaner, more spacious and more detailed, with better dynamics

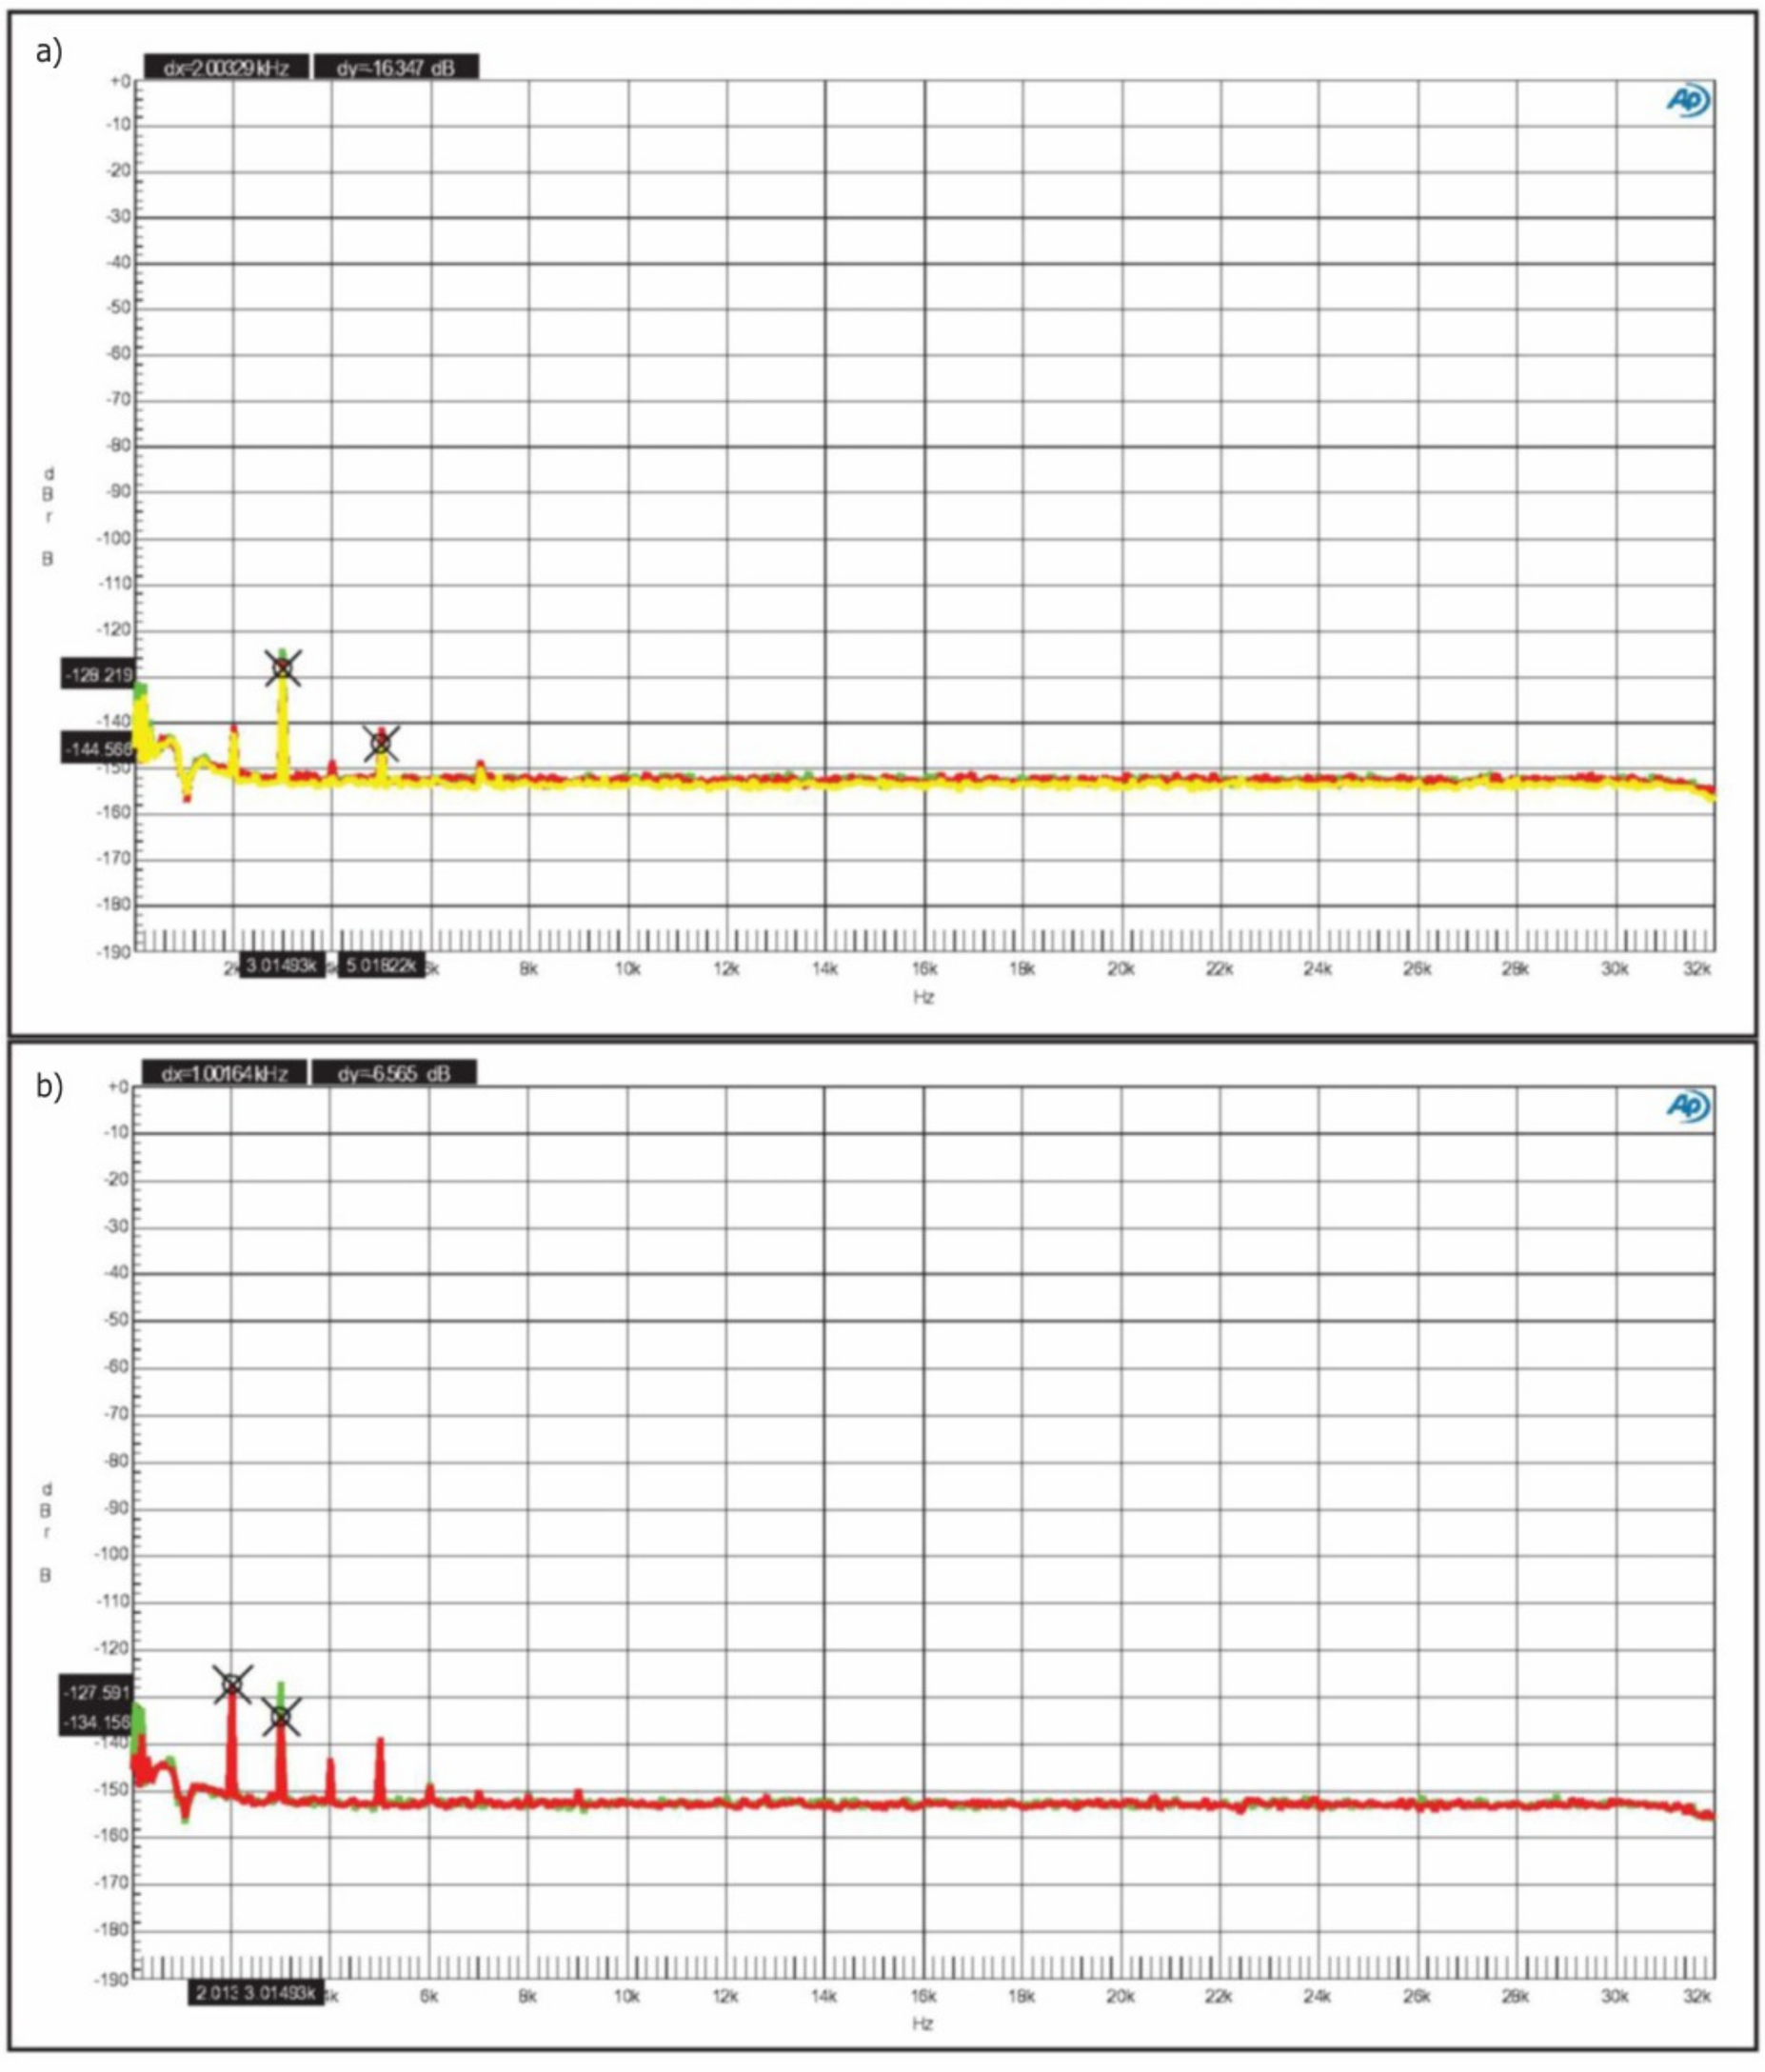

Figure 3: These graphs are FFT plots showing the THD of a 1 kHz tone, at an output level of +22 dBu. The upper graph (a) is balanced in to balanced out. Green, red and yellow are left, right, and summed mono, respectively. The third harmonic measures better than -124 dB (0.00006%) and the fifth harmonic measures better than -141 dB (0.000009%). The lower graph (b) is balanced in to headphone out, with the headphone output driving its full rated 6 W into a 16 Ω load. Green and red are left and right. The second harmonic measures better than -127.5 dB (0.00004 %) and the third harmonic measures better than -126.5 dB (0.00005 %).
(Images courtesy of Benchmark Media Systems)

and better low-frequency extension and control.

The AHB2 power amps have always left me with the impression that they have unlimited speed—no audio signal is capable of outrunning them. This, no doubt, contributes to their subjectively crisp, clean performance, free of any transient smearing. I have the same reaction to the line amp and the headphone amp in the HPA4. By comparison, my own preamp seems just a bit sluggish.

Like the other Benchmark products I use, the line amp and the headphone amplifier achieve exceptional detail without any exaggeration in the treble region. The HPA4 is tonally neutral and does not contribute any brightness or harshness to the sonic presentation—it has provided me with many hours of fatigue-free listening. The audio professional will find the HPA4 to be an accurate, reference-quality tool for evaluation of their work, whether in the recording studio or in the field. For the audio enthusiast, the HPA4 will provide listening that's as musically satisfying as the recordings in their collections. The audio it's fed from your source is what you will hear—no more and no less. And, that's exactly as I want it. Although the HPA4 makes a synergetic match to Benchmark's own DACs and power amp, it should enhance the performance of any audio system. Its flexible gain adjustments will allow it to interface easily with consumer and professional equipment.

Ear vs Measurement

Benchmark addressed the issue of measurements versus listening head-on in an application note posted on its website (posted March 19, 2016; go to Company/About Us). These four points summarize its views:

- Transparent products will usually produce an impressive set of measurements.
- Non-transparent products will usually produce a "poor" set of measurements.
- Impressive measurements do not necessarily guarantee that transparency has been achieved.
- Transparency must be confirmed with listening tests.

Benchmark notes that it has sometimes found that a design that measured well still exhibited problems in listening test. These experiences have helped the company refine its measurements so they correlate better with what their ears tell them. The audiophile world has traditionally been divided into two camps—the purely subjective, who believe that measurements are meaningless, and that our ears tell everything we need to know; and the purely objective, who believe that anyone claiming to hear things that are not measurable are delusional and lacking in scientific rigor. Benchmark occupies the sensible middle-ground, which is why its products have received acclaim by users on both sides of the fence. I agree with Benchmark's views on this still-contentious issue.

The HPA4 is the latest embodiment of its design philosophy. You'll have a tough time finding a cleaner, quieter, more transparent line amplifier and headphone amplifier anywhere, at any price. I found it to be a new reference. Once I heard it, I could not live without it, so I purchased the review sample. The HPA4 deserves my highest recommendation.

Hooke Verse Bluetooth Binaural Recording Headphones

The Hooke Verse, part Bluetooth headphones and part wireless binaural recording device, aims to empower creators-on-the-go with three-dimensional recording. We tested it to find out how it sounds and who should buy this product.

Photo 1: The Hooke Verse is a binaural recording system piggy-backing on the form factor of wireless headphones.

By
Luke McCready
(United States)

The smartphone in your pocket right now has a stunning camera: objectively poor quality images are becoming a thing of the past because Apple, Samsung, and other brands created a revolution in mobile, affordable image quality. Now, Hooke Audio intends to do the same with portable and affordable binaural audio recording with its Hooke Verse.

I recently got the chance to interview Hooke Audio founder and CEO Anthony Mattana. When I asked what problem the Hooke Verse was created to solve, he said:

"Think of the last amazing thing you saw and how easy it was to capture and relive that moment on your phone. Now think of the last amazing thing you've heard and how disappointing it sounded when you tried to capture and relive it. Advancements in video tech have brought 4K video to smartphones, but when it comes to recording sound, we're stuck in the 1980s. Most phones still capture flat, mono audio. We hear the world in 3D; we should be capturing and reliving it that way."

Hooke Verse

The debut product of Hooke Audio is the Hooke Verse, a binaural recording system piggy-backing on the form factor of wireless headphones at the price of $160. The two in-ear-monitors attach by cable to packages containing batteries and electronics that rest behind each ear. On the outside of each in-ear-monitor is an omnidirectional electret condenser microphone. Wireless stereo recording via Bluetooth is the Hooke Verse's claim to fame, and it's designed to complement smartphone use by recording to your smartphone's storage as a stereo audio or video file. Beyond smartphones, the Hooke Verse uses Bluetooth or a wired connection to act as the binaural recording rig for your DSLR, GoPro, or field recorder (see **Photo 1**).

I asked Mattana about the target market for the Hooke Verse, and he said, "For people who create premium content, Hooke Audio makes tools to capture and relive moments the way they are sonically and visually experienced. Podcasts, vlogs, music, and travel videos come to life."

The Hooke Verse features a Bluetooth 4.1 radio and records 48 kHz/16-bit stereo sound. Hooke Audio's website claims a headphone frequency response of 20 Hz to 16 kHz and a microphone frequency response of 20 Hz to 20 kHz, though these measurements mean little without notated precision. Hooke Audio's website boasts 8 to 9 hours of recording time on a full battery charge, or 9 to 10 hours of continuous audio playback.

The Story of Hooke Audio

Mattana shared that the idea for the Hooke Verse came to him when he was working as a theatrical sound designer on Broadway. "I was frustrated by the fact that both my audiences and collaborators

didn't fully understand how powerful audio is to storytelling," he said. "The directors and producers could speak intelligibly about color temperature, brightness, and focus, but when it came to sound, they were at a loss for words." Mattana theorizes that we've developed the vocabulary for what we can create, and we can create excellent visual content because of the great cameras in our pockets. But not so for audio. Mattana added, "I felt that if I could provide a tool that made changing the way you record the world as easy and professional as an iPhone camera, then maybe people would care more about sound. Thus, Hooke Verse was born; the first 360 [degree] sound recording Bluetooth headphones."

Mattana shared that the biggest challenge he and his team faced was transmitting multichannel lossless audio via Bluetooth: because it had never been done before, no existing codec supported it. "So we built that codec from the ground up."

He assembled a team that built a prototype Hooke Verse, then the team campaigned on Kickstarter for the resources to launch the product. During the crowdfunding campaign, 1,385 people donated $163,000, giving Hooke Audio the backing it needed. A few months later, Hooke Verse entered manufacturing.

Usability and Fit

The headphones have a one-button interface and a single physical port that can be used for charging or wired recording. I found the Hooke Audio app, the smartphone hub for the Hooke Verse, in the Google app store without any trouble. The app includes tutorials on navigating the app and pairing the device (see **Figure 1**).

The Hooke Audio app aims to be a single-program

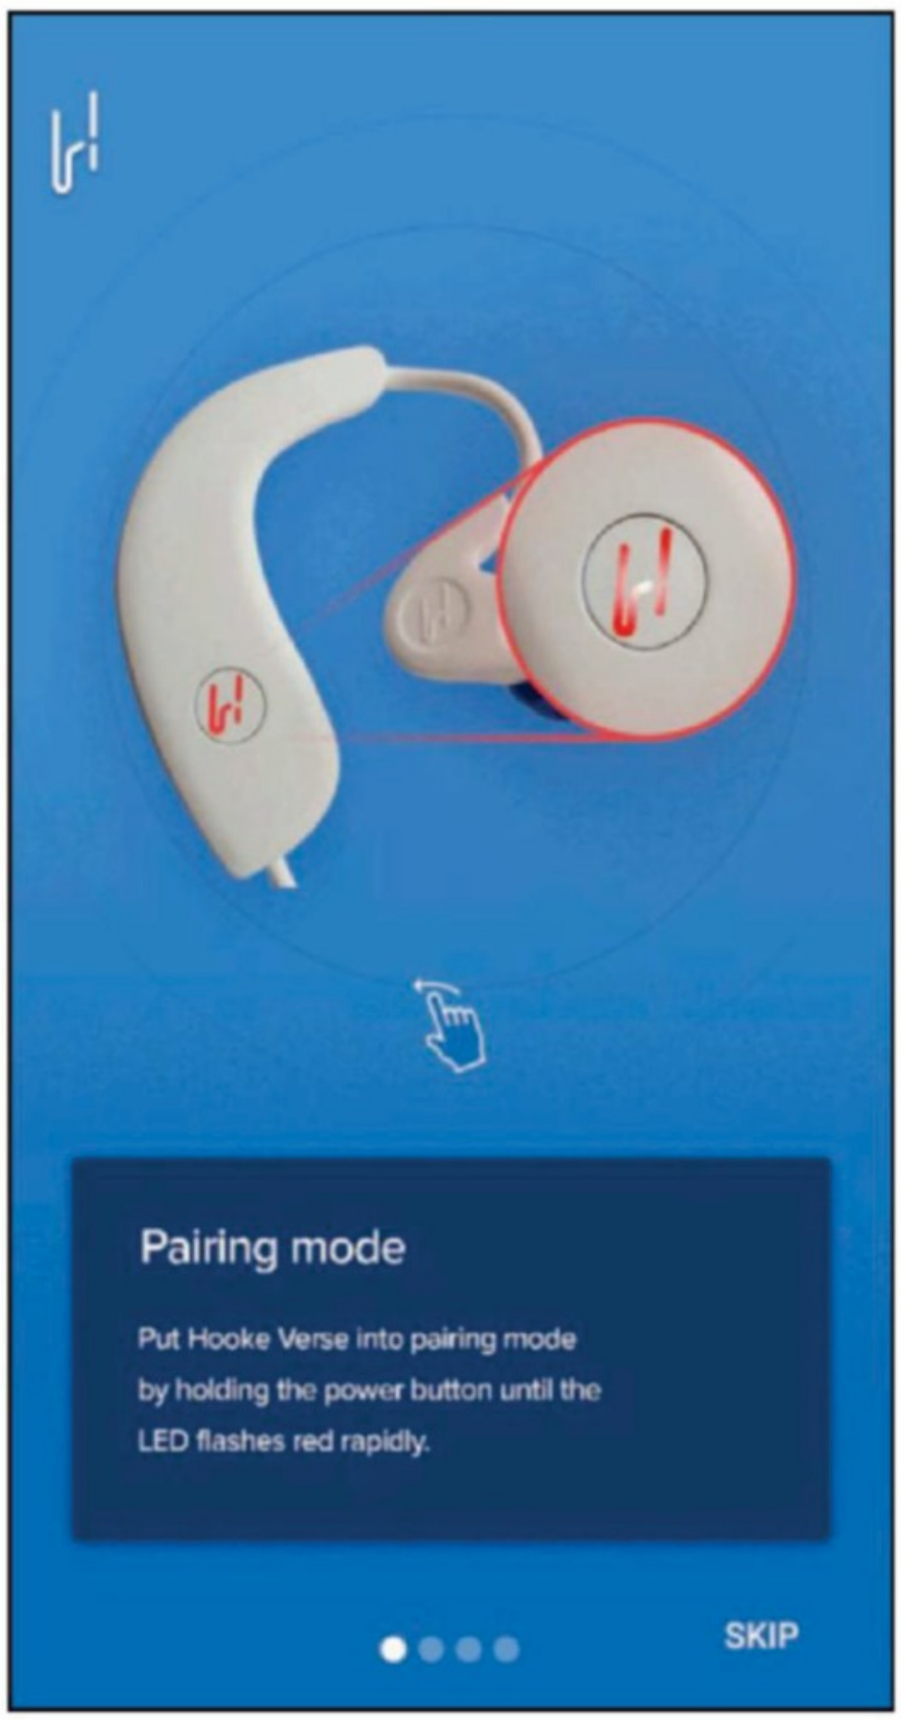

Figure 1: The app's tutorial on Bluetooth pairing helps get you up and running.

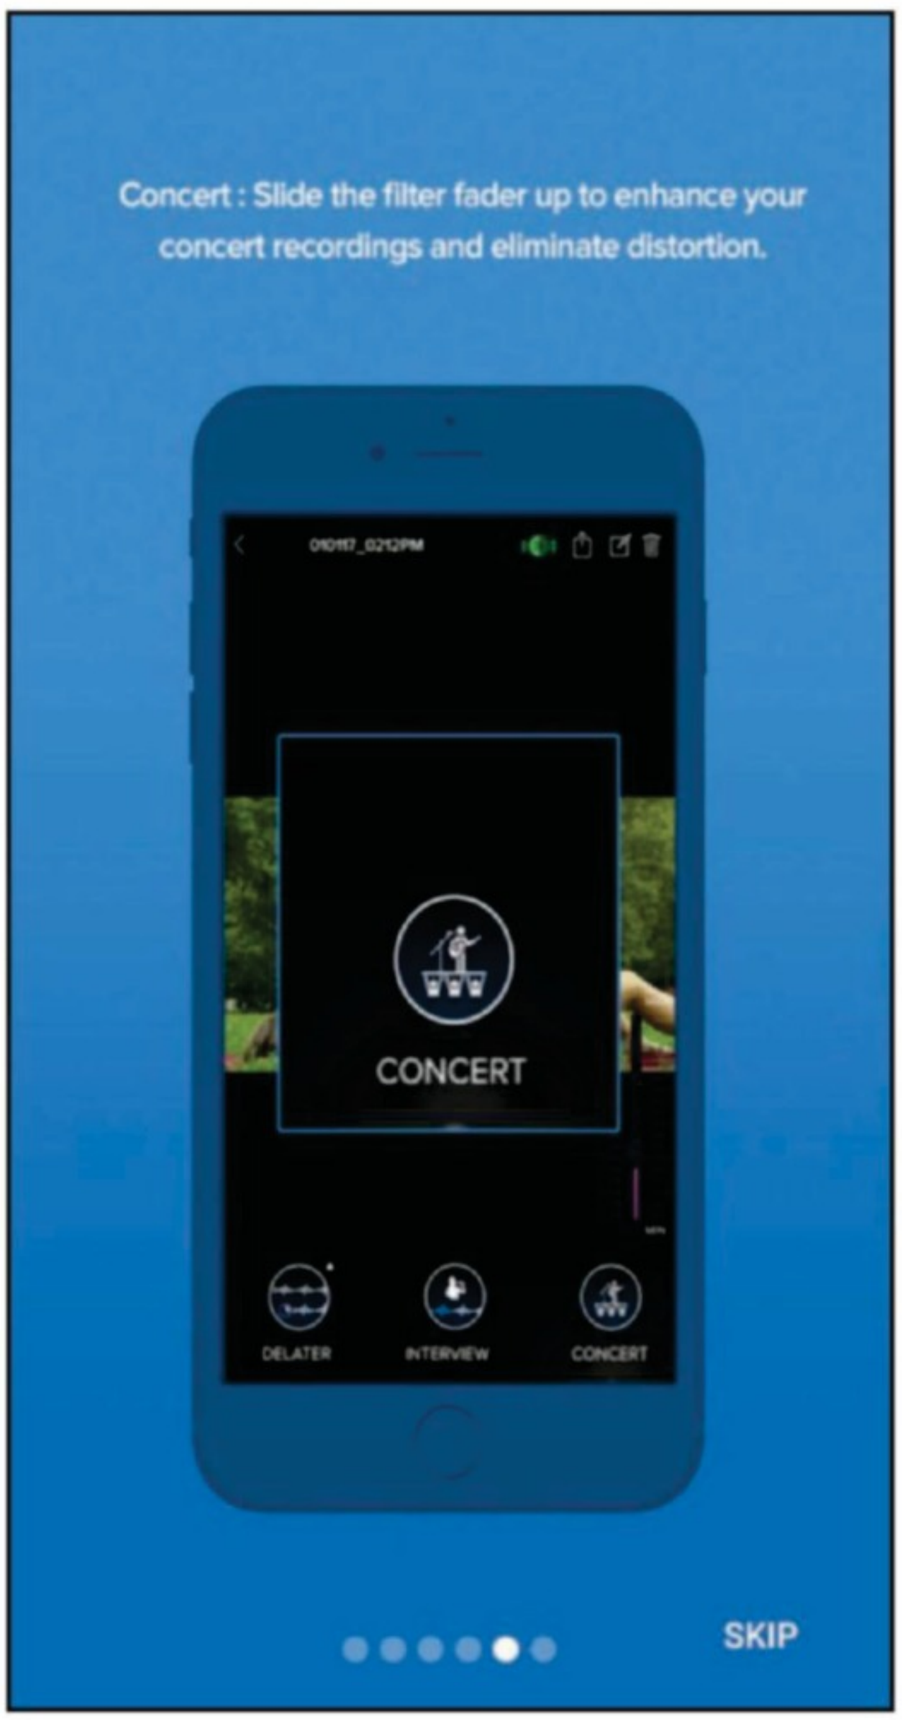

Figure 2: The Concert filter is designed to reduce distortion on previously recorded audio.

Photo 2: Hooke Verse accessories include a cable for wired audio recording, three different size foam tips, and three sizes of silicone tips.

solution for all of your Hooke Verse recordings: it is capable of recording audio, recording video combining the Hooke Verse microphones with your smartphone's camera, media storage, and playback, and even sharing clips on social media. For $5, the app will sell you a pack of "filters" for adding reverb, delay, and stereo widening, and also a gate for filtering out background noise and a concert mode that can supposedly reduce distortion on already-recorded clips (see **Figure 2**).

The headphones come fitted with large foam tips, which were not quite large enough to fit snugly in my ears. The packaging came with two smaller sizes of foam tips and three sizes of silicone tips: however, the largest of the silicone ear tips were also too small for me to consistently get a good seal (see **Photo 2**).

Photo 3: I recorded the same test simultaneously with three devices (pictured left to right): the built-in microphones on a Samsung Galaxy S8 smartphone; the Hooke Verse recorded over Bluetooth; and my Tascam DR-05, an $80 field recorder.

Playback Quality

The audio playback of the Hooke Verse is excellent. The frequency spectrum is largely balanced, provided you have a good seal, and the sound is crisp and detailed. Voices sound natural and instruments are heard with precision. There is a mid-bass bloom, as is common in headphones—a positive trait to many consumers. If I were to be picky, I'd say that I don't care for the mid-bass bloom, the sub-bass is a little lacking, the high-end frequency response isn't quite extended enough, and there's a mild brittleness to hi-hats and other dark-treble sounds. But these are small differences: the Hooke Verse sounds better via Bluetooth than 80% of the wired in-ears I've heard across all brands and price points.

Recording Quality

I compared the Hooke Verse to devices our readers may own by recording in a busy cafe with an open kitchen, for sounds of refrigerators and ovens, and clinking dishes and dialogue coming from all directions. I recorded the same test simultaneously with three devices: the Hooke Verse recorded over Bluetooth; the built-in microphones on a Samsung Galaxy S8 smartphone; and my Tascam DR-05, an $80 field recorder (see **Photo 3**).

Listening to the test results, my immediate impression was of stereo width: the Hooke Verse is very wide, even too wide. The field recorder sounded the narrowest: still stereo, but not as enveloping. The smartphone audio felt the most balanced in stereo width.

The Hooke Verse sounded quite dull and benefited from equalization with a significant high-shelf boost. The field recorder needed less aggressive EQ, and the smartphone recording, to my ears, did not need EQ.

In terms of raw fidelity—that tactile precision of good sound—the field recorder took the win, with a tie between the Hooke Verse and the smartphone.

3D Audio

The biggest strength of the Hooke Verse is stereo separation: this lends itself not only to the separation of primary sounds, but also a hyper-realistic feeling of space in that ambient noises feel widely panned. To some, this might sound amazing; to others, it's a bit too much. And while this might be the very definition of "3D Audio" to some, let's take a moment to explore this claim.

The term "3D Audio" is found so often on product packaging and websites, I'm convinced it's metastasizing. There's nothing 3D about it: It's simply two-channel binaural audio by nature of

recording from the ears. The natural baffle of the human skull creates the stereo separation, but there is no sensation of height or even much difference between sounds directly in front of and behind the person recording. The mics simply aren't deep enough in the ear canal for these benefits to appear. The microphones on the Hooke Verse can't record Ambisonic or multichannel audio beyond stereo, and the stereo headphones offer no form of multichannel up-mixing.

Some Issues

The electronics on the Hooke Verse are noisy. When recording at podcasting levels, the noise floor was at -45 dBFS: soft, noticeable, and unprofessional. And in some of my testing, the noise floor rose to -33 dBFS: certainly undesirable.

The noise floor seems fixed before adding gain in the Hooke Audio app or any external program. It would be less of an issue recording louder sounds in a louder environment. I tried noise removal tools in my digital audio workstation with less than ideal results (see **Figure 3**).

Also, the left microphone of my Hooke Verse is 2 dB quieter than the right one. This differential is

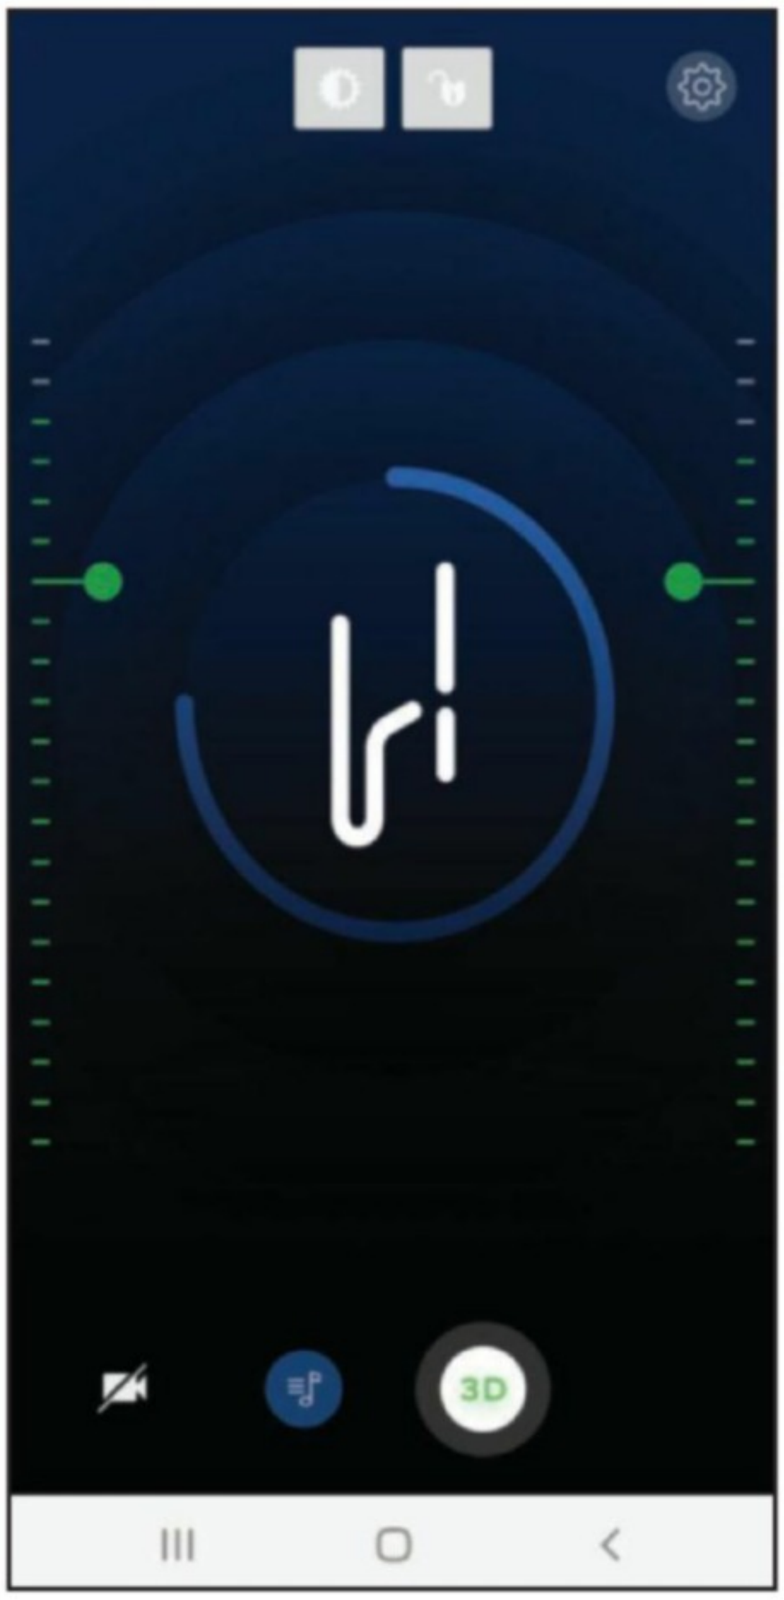

Figure 3: Slide the faders to apply gain in the Hooke Audio app.

SWANS SPEAKERS KITS

www.swanskits.com

Take your pick

Swanspeakers M3
$1699.99

OR

Swankits 3.1
$289.99

*same components in each

3.1 is avialable on amazon.com

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

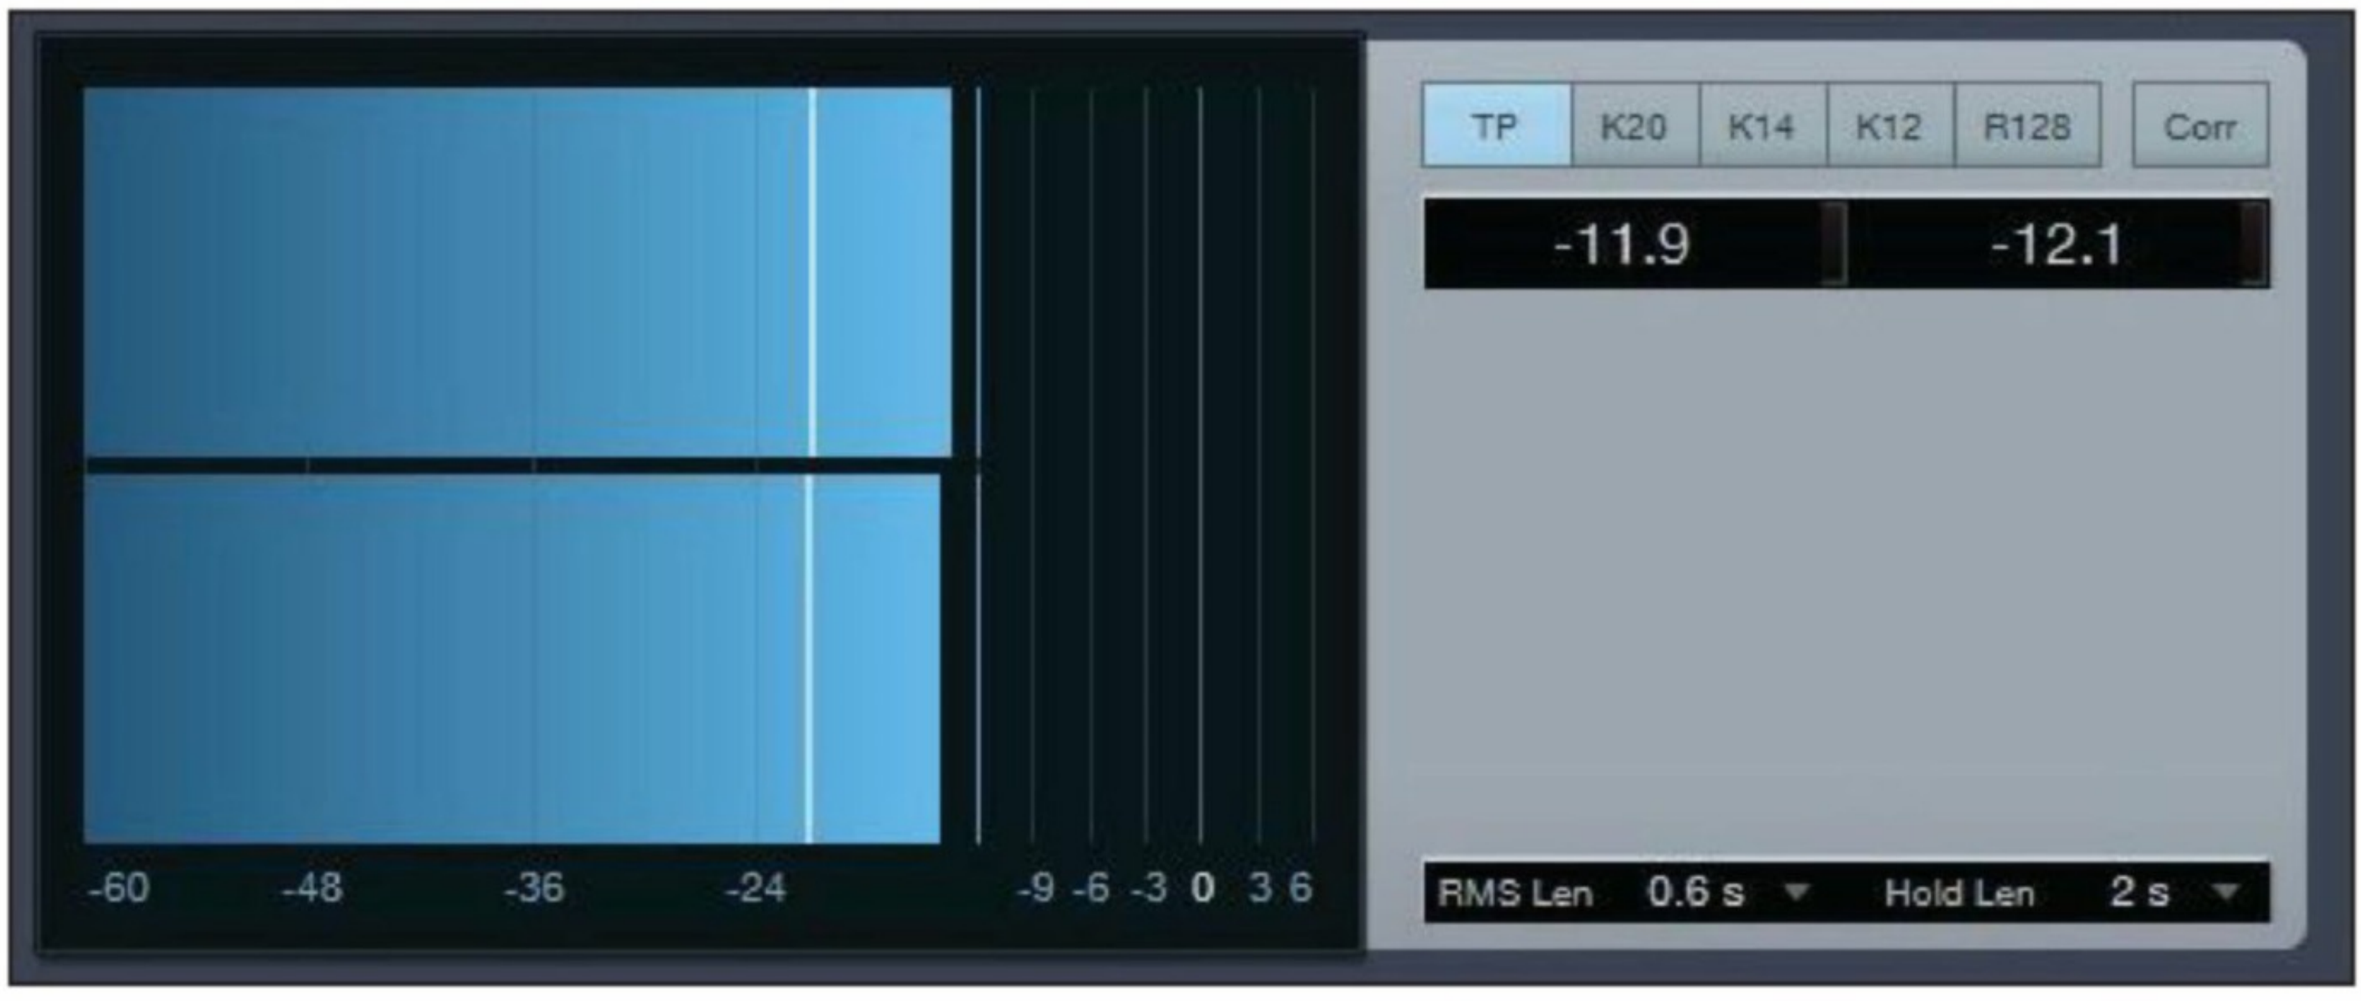

Figure 4: This is the unusable noise floor when using the Hooke Verse to record into my field recorder.

noticeable. I would have hoped for more carefully matched microphones in a product at this price.

I plugged the Hooke Verse into my field recorder using the included cable to record audio without my smartphone. My recorded noise floor was astonishingly loud (see **Figure 4**), rendering this method completely unusable, at least with the one field recorder I tested with. Neither the included documentation nor the Hooke Audio website specifies the Hooke Verse's line output voltage or sensitivity. For reference, my field recorder's manual specifies an input sensitivity of 25 kΩ and a nominal input level of -20 dBV.

While recording to my smartphone via Bluetooth, I experienced dropouts. The headset was only a couple of feet from my smartphone, but every 5 to 20 seconds, the audio would dip into muffled distortion or a full second of silence. I worked with Hooke Audio on troubleshooting the audio dropouts, and we discovered that the dropouts disappear when I turned on my smartphone's airplane mode. Mattana suggested that "interference from cell signal is most likely due to harmonic frequencies of the cell bands causing interference with the Bluetooth channel."

Though I am pleased the airplane mode helps, this issue very likely would be solved by upgrading the Hooke Verse Bluetooth radio: the advantages include potentially increasing the data rate and a theoretical 4x improvement in range. While Bluetooth 4.1 includes error detection, Bluetooth 5.0 introduces error correction, allowing for accurate signal decoding at a significantly lower signal-to-noise ratio (SNR). This, combined with Bluetooth 5.0's newly robust Adaptive Frequency Hopping algorithm and nature of sending signal over multiple Bluetooth channels simultaneously, contributes to greater stability when using Bluetooth to transmit data, particularly audio.

Looking to the Future

When asked about the future of Hooke Audio, Mattana said, "We're looking into solutions for live streaming in binaural 3D audio as well as experiencing binaural 3D audio without the use of headphones. I can't say much yet, but we have some very exciting product launches on the horizon."

Conclusion

The reason for the Hooke Verse to exist is that supposedly all smartphone audio recordings sound awful and most are limited to mono. I tested on one- and two-generation old smartphones that both record in stereo, and the fidelity of each surprised me. In my opinion, a stereo recording from a recent smartphone is about equal in quality to the Hooke Verse, and even an affordable field recorder sounds superior.

Resources

Hooke Audio, https://hookeaudio.com

"Hooke: Wireless 3D Audio Headphones," Kickstarter, www.kickstarter.com/projects/124900576/hooke-wireless-3d-audio-headphones.

M. Woolley, "Bluetooth 5: Go Faster, Go Further," Bluetooth, https://www.bluetooth.com/bluetooth-resources/bluetooth-5-go-faster-go-further/?utm_campaign=specification&utm_source=internal&utm_medium=blog&utm_content=exploring-bluetooth-5-how-fast-can-it-be.

However, absolute quality doesn't always lead to the best-selling, most impactful product. The revolution of consumer image quality in the last five years isn't because top-end DSLRs achieved a breakthrough in image quality: It's because Apple, Samsung, and others decided to go all-in on squeezing the highest imaging performance possible out of the smartphone camera. Likewise, the Hooke Verse doesn't need to be the world's best to make a splash. But the question remains: Is it easy enough to use—as easy as a smartphone camera? If you're wearing your Hooke Verse all the time, maybe. But if not, it's just another gadget to remember to charge, carry with you, and power up at the opportune moment. To make the biggest difference, perhaps Hooke Audio needs to partner with smartphone makers to better the audio recording in the devices people are sure to carry with them.

Stereo Bluetooth recording is a step forward, but the Hooke Verse might seem more revolutionary if it achieved Ambisonic or multichannel recording beyond stereo in a subtle, wearable form-factor.

If the purpose of the Hooke Verse is to enable professionals to create, I can't see the angle. The audio isn't high enough quality and the omnidirectional mics pick up too much ambient sound to be useful for mobile filmography. Music videos are generally better off sounding anthemic versus present. And podcasts are a poor option since the Hooke Verse picks up the wearer's voice in a strange, muffled way, as if standing behind a camcorder.

In my opinion, the best use-case for the Hooke Verse is for travel videos without narration, as an easy-to-pack product for trips, using sound to give context and environmental space to videos where the perception of being there is possibly more important than quality. That said, the headphones sound quite nice on their own, and there's reason to enjoy the product as a consumer rather than a creator.

About the Author

Luke McCready is an audio engineer, a technical writer, and a music lover based out of Denver, CO. He graduated with a degree in audio engineering from North Central University with a passion for audio and a fascination for speaker setup and acoustic treatment.

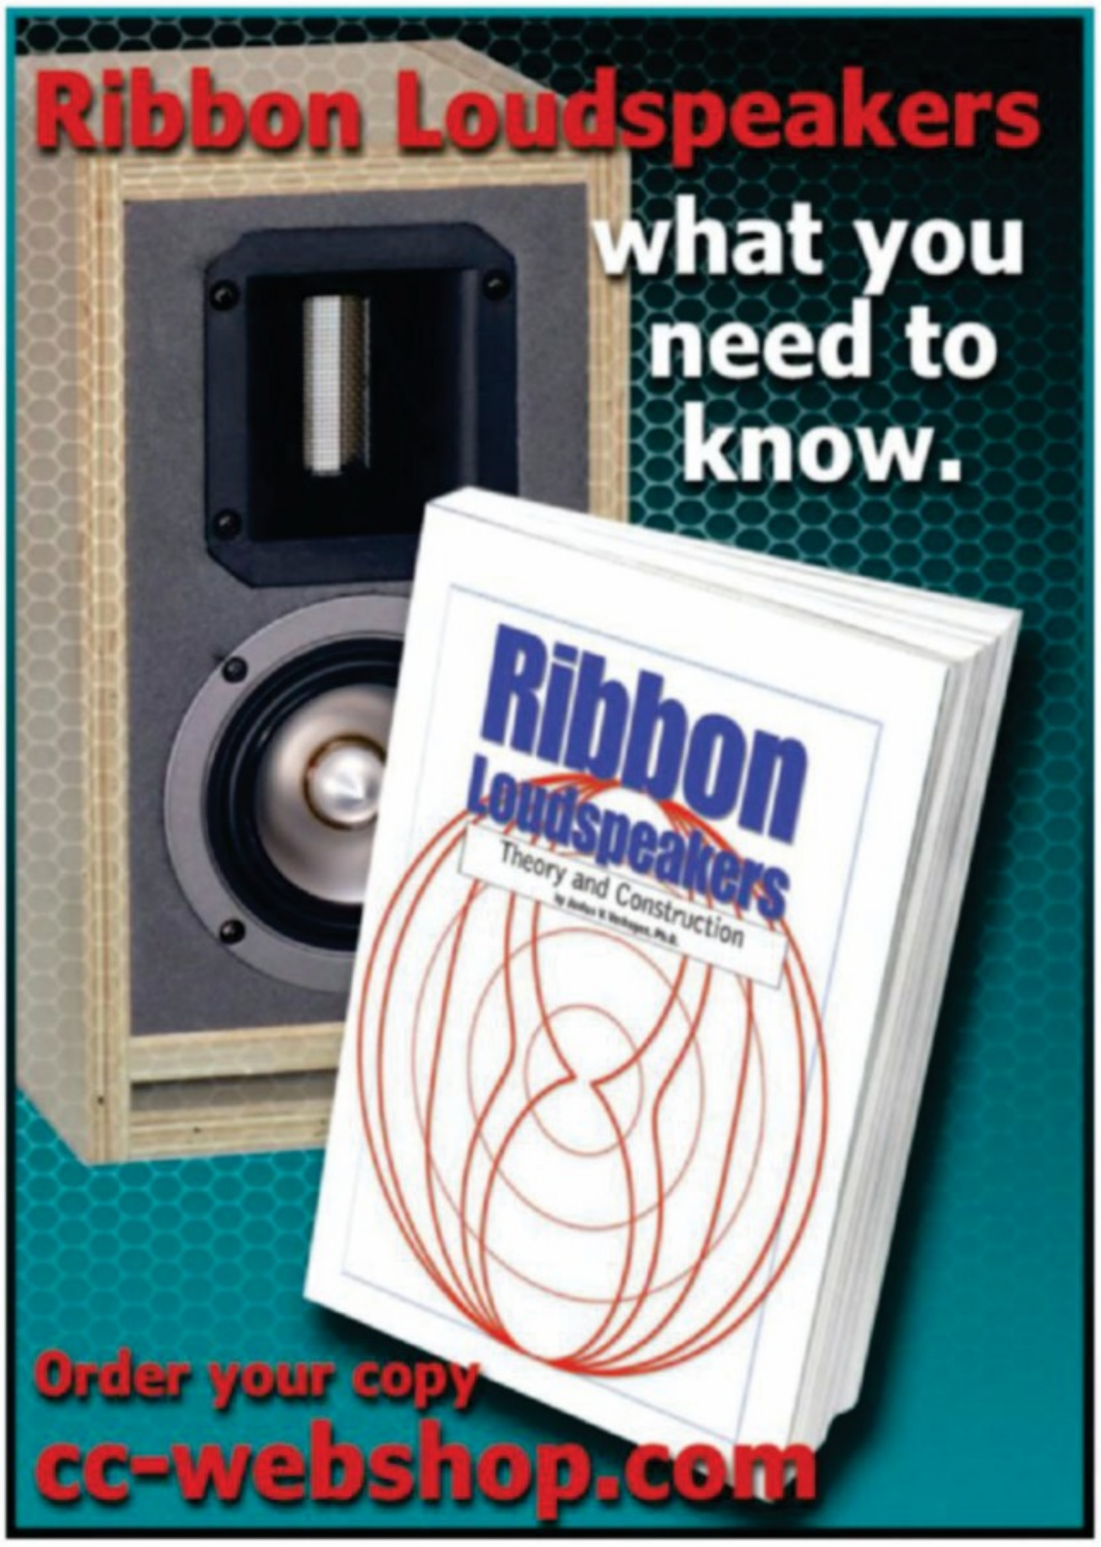

Is Your Neighborhood Acoustically Friendly?

By
Richard Honeycutt
(United States)

The field of acoustics deals with sound in buildings, sound in outdoor environments, the effects of sound on people and other animals, and many other aspects of sound. The Sound Control article printed in the March 2017 issue of *audioXpress* discussed community noise. In this article, we will look more deeply into the personal and political issues surrounding community noise.

Ever since the establishment of the Occupational Safety and Health Administration (OSHA) in 1970, there has been a growing awareness of the importance of controlling community noise. Unfortunately for those of us in the United States, this awareness is less prevalent here than in Europe and some other parts of the world. But in addition to the annoyance produced by excess noise, there are very real personal, social, and health concerns with excess community noise. OSHA has dealt primarily with hearing protection in the industry. While the importance of this focus cannot be overstressed, we must also be aware that noisy environments are associated with two other classes of negative response in humans: stress and annoyance.

Physiologically, the stress response results in release of cortisol into the bloodstream, causing vasoconstriction and consequent increase in blood pressure. The World Health Organization (WHO) has reported a correlation between long-term sound exposure above 67 dBA to 70 dBA and hypertension, and more recent studies suggest that sound levels above 50 dBA at night can increase the risk of myocardial infarction (heart attack). Other studies have emphasized the negative effects of sleep disturbance caused by sound.

An extreme stress response is observed in sensitive persons such as children or those with PTSD, in which the activation of the "fight-or-flight response" can be severe. Such effects have been documented in wild and domestic animals as well.

Psychologically mediated responses to noise include annoyance leading to irritability and moodiness, inability to concentrate, decreased cognitive performance, decreased creativity, and increased errors in work. Persons with latent mental/emotional instability may become destabilized as a result of sound-induced annoyance or stress.

Research on Sound

People can become accustomed to some sounds, with a consequent reduction in stress effects; but stress reaction to other sounds, notably gunshot sounds, cannot be ameliorated by repeated exposure. Ragner Rylander has determined that involuntary muscle contractions indicative of the "fight-or-flight response" have been observed even among policemen who regularly train with guns and are accustomed to the sound. This same research indicates that a person's attitude toward the type of sound has little if any effect upon the sound-induced stress response (see Resources).

Most studies of the non-auditory effects of sound have concentrated on annoyance more than on physiological responses. The assessment instruments used in sound-annoyance research

typically ask for yes/no answers to the question: "Do you find this sound very annoying?" The responses are highly variable from person to person, with some people finding sounds as low as 45 dBA very annoying, and others having annoyance thresholds as high as 75 dBA to 80 dBA.

Reasons for this variability include the use of hearing aids by some listeners and personal attitude toward the type of sound. Hearing-aid users often experience increased difficulty with speech comprehension in the presence of noise, as compared to those who do not use hearing aids.

Although attitude toward the type of sound does not affect the physiological response, it has a strong effect on the annoyance experienced by a given listener. For example, a classical music lover might be more annoyed by the presence of a loud rock music club next door than would a person who enjoys rock music. Or a person who is accustomed to a quiet rural environment may be more shocked by gunshot sounds than someone who grew up where gunshots and city noise are more the norm.

Usable Guidelines

In order to provide usable guidelines on sound levels for land-use compatibility decisions, researchers and governmental agencies have focused on sound-level limits that represent an average response in terms of annoyance, knowing that there will be listeners who will be annoyed even at the permissible sound level. These have resulted in numerous community noise regulations. Typically, these specify a maximum permissible sound level at a property boundary. Unlike many urban noise regulations that specify maximum sound level at the property boundary, the Environmental Protection Agency (EPA), the US Army, and the National Rifle Association (NRA) guidelines specify sound on a nearby owner's property, not simply at the boundary of the property on which the sound is being generated.

On the other hand, outdoor noise levels are not entirely under human control. One community noise ordinance passed in the late 1970s specified a maximum level of 55 dBA. When engaged to investigate a noise complaint, an acoustical consultant determined that the cricket noise at the time of measurement (a July evening) exceeded the maximum. Traffic noise, especially that associated with vehicles having improper mufflers, may easily exceed the legal maximum. If a property owner initiates a noise complaint against the owner of a music club, legal noise limits that are lower than the normal ambient levels from natural and traffic sources make the noise regulation unenforceable and essentially useless. In all too many cases, investigation of noise complaints is left to police personnel who have little or no training in making SPL measurements, and whose instruments may not be regularly calibrated (see **Photo 1**).

Controlling community noise is best commenced at the design stage for the possible generators of the noise, and also for residential and commercial buildings in which people must be protected from noise annoyance. Unfortunately, all too often community noise problems are addressed only after they have begun to cause annoyance. In either case, preventing or correcting a community noise problem often necessitates careful measurements. These measurements must be made using appropriate equipment, and in some cases, may require a schedule of measurements spanning hours, days, weeks, or

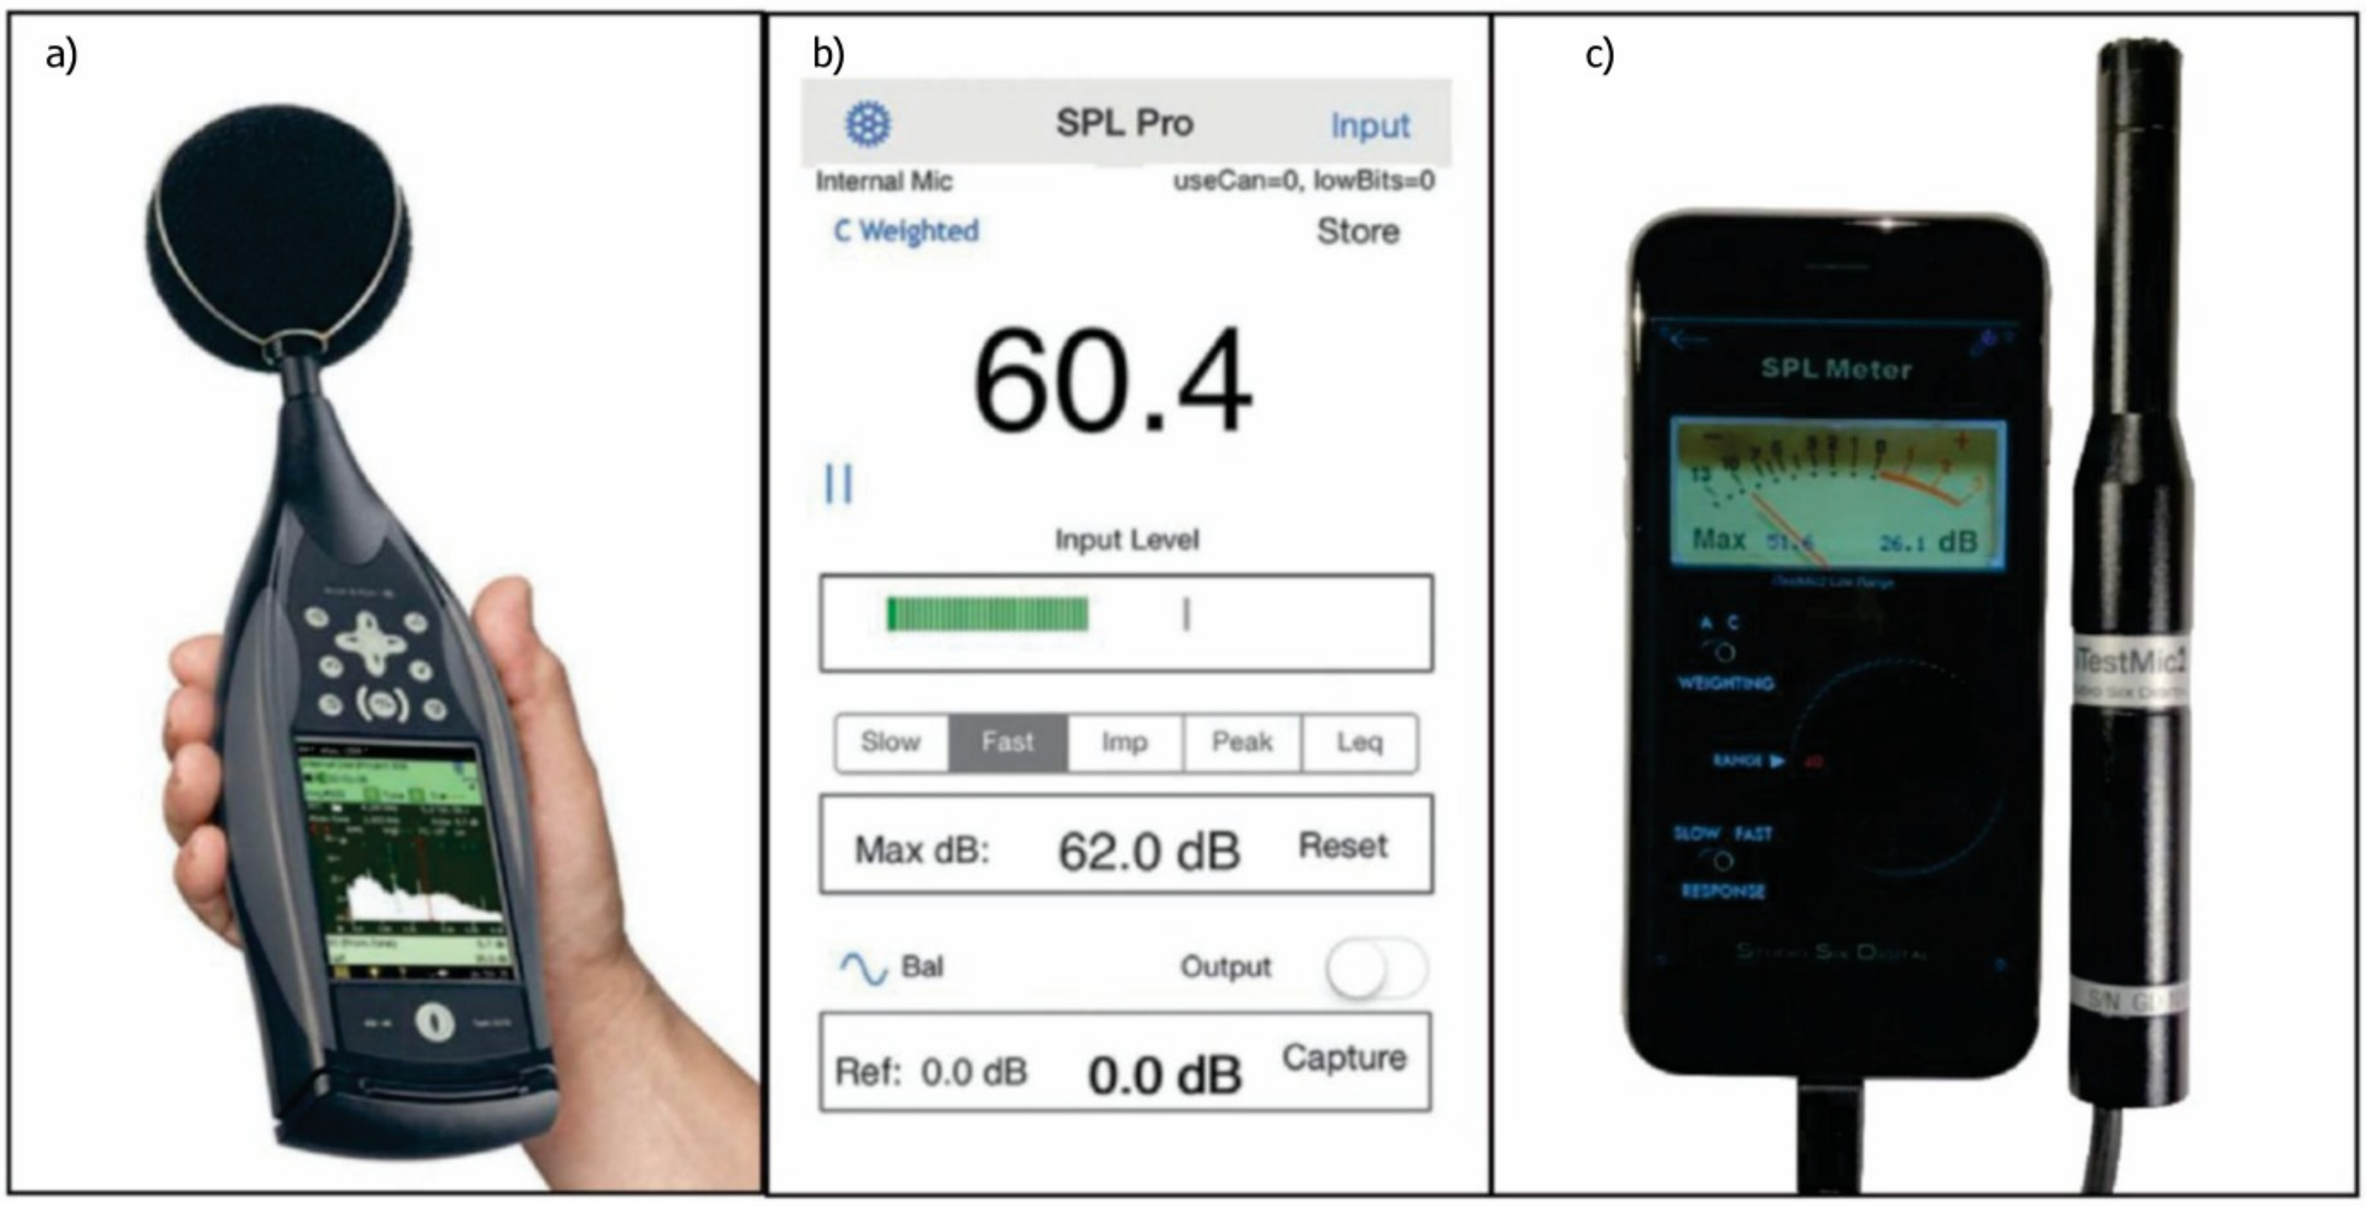

Photo 1: The Brüel and Kjaer Model 2270 (a), an SPL Pro-App (b), and the Studio Six Digital new Ligntning iTestMic 2 microphone (c) are all examples of ways to measure sound pressure levels (SPLs).

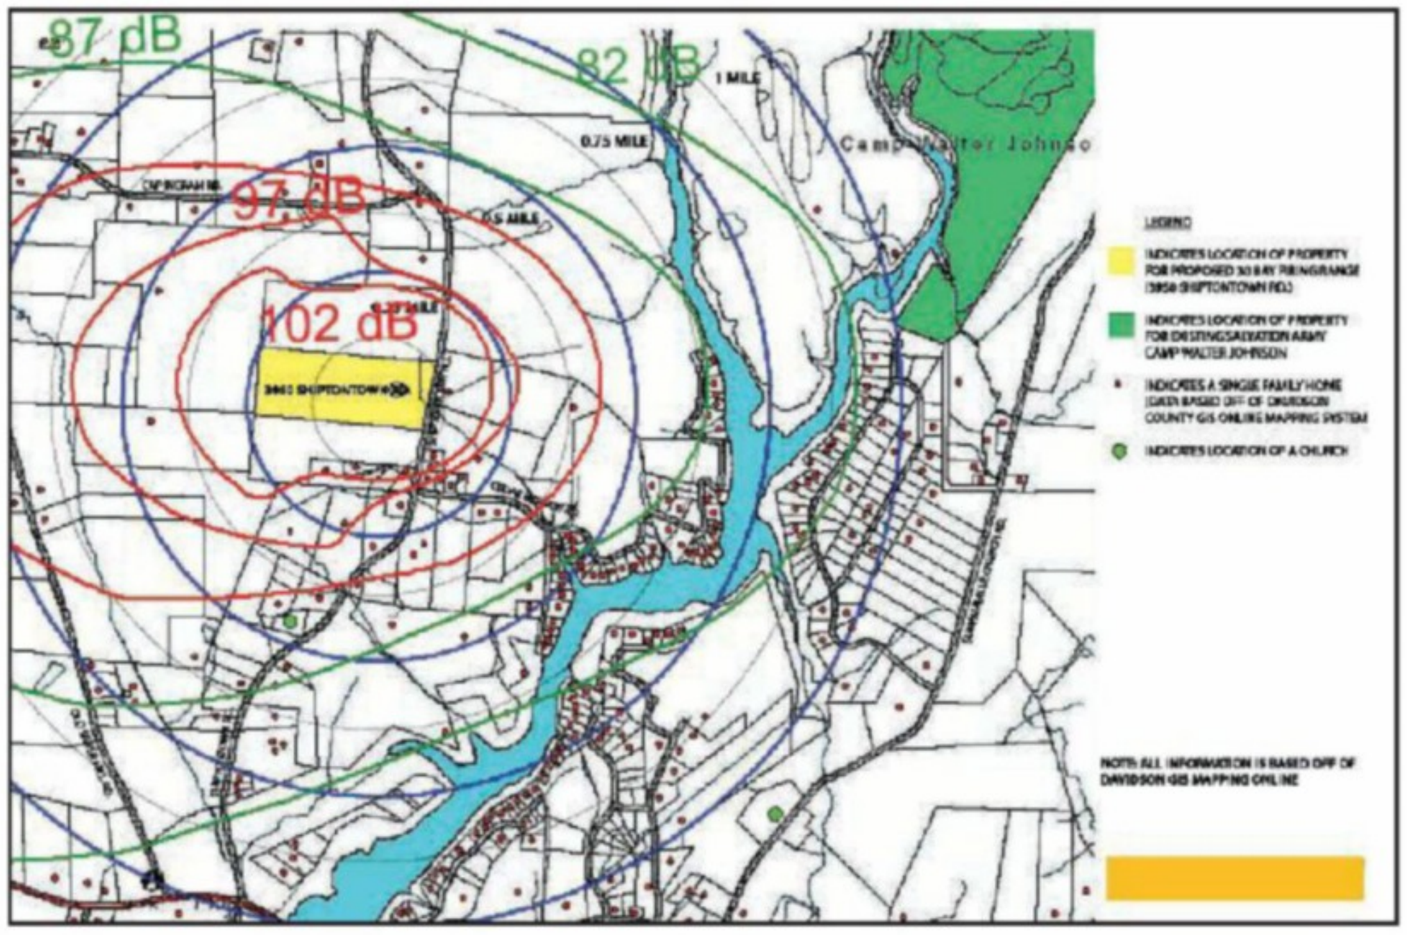

Figure 1: Calculated noise contours for a proposed shooting range

even months, to account for diurnal, seasonal, or other time-related effects of noise generation and propagation.

Practical Noise Examples

Two community noise cases in which I have been involved may prove interesting and instructive. In the first, I was engaged by a homeowners' group to oppose the grant of a special-use permit for a combination indoor/outdoor shooting range in their neighborhood. Since the range did not yet exist, no relevant noise measurements were possible, I used SARNAM, a software developed by the US Department of Defense, to calculate the expected noise exposure at the nearest property-owner's location. The resulting noise map (see **Figure 1**) indicated a significant number of residences within high-noise-level contours. By contrast, the attorney for the proposed shooting range owners engaged a consultant in forensics, who attempted to make his point that the range would not cause problems by measuring the SPL from a shotgun blast at an unmeasured distance, using an uncalibrated sound level meter. Since weather conditions affect sound propagation outdoors, his measurement would at best have been valid only for the conditions of temperature, humidity, and presence/absence of temperature inversion that may have existed at the time of any measurement.

The local governing board decided that they could not understand the testimony and, therefore, chose to rule in favor of the special use permit by ignoring the requirement that the value and usability of the adjacent property be considered in making their decision. However, the owners' group appealed the decision, and the court ruled that the board had acted illegally and the permit must not be granted.

This example illustrates the need for governmental officials to pay attention to good acoustical science in making laws, regulations, and zoning decisions. Actually, the forensics consultant, though clearly acting outside his area of expertise, measured pretty much as most local police personnel would have, and the board demonstrated the effects of ignorance of acoustical principles that are involved in making a legitimate decision.

In the second case, I provided testimony for a sawmill owner who had operated the sawmill for many years and had not known that he needed a special use permit. A neighbor who had previously operated a motorcycle race course on his property opposed the granting of the sawmill permit. My job was easier in this case.

Since the sawmill was in operation, I could make measurements at the plaintiff's property boundary. These involved four pieces of equipment: the sawmill itself, a wood processor that turns unusable logs into firewood, a log loader having a three-cylinder diesel engine, and a skid-steer used for loading firewood onto trucks. Surprisingly for many people hearing that the operation involved a sawmill, the mill itself was not the worst noisemaker. Being a small bandsaw mill, it produced less noise than the average lawnmower.

I set up equipment to measure the noise levels in dB(A) at the southern property line of the logging company. (See the red arrow just left of center, about one-third of the way up from the bottom in **Photo 2**.) The measurements were performed using an NTi Instruments Model ML1 analyzer with firmware crossgraded to Model AL1. The Sabine SQ-1001 measurement microphone was calibrated

Photo 2: An aerial photo of the properties discussed in the Sawmill Case

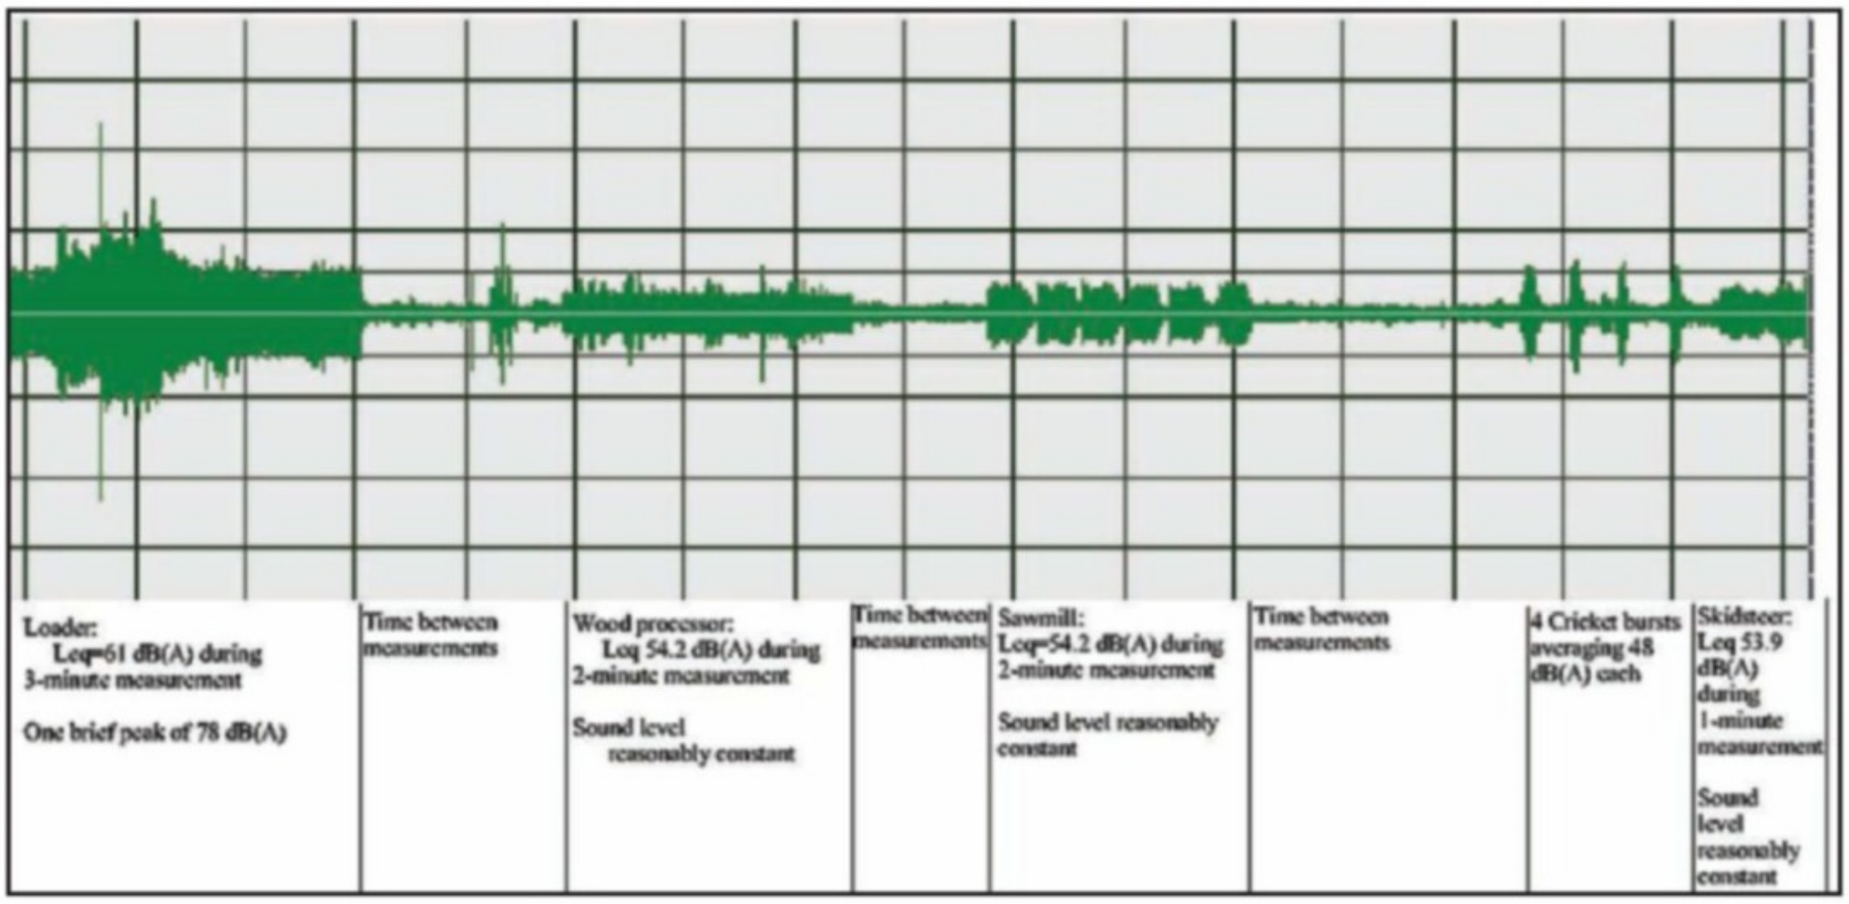

Figure 2: A .wav file of noise measurements

one month before the measurement, using a Quest QC-10 calibrator having NIST-traceable calibration.

The measurement microphone was mounted on a tripod at approximately ear height for a standing adult. Instantaneous levels using the analyzer's FAST setting were observed while the analyzer computed the L_{eq} for the full measurement time. (L_{eq} is the time-averaged sound pressure level.) The longest measurement time (3 minutes) was used for the loader, since it had the greatest sound-level variation with time. For the sawmill and wood processor, 2-minute measurements were made; and for the skid-steer, a 1-minute measurement.

Figure 2 contains a green trace corresponding to the recorded waveform of the measurement session. The horizontal scale indicates time: about 1 second per division. The height of the trace at any given time indicates the output voltage of the measurement microphone. (Of course, this voltage was automatically converted to sound pressure levels in dB(A) inside the analyzer. Thus, a trace that is two divisions high is not twice the decibel level of a trace one division high, although it represents twice the voltage from the microphone. Twice the voltage corresponds to a 6-dB increase.)

A cricket near the measurement microphone contributed an interesting feature to the .wav file, as you can see from Figure 2. Only the loader produced sound levels that varied significantly with time. There was one brief peak at which the level reached 78 dB (roughly the level of average street traffic in a city). The sound levels produced by the sawmill, the wood processor, and the skid-steer were essentially constant with time, with an L_{eq} just under 55 dB(A), which was the maximum legal limit.

In the hearing before the board, the neighbor attempted to illustrate how loud the sawmill noise was by using an iPhone recording of two people talking nearby, followed by a pause, then the sound of the sawmill. He thought the comparison between the loudness of the talking and the loudness of the sawmill would provide something akin to calibration. Of course, he failed to take into account the effect of the iPhone's automatic level control, which adjusted the gain of the iPhone higher for the sawmill, whose level was in fact lower than that of the talker. In this case, the board recognized the difference between careful science and emotionally based opinions, and granted the permit.

Additional Noise Issues

Aside from considerations of noise ordinances, zoning permits, and the need for enforcement personnel to be properly trained and equipped, there is another noise issue attracting growing attention: noise inside restaurants. Some groups are attempting to use crowd-sourced data on restaurant noise to help rate restaurants for quality.

The problems here are essentially the same as those illustrated in the cases recounted earlier. The usual measuring instruments are cell phones using apps whose accuracy varies greatly (see the Sound Control articles in November and December 2015 issues of *audioXpress*), and which have not been calibrated. And the operators may not be aware of a possible need to average noise measurements or how to perform the averaging, or even how to select an appropriate location for the measurement. As usual, the GIGO principle applies! ax

Resources

R. Honeycutt, "Community Noise Effects and Measurements," *audioXpress*, March 2017.

———, "Audio Tools from Studio Six Digital, "*audioXpress*, December 2015.

———, "Smartphone Apps for Sound Measurement, "*audioXpress*, November 2015, https://www.audioxpress.com/article/smartphone-apps-for-sound-measurement.

R. Rylander, "Noise, Stress, and Annoyance," *Noise and Vibration Worldwide*, June 2006.

Can Audio Experts Learn from High-Speed Amplifier Developments?

High speed amplifiers offer very high gain bandwidth, high slew rates, an excellent speed-to-current-consumption ratio, and very low noise characteristics. In his first article under the Audio Scope heading, our newest columnist, Michael Steffes, addresses some of the advanced characterization and application tricks coming out of the high-speed groups that may not have filtered out into the audio world.

By
Michael Steffes
(United States)

In the course of my 32 years of high-speed amplifier developments, we often designed for the next-generation test equipment. The irony is that there was usually nothing we could buy to perform some of the necessary parametric measurements. Hence, a great deal of advanced characterization and application tricks came out of the high-speed groups that may not have filtered out into the audio world. For instance: How do the signal harmonics coming into an ADC combine with the ADC-generated harmonics for a combined Fast Fourier Transform (FFT) output harmonic spectrum?

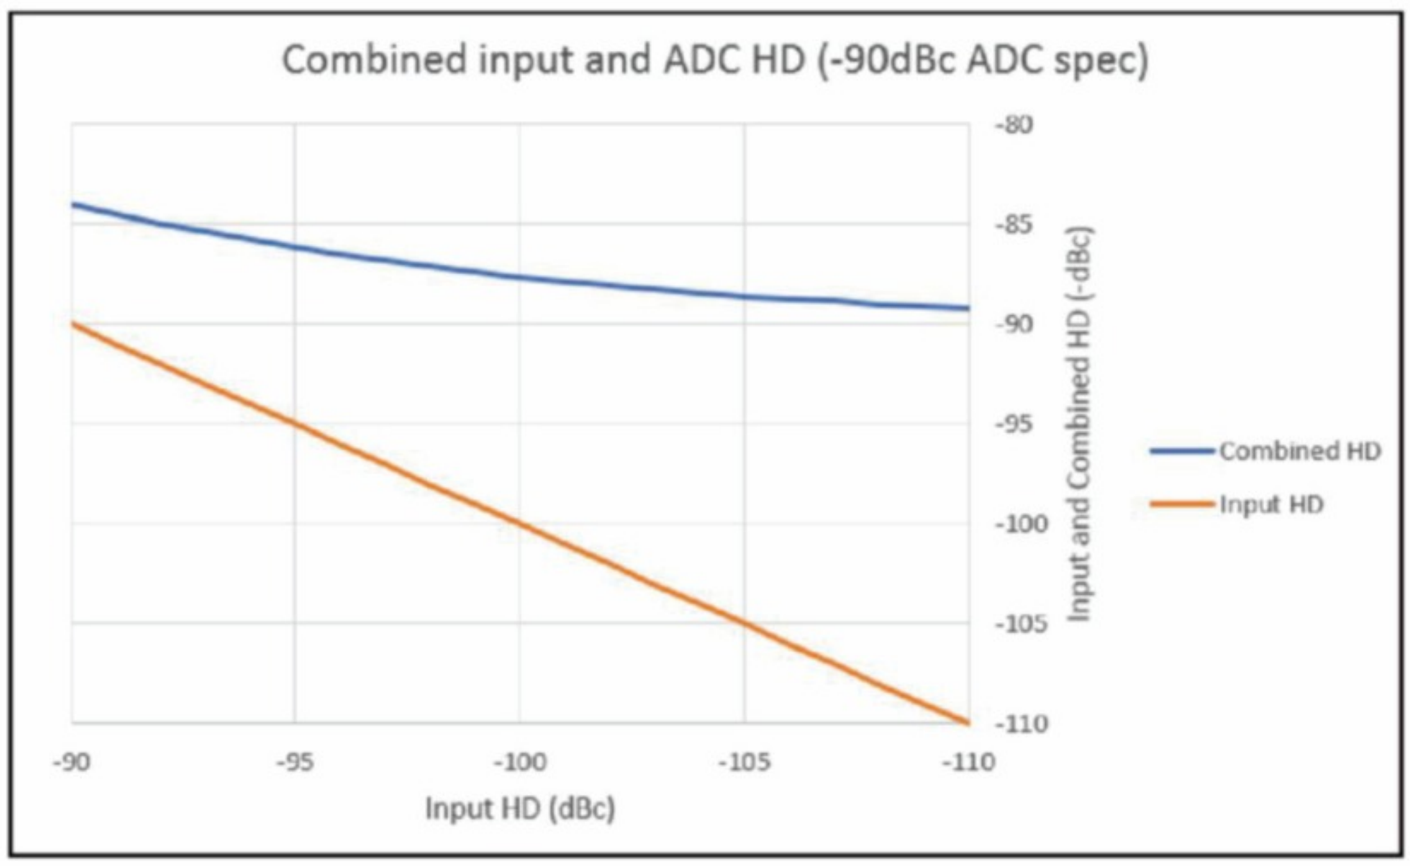

Figure 1: Combining input signal and ADC only harmonic distortion -dBc for a combined result

This kind of fundamental question leads into the allowed input harmonic levels to meet an allowed degradation through the ADC. To combine distortion terms, you can work in absolute voltage swing levels for the harmonics, more commonly the -dBc specifications, or the negative of that specification—the Spurious Free Dynamic Range (SFDR) for the highest harmonics. Working in -dBc specifications will yield an expression that applies more globally.

Some of the legacy literature simply repeats the total harmonic distortion (THD) equation shown as Equation 1. This is converting the -dBc specification for harmonic distortion terms (second and third usually dominant) at the ADC input pins and the ADC -dBc specification into a power specification, adding those and taking the square root, then converting back to a -dBc specification.

$$HD_{system} = 20 \times \log \sqrt{\left(10^{\frac{HD_{amp+filter}}{20}}\right)^2 + \left(10^{\frac{HD_{ADC}}{20}}\right)^2} \quad [1]$$

So, for instance, if both the incoming signal and the ADC have -90 dBc HD2, this equation predicts -87 dBc combined. But really?

Resources

"Low Noise, Precision, 150MHz, Fully Differential Amplifier," Texas Instruments THS4551 Datasheet, revised July 2017, www.ti.com/lit/ds/symlink/ths4551.pdf.

"High Speed, +/-0.1µV/°C Offset Drift, Fully Differential ADC Driver," Analog Devices ADA4945-1 Datasheet, www.analog.com/media/en/technical-documentation/data-sheets/ADA4945-1.pdf.

The ADC is accepting the input harmonic distortion as a signal, then generating its own harmonic at the same frequency to combine with the input harmonic. Equation 1 is essentially assuming those two terms add in quadrature (the assumption of Equation 1 is like you think of little sine waves at the harmonic frequencies -90° out of phase). There is no reason for that assumption to be accurate—and, in fact, the two harmonic sine waves add in phase as a V_{pp} addition—they stack up on top of each other increasing the voltage swing at the harmonic frequency.

Correcting Equation 1 to add voltages instead of powers, gives Equation 2:

$$HD_{system} = 20 \times \log\left(\left(10^{\frac{HD_{amp+filter}}{20}}\right) + \left(10^{\frac{HD_{ADC}}{20}}\right)\right) \quad [2]$$

Putting two -90 dBc terms into Equation 2 will give a -84 dBc combined result coming out of the ADC in the FFT—a 6 dB degradation. But even more generally, if we ramp the input harmonic distortion down how does the net result behave. **Figure 1** shows this where both ADC and input signal start out at -90 dBc and the harmonic distortion at the input pins is ramped down going to the right. The combined result slowly approaches the ADC only number as the input harmonics go down.

In fact, to degrade the ADC by no more than +1 dBc, Equation 2 can be solved to show a 19 dBc better input harmonic distortion is required. This plot in dBc looks the same for any starting point on the ADC—from any ADC only spurious specification tone by tone, the input signal would need to have a harmonic distortion term that is a least 19 dB below that ADC term to yield <1 dB degradation. Or, Equation 2 can be used with the actual levels to get the combination. A common target in the literature has been 10 dB better than the ADC coming into the input pins. That is fine, but note that will degrade the ADC only specification by 2.4 dBc.

This much more extreme requirement explains the much lower harmonic distortion targets in the most recent Fully Differential Amplifier product introductions such as that shown in **Figure 2**. Similarly, <-145 dBc harmonic distortions show up in the newest ADA4945-1 April 2019 FDA introduction (see Resources).

The next two technical briefs will show bench data reconfirming this V_{pp} additive distortion effect for a recent two-stage amplifier product then go on to describe how you might actually take the data of Figure 2 (hint, it is not an Audio Precision solution). ax

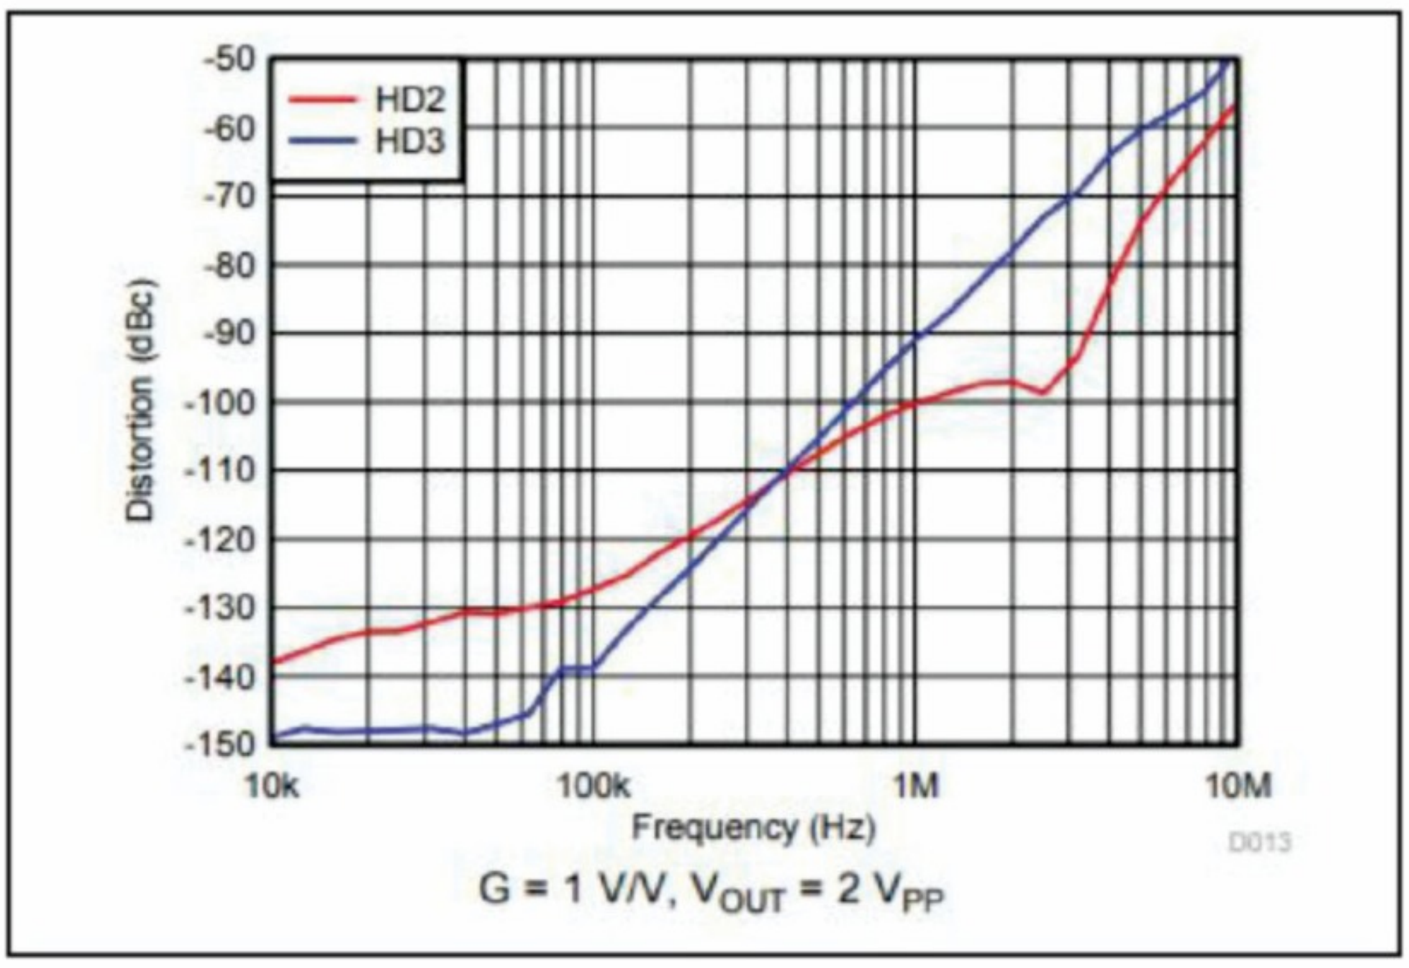

Figure 2: An example of very low harmoic distortion performance for the recent THS4551 FDA device

Two Mu Valve Topologies

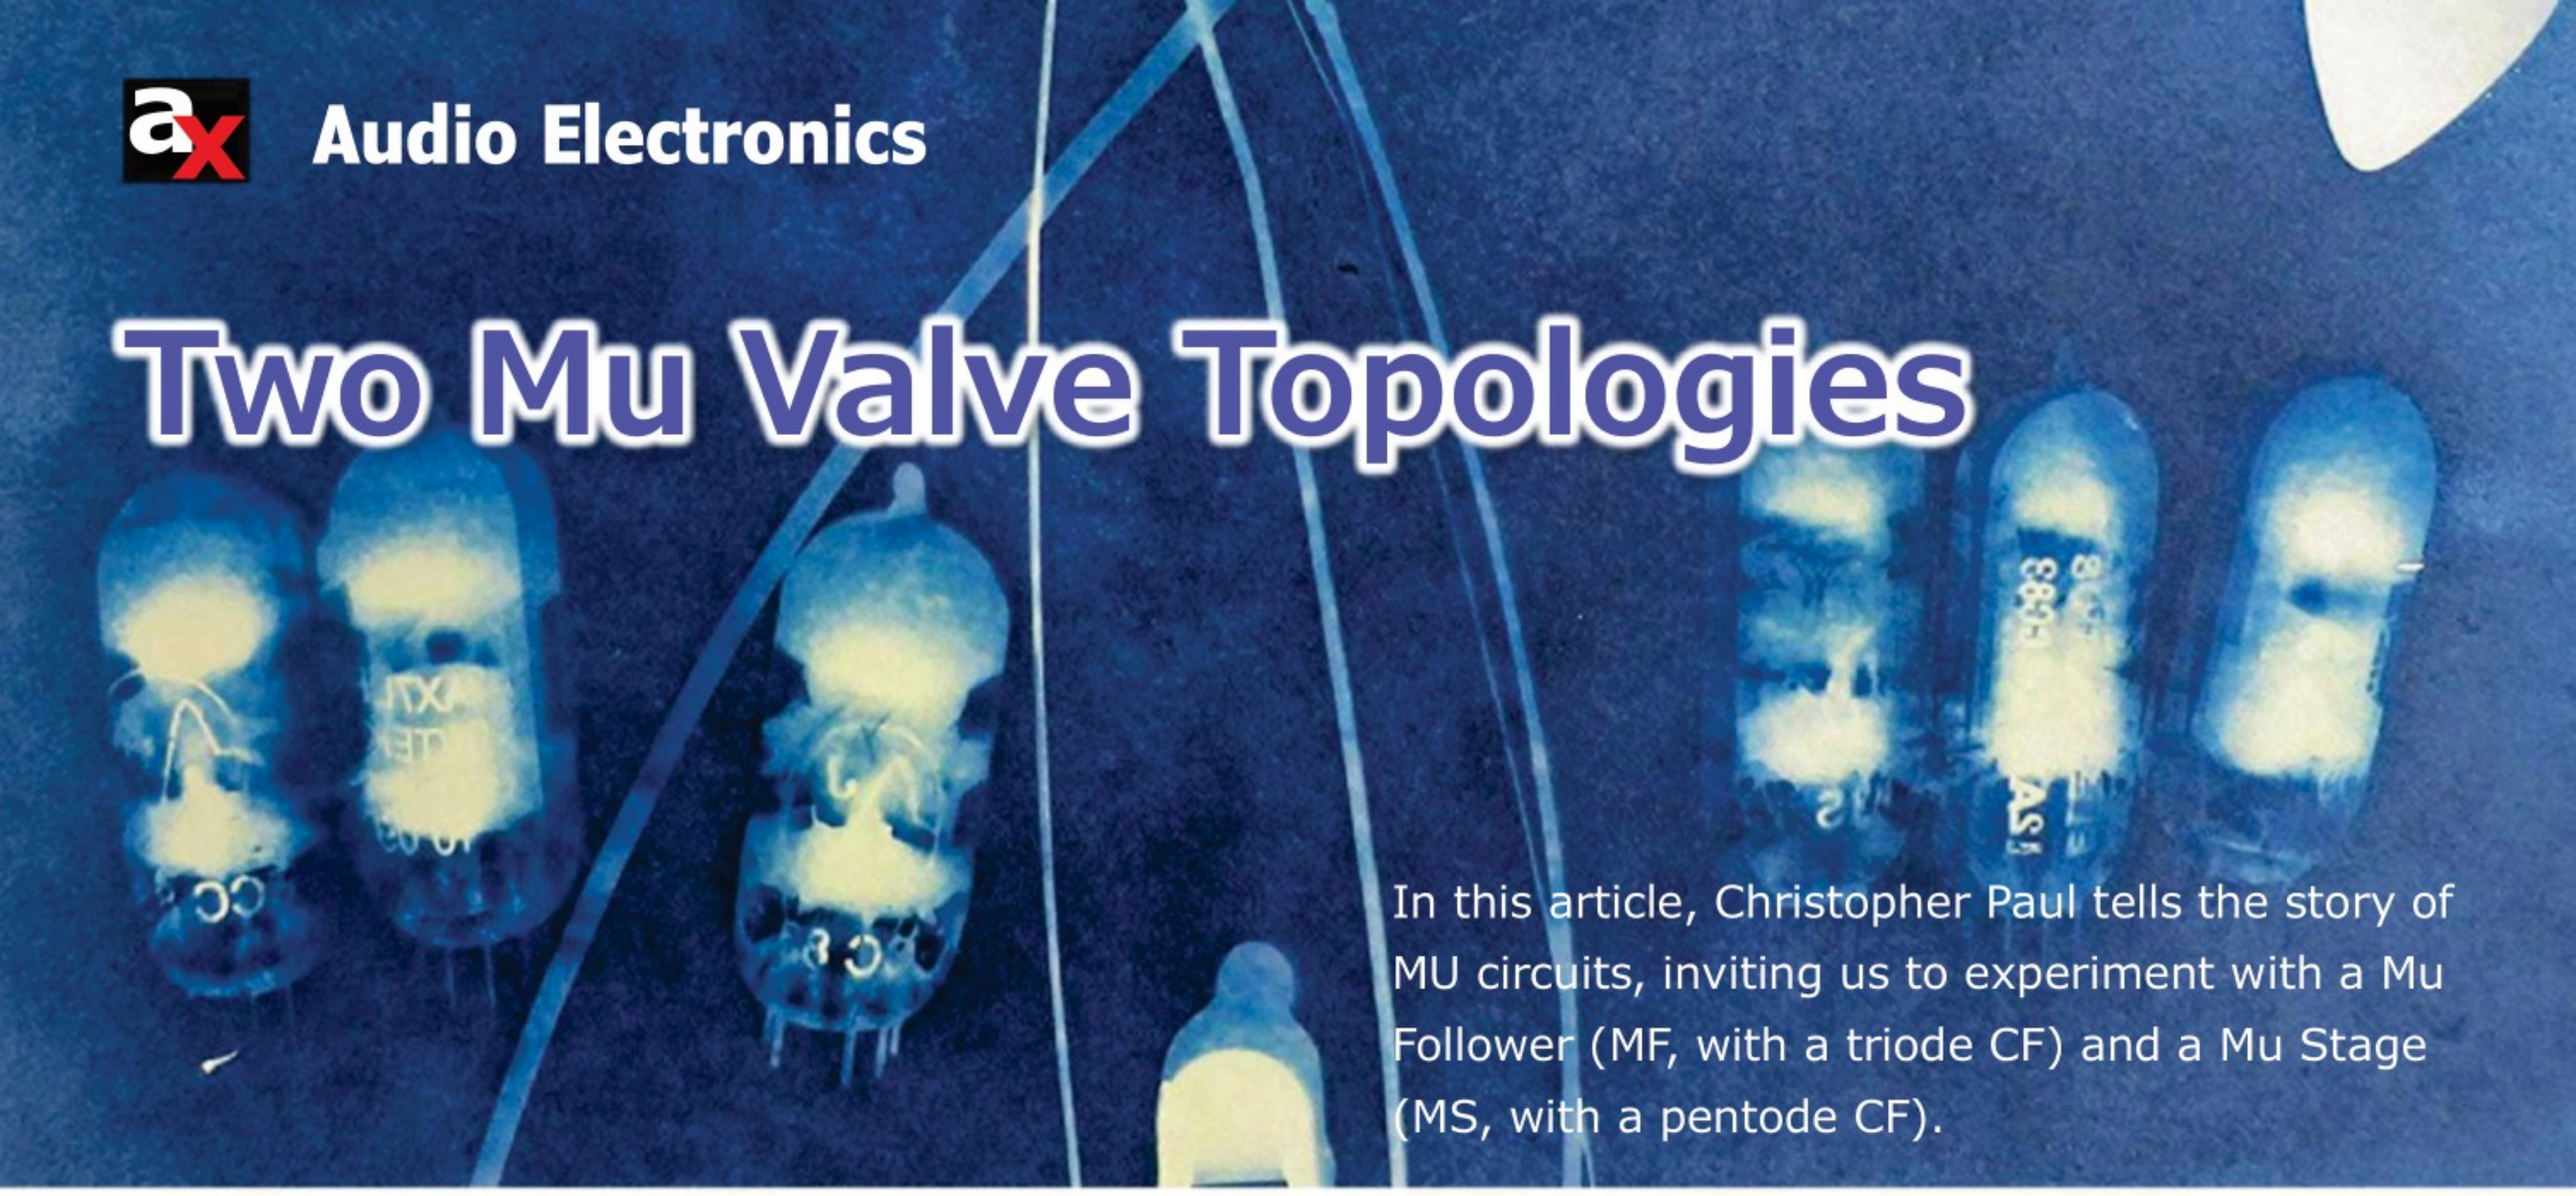

In this article, Christopher Paul tells the story of MU circuits, inviting us to experiment with a Mu Follower (MF, with a triode CF) and a Mu Stage (MS, with a pentode CF).

By
Christopher Paul
(United States)

Mu circuits are among the more interesting simple hollow-state voltage amplifiers. These two-valve topologies achieve high linearity/low distortion, admirable power supply rejection ratio (PSRR), and low output impedances which approach those of cathode followers—all without the application of valve-to-valve feedback.

In this article, I'll use "valve" to denote a single triode or pentode and "tube" to refer to an encapsulated component that contains one or more valves. In general, a Mu circuit consists of a common cathode triode voltage amplifier (VA) whose anode is coupled to the grid of a cathode follower (CF). If the description of the circuit ended there, there would be nothing special about the topology. But the CF's cathode is coupled to the side of the VA's anode load which would otherwise be connected to the B+ supply. This has two salutary effects. Because the CF's gain is approximately unity and its grid and cathode are connected to opposite sides of the load, the AC voltage across that load is greatly reduced. This increases the load's effective AC impedance as seen by the VA triode. ("Bootstrapping" is the common term for this type of arrangement.) And triodes typically exhibit their highest linearity (and voltage gain) when that impedance is maximized. The second effect is the CF's buffering of this load from B+ power supply noise, enhancing the circuit's PSRR.

Let's jump straight to some exemplary Mu circuits: a Mu Follower (MF, with a triode CF) and a Mu Stage (MS, with a pentode CF). All the total harmonic distortion (THD) tests I'm about to discuss employ my Simple Sineman sine wave generator (audioxpress, November 2013)[1] as a signal source, with a 3 Vrms maximum output and a THD better than -130 dBc.

Exemplary Mu Circuits

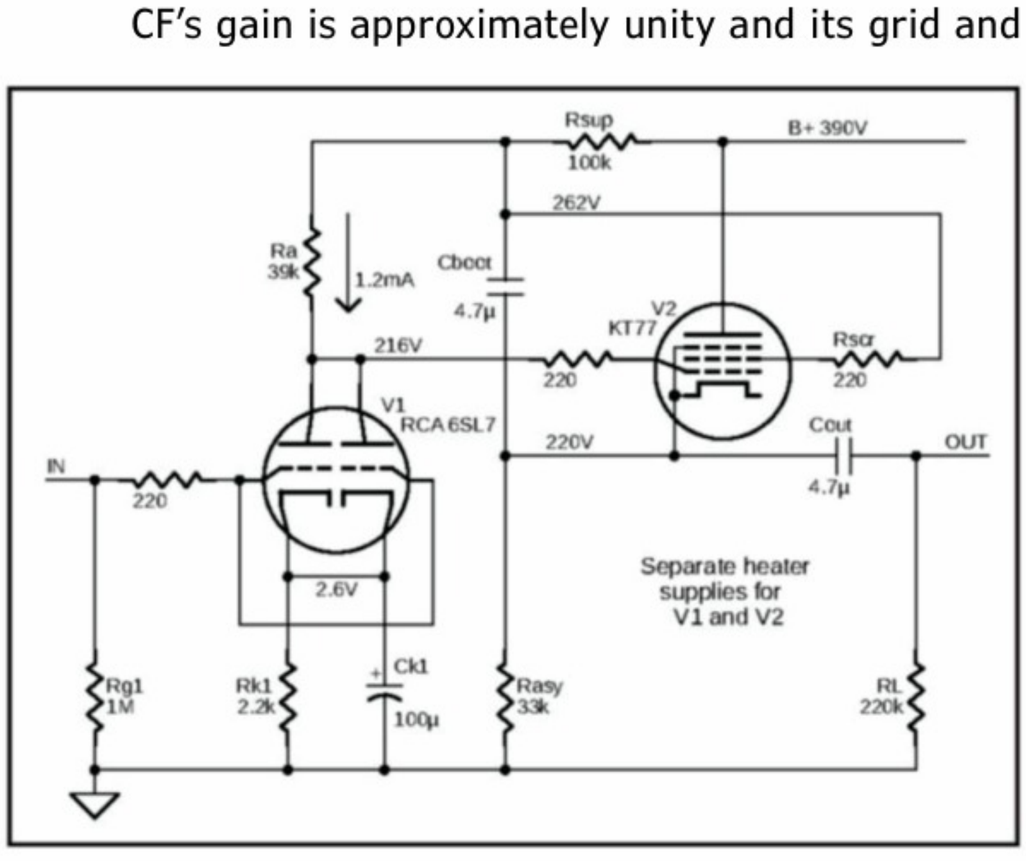

Figure 1: A new variant of the original Mu Stage introduced in 1993, this exemplary Mu Stage appeared in the May 2019 edition of *audioXpress*.

Figure 1 depicts an MS, which was recently discussed in the May issue of *audioXpress*[2] (the DC voltages shown reflect the measurements of the circuit I built.) A THD of 0.02% at 30 Vrms was claimed for this circuit, impressive performance indeed! I purchased two used RCA 6SL7s and two new JJ Electronics KT77s and built and tested this circuit using combinations of these units. The KT77s performed similarly, but one of the 6SL7s contributed noticeably more distortion than the other. However, circuits with the "good" 6SL7 exhibited an even lower (0.0080%) THD (counting harmonics 2 through 9 with a 2400 Hz fundamental.) I wondered whether I could replace the KT77 with a physically smaller, lower power triode and obtain a

similar THD. Sure enough, the MF shown in **Figure 2** achieved a 0.0105% THD with either of the two Tung Sol (Russian-made) 6SN7s I bought.

A summary of the performance of these two designs appears in **Table 1**. The Mu Stage has 18 dB better PSRR and a 1% THD voltage swing of 268 V_{PP} in comparison with 194 V_{PP} for the Mu Follower. (Note however that an optimally biased cathodyne can accept at most a B+/2 V_{PP} swing before saturating or being cut-off, and the Mu Follower provides that—larger swings are not useful.) But the MS also requires one additional tube (the 6SN7 offers two CFs, one for each stereo channel), dissipates more than twice the power, and the complete circuit costs approximately $40 US more (4.7 µF, 600 V capacitors are expensive!)

The remaining characteristics of the two Mus are similar, except for one thing. It's striking that the circuits' distortion performances are practically identical even though the effective mu of the Mu Follower CF is under 20 while that of the Mu Stage is almost 110. In this circuit at least, even a low to moderate Mu triode suffices as a CF for low distortion operation; the pentode CF's higher gain is unnecessary.

A Brief History of Mu Followers and Mu Stages

Some people think that interest in the audio applications of Mu circuits began with the appearance of the Mu Stage[3] in the 1990s. But in fact, it started in the 1980s with the introduction of the Mu Follower (as indicated in a letter published in *The Audio Amateur* in 1985)[4]. Morgan Jones is among those who mis-attributes the inception of this enthusiasm in one of his popular books. But he does cite a reference by JW Horton to a Mu circuit from 1933[5]. Unfortunately, I have been unable to obtain a copy of that reference. But you can find US Patent 2,661,398 online granted in 1953, which has a schematic of a Mu Follower on its front page, although it was neither so named nor were all its benefits disclosed.

The same circuit was then "re-invented" in the mid-1980s and presented in the Letters column of *The Audio Amateur*, a predecessor of *audioXpress*. (As *audioXpress* Technical Editor Jan Didden once wryly observed, the "ancients" are always stealing our ideas from us.) The active elements of this late comer were twin triodes in a single tube. You can find its schematic in **Figure 3a**. (Note that this disclosure does not employ the resistor R_{asy}; R_{asy} is used instead in a variant, which I shall describe next.)

A later letter to *The Audio Amateur*[6] presented a further discussion of the circuit and assigned it

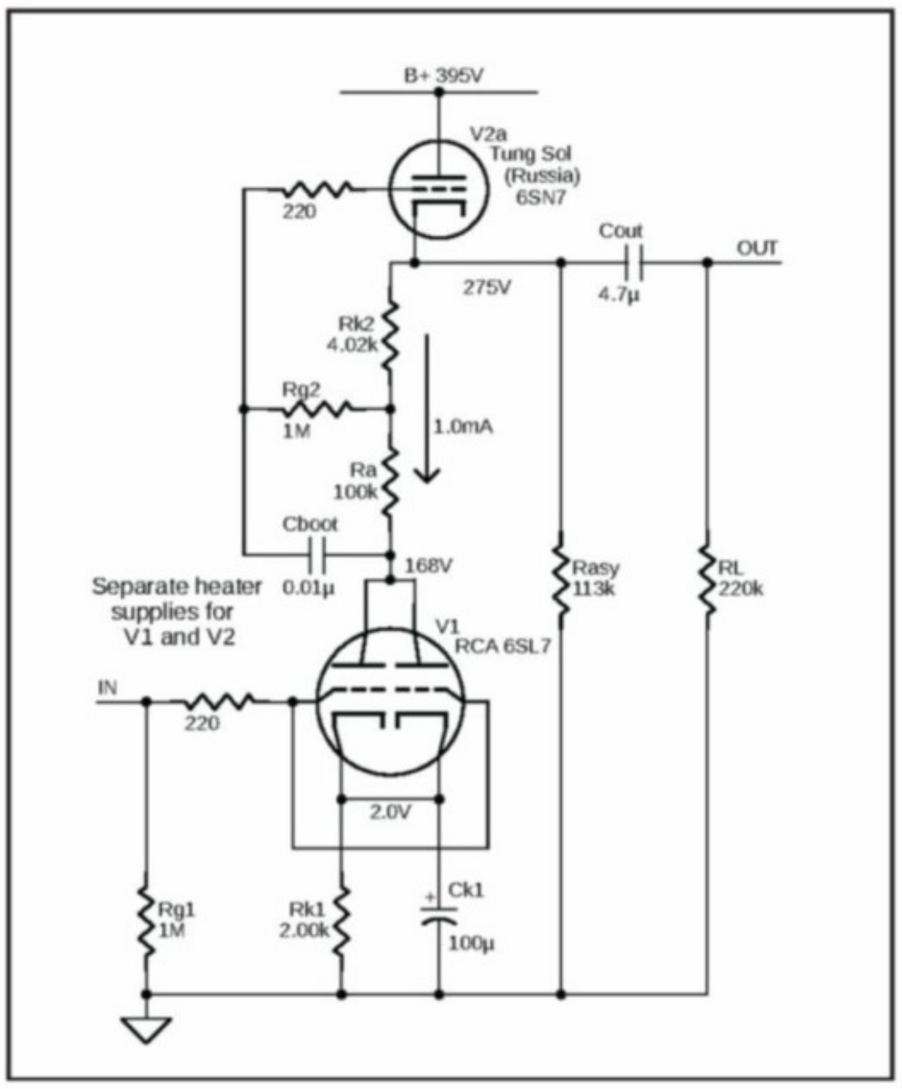

Figure 2: This exemplary asymmetric Mu Follower exhibits a low distortion virtually identical to that of the Figure 1 Mu Stage.

the moniker "Mu Follower." Sometime in the late 1980s or early 1990s, Rickard Berglund realized that there could be an advantage to using different triodes for the circuit's two valves or biasing twin valves at different currents. Perhaps the CF could be optimized for driving moderate loads and the VA for gain, linearity or otherwise. And thus, the Asymmetric Mu Follower was born (see Figure 3a again, this time using R_{asy} to draw extra bias current from the CF.)

	Mu Stage (Figure 1)	Mu Follower (Figure 2)
THD, Harmonics 2-9	0.008%	0.0105%
Effective mu of cathode follower	109.4	18.5
PSRR	41.7 dB	23.4 dB
Output impedance	382 Ω	600 Ω
Bandwidth	-3 dB at 160 kHz	-3 dB at 124 kHz
Output voltage swing	268 V_{PP} at 1% THD	194 V_{PP} at 1% THD
	273 V_{PP} at 3% THD	202 V_{PP} at 3% THD
	291 V_{PP} at 10% THD	233 V_{PP} at 10% THD
Gain	36.6 dB	36.2 dB
Cost (stereo operation, CF tubes, plus Digi-Key boost capacitors)	$61.88	$23.78
Number of tubes for stereo operation	4	3
Power, B+ and heater supplies, stereo operation	26.88 W	10.29 W

Table 1: Performance summary comparing the Mu Stage to the Mu Follower

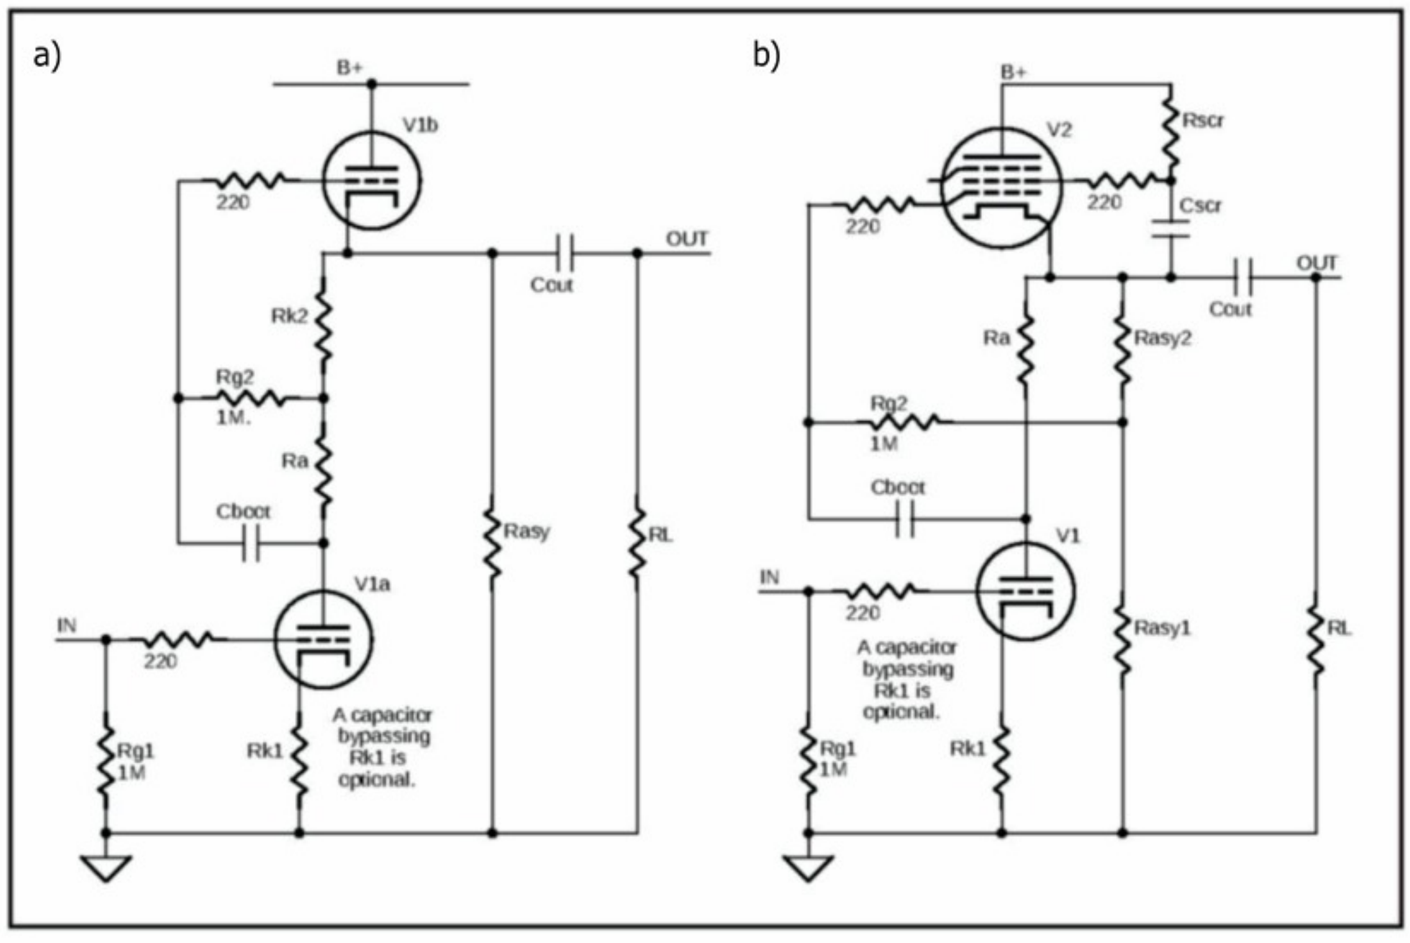

Figure 3: The (a) circuit was disclosed in 1985 in the "Letters to the Editor" column of *The Audio Amateur*. The (b) circuit is the original Mu Stage, presented in 1993 in *Glass Audio*.

Continuing with this idea of asymmetry, in 1993 Alan Kimmel suggested a pentode in place of the triode CF. He noted its generally higher grid-to-cathode gain. Since the multiplicative bootstrapping factor for the VA anode's effective load impedance is proportional to 1/(1 – gain), a pentode CF could generally make that impedance appear larger than a triode could. **Figure 3b** shows such a circuit, which was dubbed a "Mu Stage." Of course, a screen grid network had to be added for the pentode. Recently an MS variant was introduced (see Figure 1). It achieves bootstrapping by introducing resistor R_{sup} and driving its junction with R_a through a C_{boot} requiring a much higher capacitance than that of Figure 3. It should be noted that either a pentode and its screen grid network or a triode without it could be employed as a CF in all three topologies.

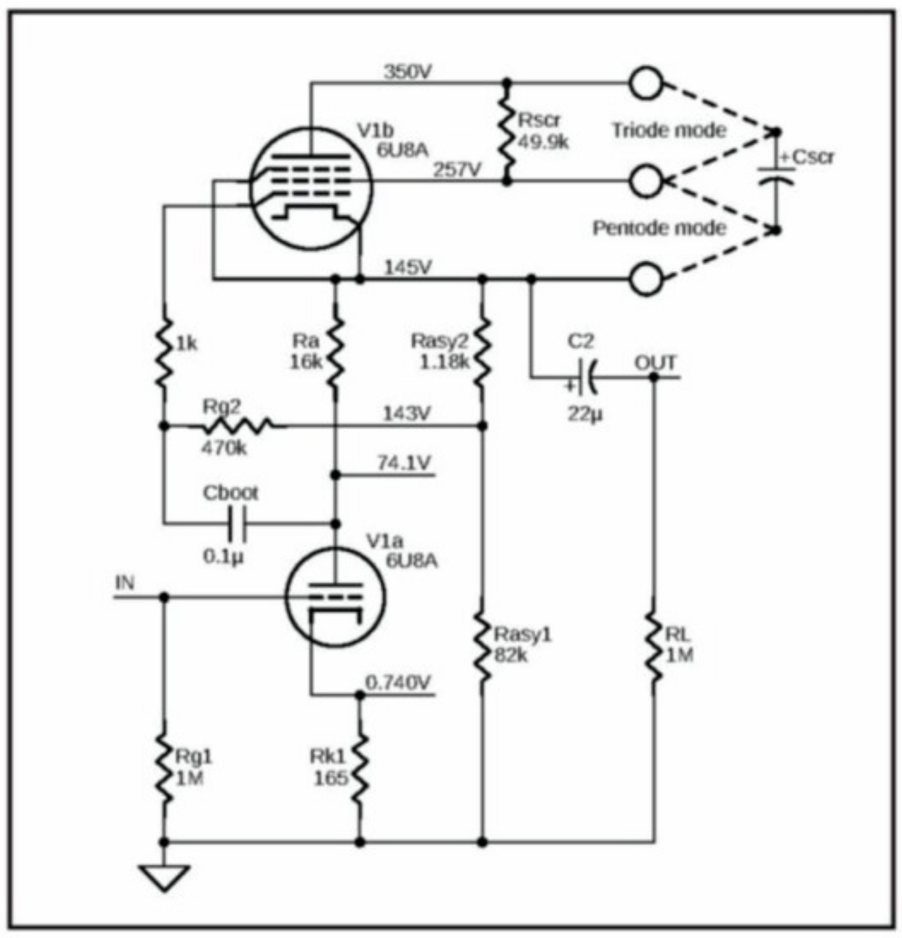

Figure 4: This circuit was characterized for both triode (MF) and pentode (MS) mode operation by simply appropriately positioning the screen grid capacitor C_{scr}.

Before launching into a consideration of the comparative merits of these approaches, let's take the opportunity to distinguish between Mu circuits and SRPPs, something not often appreciated. If you set R_a to zero in Figure 3a, eliminate R_{asy} and C_{boot}, employ twin triodes and set R_{k1} equal to R_{k2}, you have an SRPP. The SRPP has a higher Zout and lower PSRR than a Mu circuit and negligible bootstrapping, but its twin triode topology lays claim to a cancellation of even (only) harmonics. John Broskie gives a detailed discussion on the website *Tube CAD Journal*[7] of the differences between these circuits and introduces an interesting variant, the SRPP+. Merlin Blencowe, in a May 2010 *audioXpress* article[8], also delves into the SRRP in good detail.

Comparing Mu Circuit Topologies and Pentode and Triode CFs

Let's consider the topologies first. Other than the screen grid network, the only difference between the Figure 3a and Figure 3b circuits is the biasing of the CF grid. Note that the highest AC voltage level in these circuits is found at the VA anode. The AC voltage at the CF cathode is somewhat smaller, since the CF gain is less than unity. Referenced to the cathode, the voltage at the junction of Figure 3a's R_a and R_{k2} is a bit higher, but it's a bit lower at the junction of R_{asy1} and R_{asy2} shown in Figure 3b. So at low frequencies when the impedances of C_{boot} and R_{g2} are similar, the bootstrapping effect shown in Figure 3b is less than that of Figure 3a. This is not a big effect, but it costs nothing to bias the CF as in Figure 3a rather than 3b. This holds whether the CF is a pentode or a triode. I am unaware of any benefit of the Figure 3b CF biasing scheme.

Moving on to the Figure 1 topology, one possible advantage it has over that of Figure 3 is that its valves needn't share the supply's DC voltage with that across R_a. This means that a larger output voltage swing should be possible, which might be useful in comparison to some other implementations.

Another advantage is that C_{boot} eliminates the need for the separate screen grid capacitor of Figure 3b—but at a steep cost. The C_{boot} in Figure 1 must have a much higher capacitance than that in Figure 2 and Figure 3 by several orders of magnitude. The reason for this is the introduction of the (AC ground) B+ referenced resistor R_{sup} which must be driven through C_{boot}. There is no analogous resistor in Figure 2 and Figure 3.

As to what type of valve to use for a CF, there is no question that the generally higher anode resistance of a pentode than that of a triode leads

Engineering Consulting for the Audio Community

Celebrating 36 YEARS 1983-2019

MENLO SCIENTIFIC

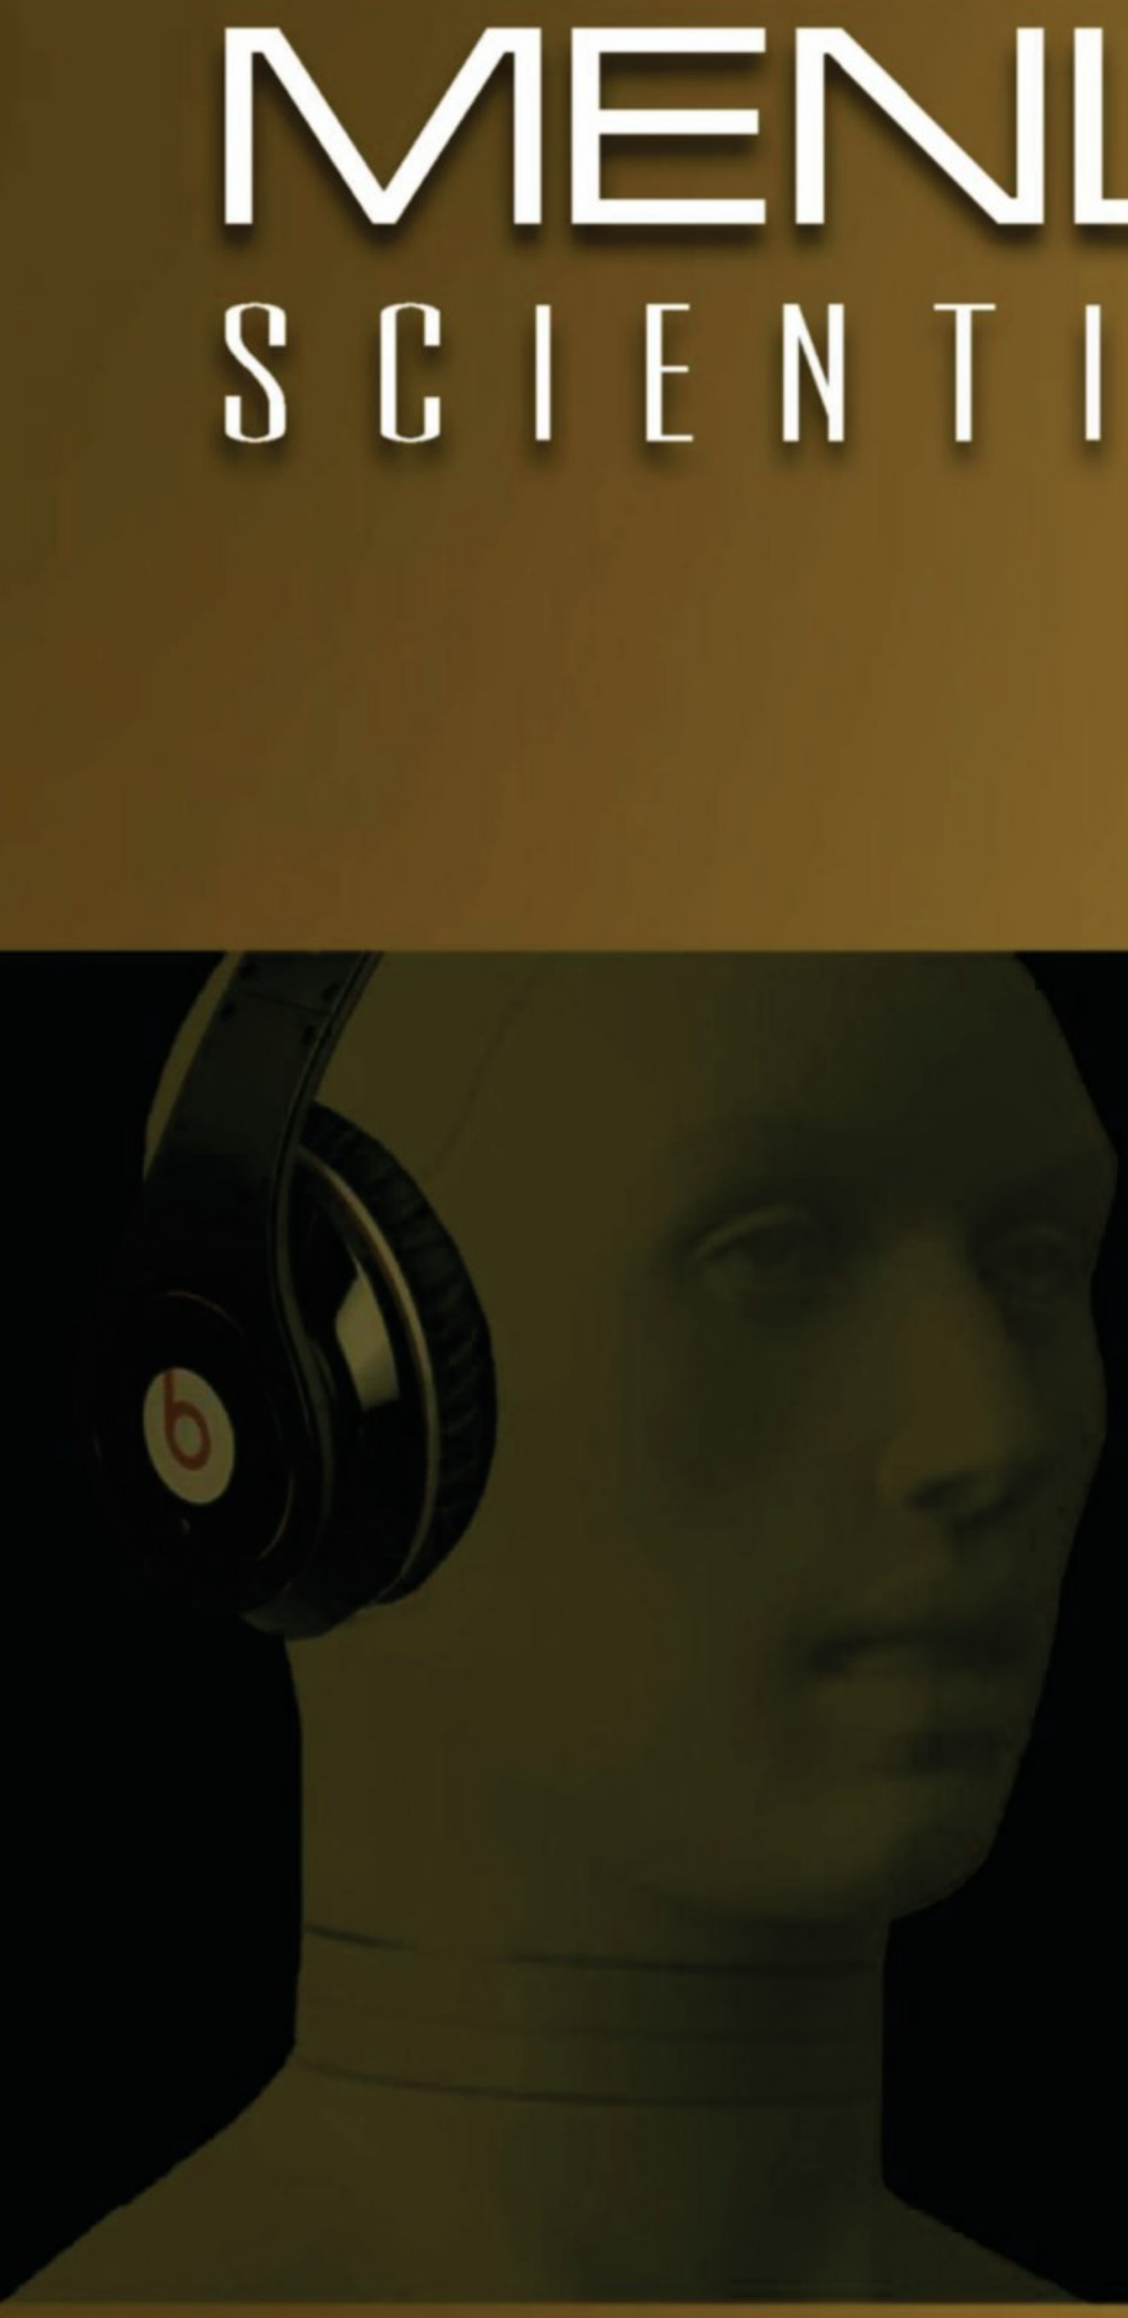

- Liaison Between Materials, Suppliers, & Audio Manufacturers
- Technical Support for Headset, Mic, Speaker & Manufacturers
- Product Development
- Test & Measurement Instrumentation
- Sourcing Guidance

MENLO SCIENTIFIC

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

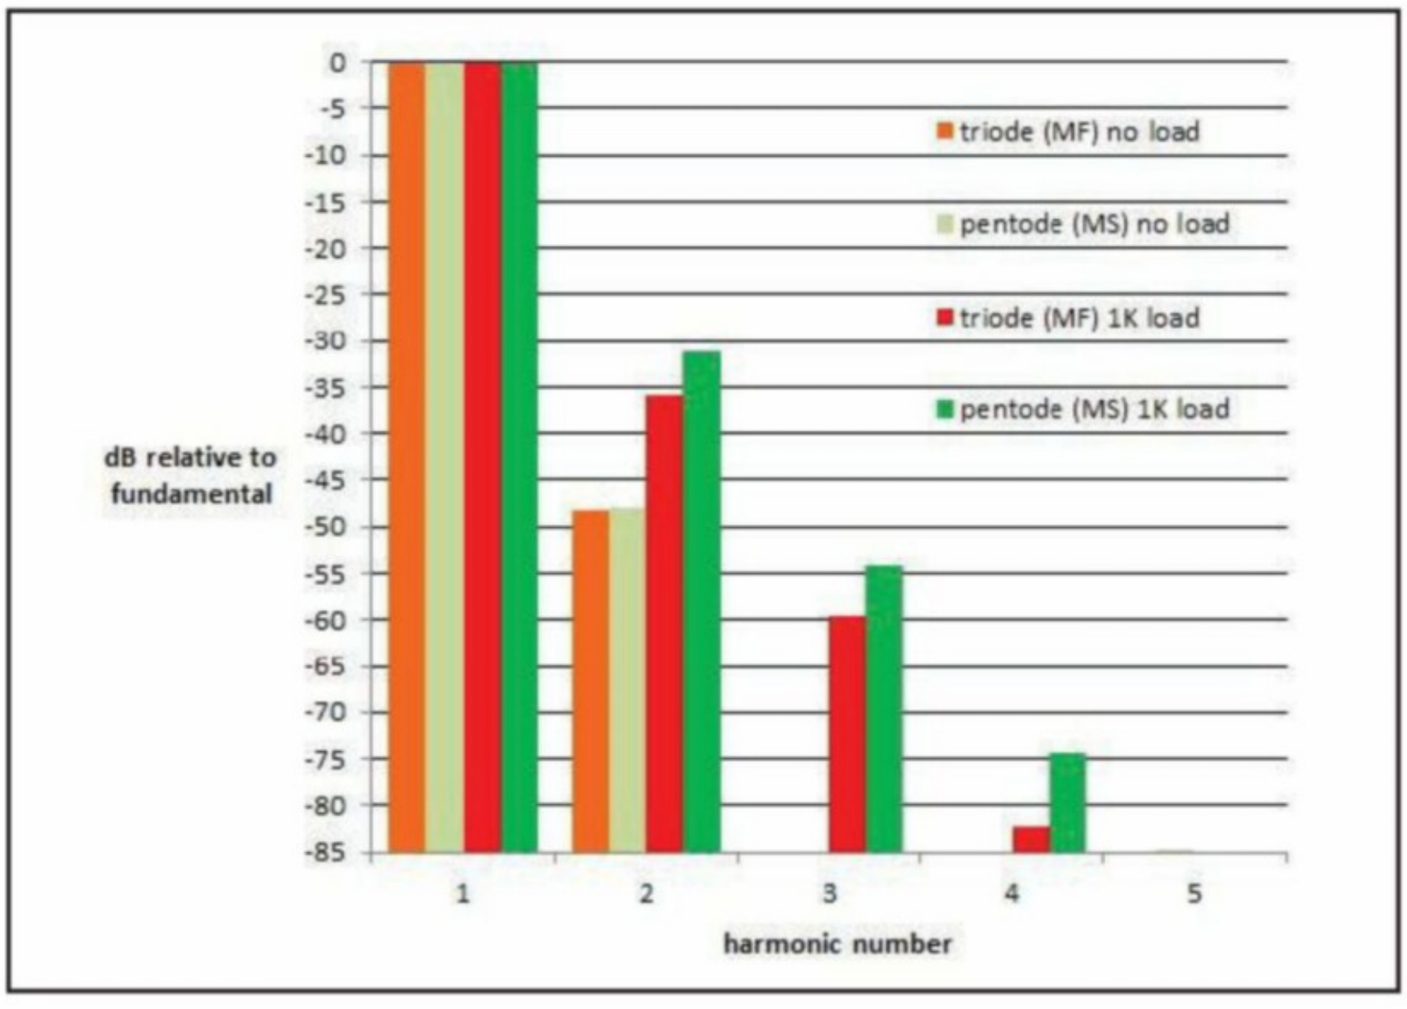

Figure 5: Here we see the distortion performance of the Figure 4 circuit with and without a 1 kΩ load connected to the output. The output voltage is 2.4 kHz at 2 Vrms.

to a higher gain CF. As discussed earlier, this offers more "bootstrapping" and therefore, a higher effective impedance of the triode VA's anode load. This same higher anode resistance reduces the coupling between the cathode and the B+ supply, affording a better PSRR. This higher effective resistance can be traded off in Figure 3b by reducing the value of R_a and the DC drop across it, which can allow higher peak voltage swings. This may or may not be advantageous in certain power amplifier designs. However, we must also beware that pentodes typically exhibit higher distortions than triodes.

Finally, it has been noted that a pentode's screen grid network acts as an added load on its cathode and that this is absent with a triode. In the very few cases in which this is of any significance, the screen grid resistor can always be replaced with a semiconductor-based current source to render the additional load negligible.

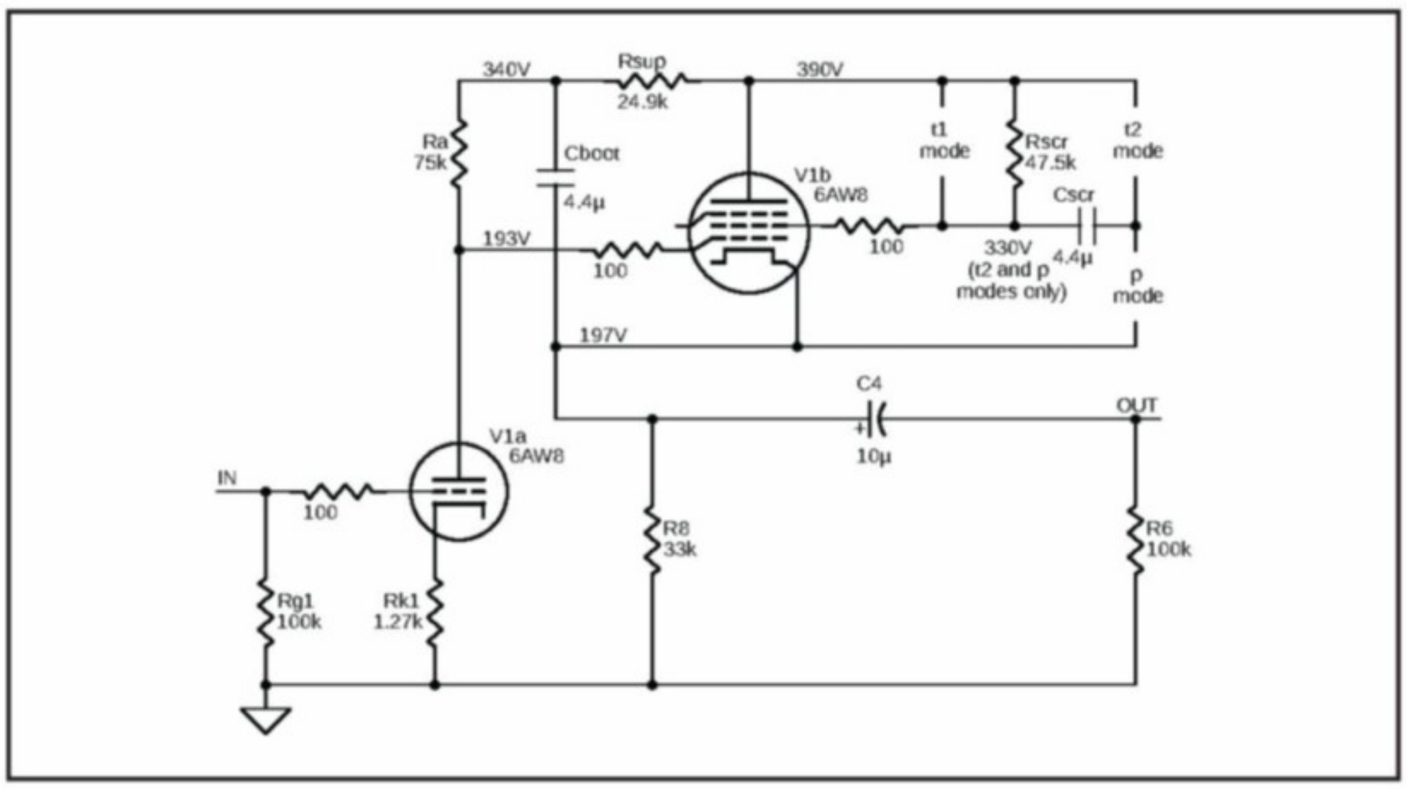

Figure 6: This circuit was tested in pentode mode and in two different triode mode configurations. One triode mode had a screen voltage identical to that of the pentode, and the other set to B+. The output voltage was 30 Vrms at 2.4 kHz.

Comparisons of Some Non-Exemplary Mu Followers and Mu Stages

It would be valuable if we could establish some sort of general comparison of MFs and MSs. What could be the ground rules for such a comparison? Devices of similar physical sizes? Same number of tubes? Same cathode currents? It seems to me that these choices are arguable and somewhat arbitrary.

But then it occurred to me: practically any pentode can be connected to operate in what might be called triode mode. In the case in Figure 3b, just move the screen grid capacitor so that it shunts the screen grid resistor. For Figure 1, the screen grid would have to be connected solely to a separate R-C network powered from the B+ supply as shown in Figure 3b, with the resistor selected to provide the same DC bias as before. Again, the screen grid capacitor could be in either its normal position (pentode mode), or across the screen grid resistor (triode mode). Not only would such tests use the same valve for both alternatives, but the valves would also have the same bias. Let's look at some tests of this type.

Figure 4 shows a Mu circuit made with the triode-pentode 6U8. In its tests, the AC input voltage was kept at 2.4 kHz at a constant level. With C_{scr} connected for triode mode operation and R_L set to 1 MΩ, 2 Vrms was present on the output; with C_{scr} in the pentode mode position, 2.10 Vrms was measured. Next, R_L was reduced to 1 kΩ. The triode mode output level measured 1.36 Vrms while the pentode mode level fell to 1.256 Vrms due to its higher output impedance, as the screen grid current was shorted by C_{scr} and was unavailable to lower the output impedance.

Figure 5 depicts the level of the first three harmonics (2 to 4) referred to the fundamental. The chart was cut off at -85 dBrc, below which my spectrum analyzer gives spurious responses. As can be seen, with a 1 MΩ load, there was no difference between the two modes in harmonic levels. But when the 1 kΩ load was added, we saw a 4 to 6 dB advantage in distortion levels for the triode mode. With the 1 MΩ load, the pentode-mode variant had a PSRR of 38 dB and an output impedance of 628 Ω. The triode-mode alternative weighed in at 24 dB and 434 Ω.

Suppose we wanted to use a Mu circuit as the front end to a power amplifier, where it had to

swing a significant voltage but needn't drive any significant load? I investigated triode-pentode 6AW8 Mu circuits (see **Figure 6**). Here, I had improved the dynamic range of my measurement system and was able to accurately measure signals lower than -85 dBc. The output voltages are 30 Vrms, again at 2.4 kHz. I considered not only the triode (t2) and pentode modes with identical DC bias, but also a triode mode (t1) where the screen grid was connected directly to B+.

Figure 7 shows the distortion performances of the first seven harmonics. You can see that the triode modes' harmonics from the second to the seventh were always lower than those of the pentode mode. The pentode MS had a PSRR of 32 dB and the triode MF's 21 dB. The MF triode t2 mode had the lowest output impedance, 302 Ω, followed by 354 Ω for the triode t1 mode and 392 Ω for the MS pentode. The gains averaged 34 dB and differed at most by a mere 0.6 dB.

So where is the examination of the KT77 in an MF? The first thing I noticed when I decided to investigate this was that the cathode voltage approached that of the screen grid at 30 Vrms out, leading to horrible distortion. So I tied the screen grid to B+. The result was an unimpressive 0.10%. Why, I wondered? Then I measured the KT77 triode mode CF gain: The effective Mu was a mere 8.4! This meant that the multiplier of R_a was smaller than even that associated with the 6SN7. Was the 6SL7 now the major contributor of distortion? Or was this a case of the known phenomenon that pentode distortion benefited from a moderate load (R_{asy}), whereas triodes prefer as light a load as possible?

I decided to eliminate the 6SL7 and test the distortion of the KT77 CF's alone. To do this, I had to design and test a low distortion voltage amplifier, which can be seen in **Figure 8**. Its THD at an unloaded 30 Vrms output is 0.002%, with second, third, and fourth harmonics at -94 dBc, -102 dBc, and -116 dBc, respectively. Higher harmonics were better than -130 dBc.

Figure 9 shows three different KT77 CF configurations that were tested with the help of this circuit. Case p shows the loads that the pentode-mode valve sees in the Figure 1 MS circuit. Case t_a retains these loads and converts the pentode to triode mode. Case t_b replicates the light loads that a triode would see in the original MF topology of Figure 3a, using the Q1 circuit to implement a high impedance current source. If you object to "sand," you could replace it with a "hollow-state" current source, but I do believe that Q1 is innocuous in this application.

As suspected, the highest distortion was exhibited by the triode with the heavy loads (case t_a) and the

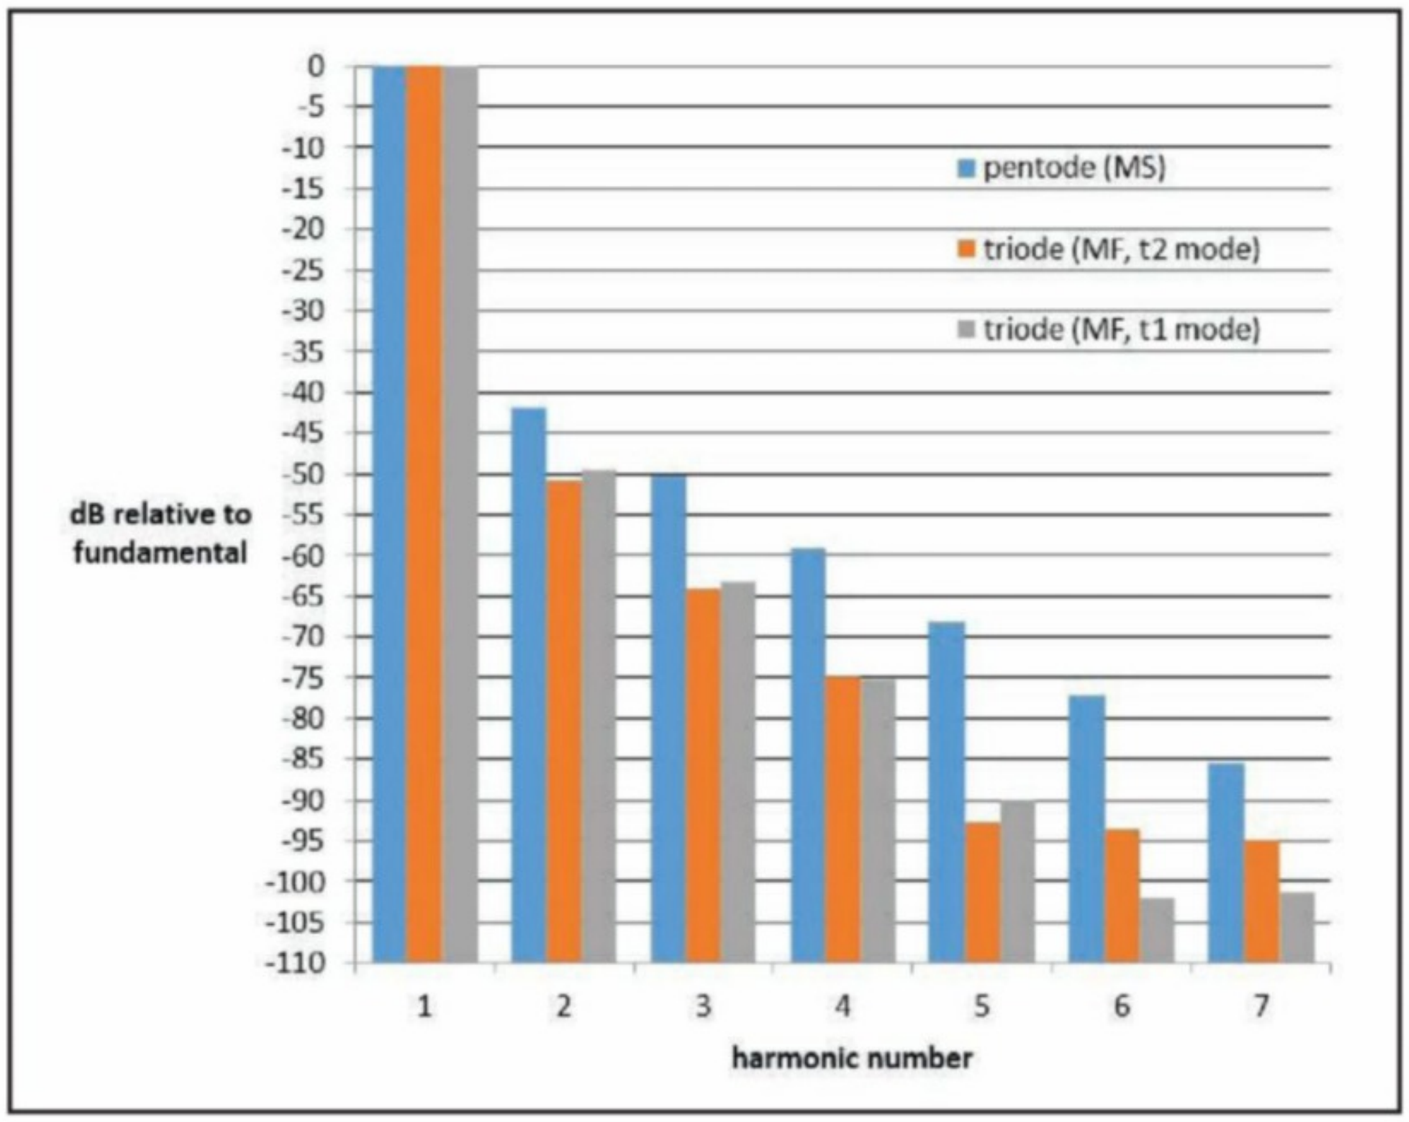

Figure 7: The distortion performance of the Figure 6 circuit is presented. Pentode mode operation leads to noticeably higher distortion than either of the triode modes.

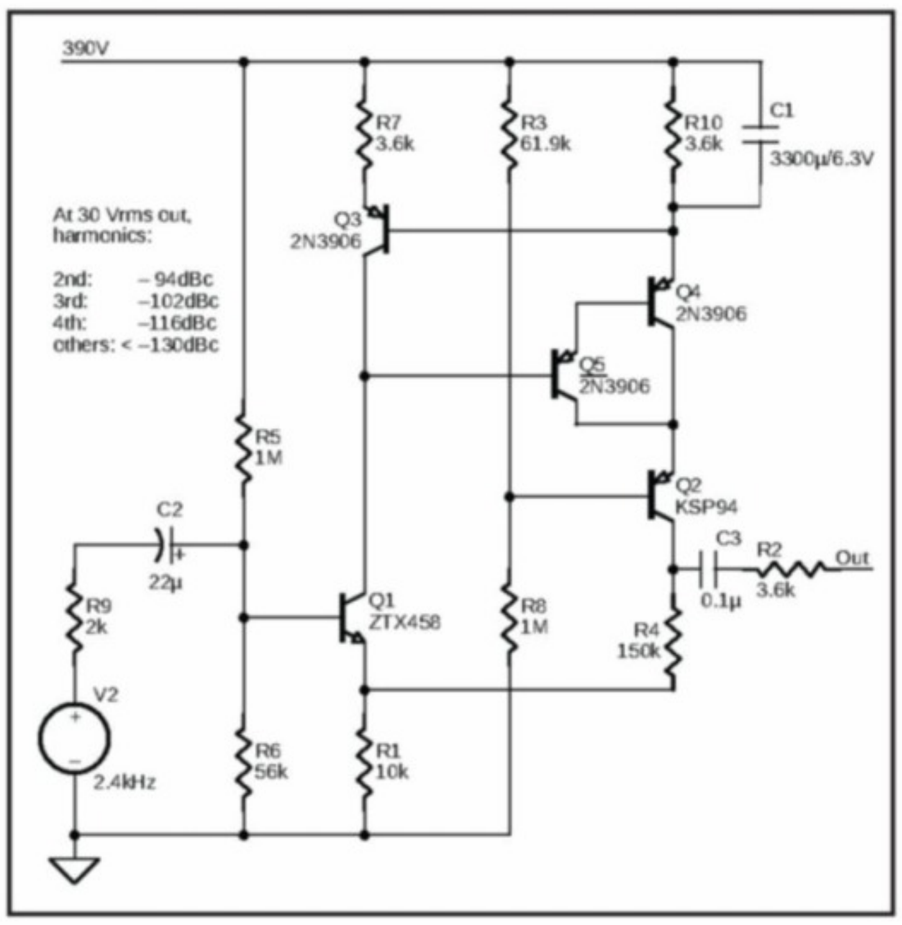

Figure 8: A low distortion amplifier is required to convert a sub-3 Vrms signal to 30 Vrms and more to drive the grid of several different CFs.

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.

FREE SUBSCRIPTION

Voice Coil is *the* Publication for Loudspeaker Industry Professionals!

Sign up for your copy at: **VOICECOILMAGAZINE.COM**

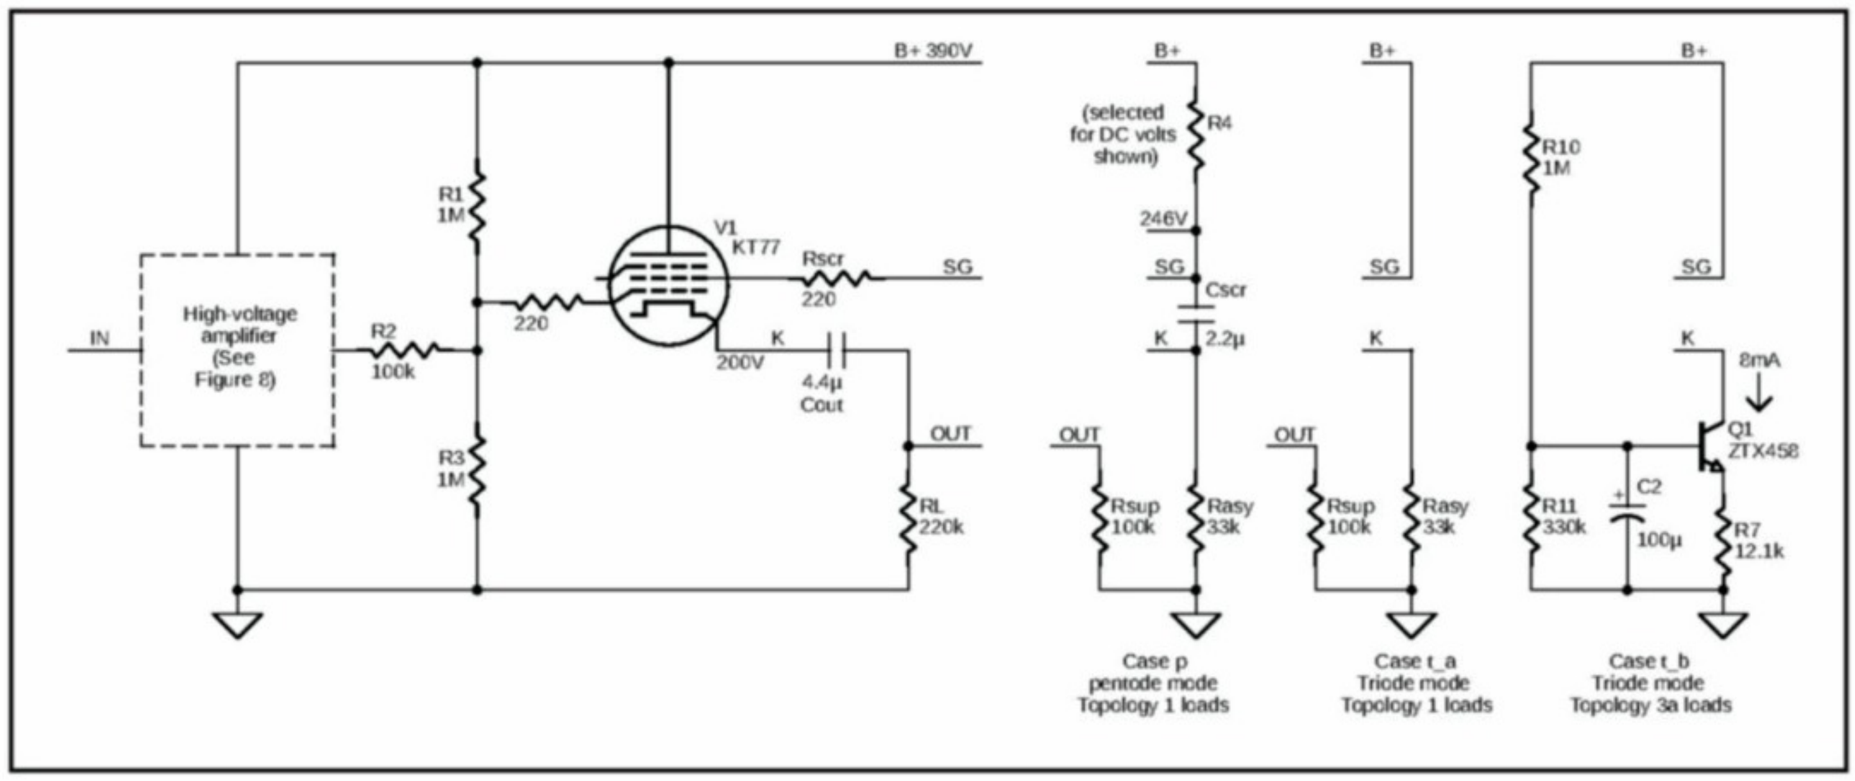

Figure 9: The circuit of Figure 8 provided the excitation to produce 30 Vrms out at three differently configured KT77-based CFs to analyze their distortion performance.

lowest by the lightly loaded triode (case t_b). The pentode (case p) was somewhere in the middle. The results appear in **Figure 10**.

Observations and Generalizations

The exemplary circuits at the beginning of the article demonstrate that even the relatively low CF gain of the low/moderate Mu Tung Sol 6SN7 now in production is sufficient to bootstrap an anode load resistor in an asymmetric Mu Follower and yield excellent distortion performance. The distortion is virtually identical to that of a Mu Stage using a KT77 in a more costly, higher power circuit that requires an additional tube for stereo operation. There are a few tubes that contain twin pentodes (e.g., the 6AR11), which can be found in a list of GE Compactron tubes[9]. Its suitability as a CF in an MS is yet to be examined, but on the bright side, as new old stock (NOS), it appears to be very inexpensive and readily available.

Mu Stages generally offer a higher PSRR than do Mu Followers, a benefit of their higher anode impedances which also lead to their higher gains as CFs. A given pentode can be operated in pentode or triode mode. Triode mode operation will generally result in a lower output impedance (because it benefits from the screen grid current which is not shorted by C_{scr}) and lower distortion if lightly loaded.

If you have selected a pentode for a Mu circuit, you would do well to investigate whether it performs better as part of an MF or an MS. Generally, the light loading of the Figure 2 and Figure 3a topology will favor triodes, whereas the heavier loads of the Figure 1 and Figure 3b topologies afford an opportunity to operate in the distortion null that many pentodes exhibit. Of course, you can always reduce R_L found in the Figure 3b MS topology and benefit from the reduced cost and size of the large C_{boot} capacitor that is required by the Figure 1 circuit.

Please don't misconstrue any of these generalizations as hard and fast rules. Many fine MFs and MSs are awaiting discovery. Happy hunting! ax

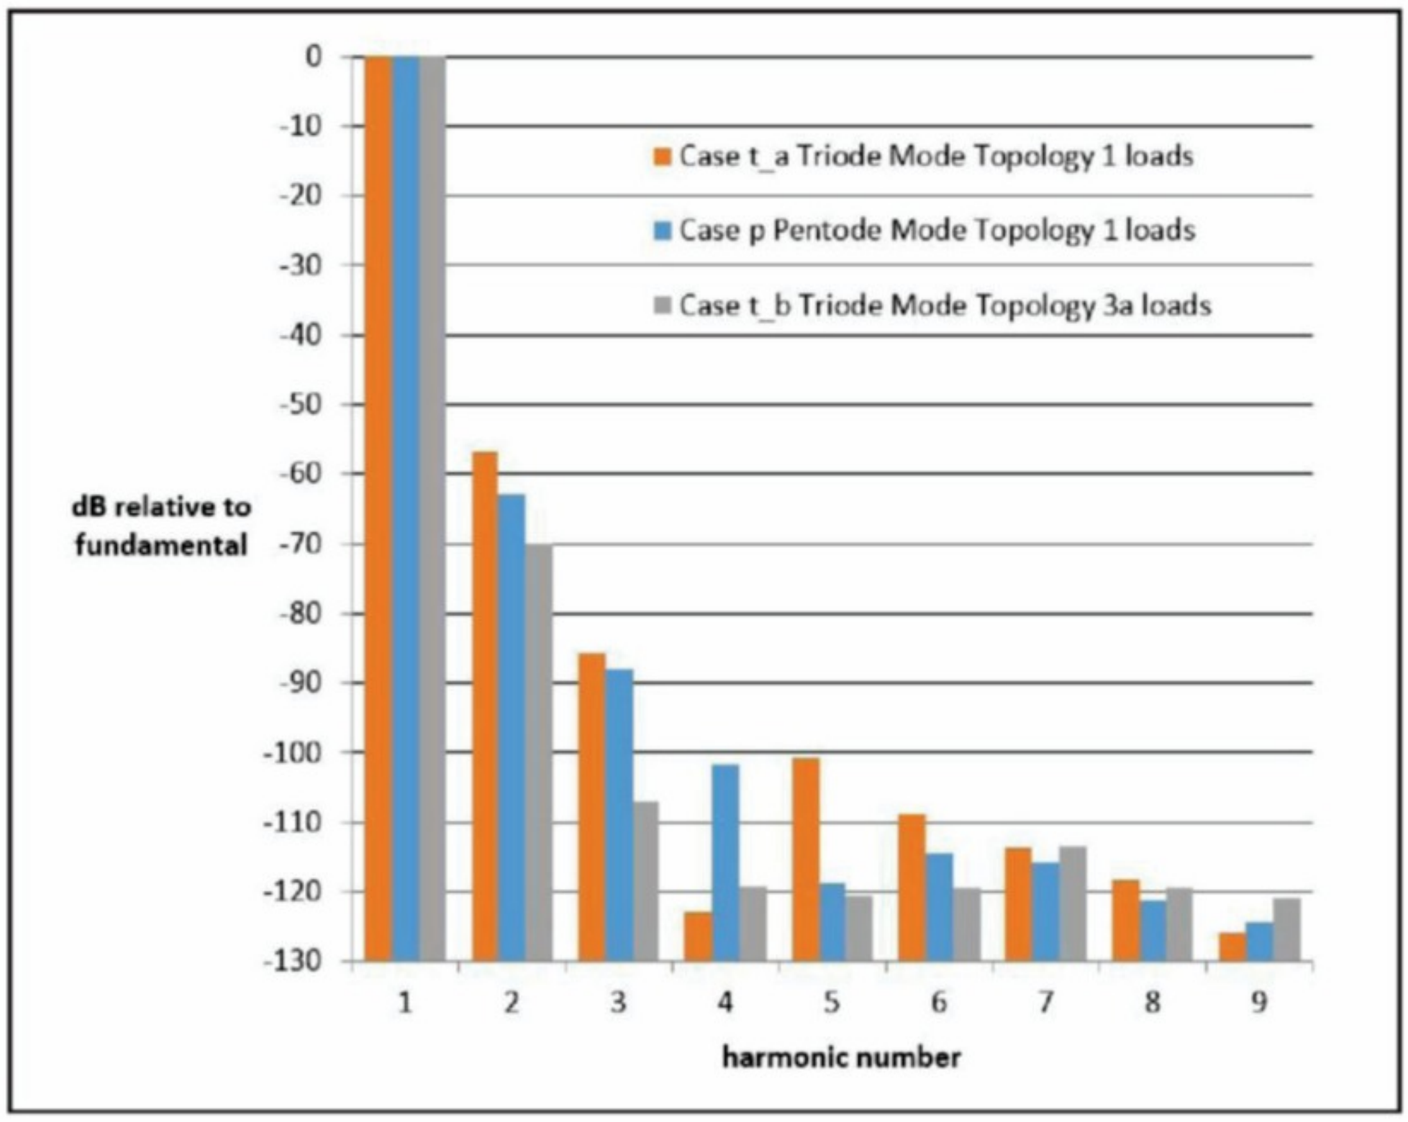

Figure 10: Three different KT77 CF configurations were tested. The highest distortion was exhibited by the triode with the heavy loads (case t_a) and the lowest by the lightly loaded triode (case t_b). The pentode (case p) was somewhere in the middle.

References

[1] C. Paul, "Simple Sineman: A Sine Wave Generator for the Rest of Us," *audioXpress*, November 2013.

[2] A. Kimmel, "An Outrageous Mu Stage: An Update of My 1993 Mu Stage Design," *audioXpress*, May 2019.

[3] A. Kimmell, "The Mu Stage," *Glass Audio*, Volume 5, No. 2, 1993.

[4] C. Paul, "Letters to the Editor," *The Audio Amateur*, Issue 2, 1985.

[5] J.W. Horton, *The Journal of the Franklin Institute*, Volume 216, December 1933.

[6] C. Paul, "Letters to the Editor," *The Audio Amateur*, Issue 3, 1991.

[7] J. Broskie, "John Broskie's Guide to Tube Circuit Analysis & Design," *Tube CAD Journal*, www.tubecad.com/2009/10/blog0173.htm.

[8] M. Blencowe, "The Optimized SRPP Amp: Part 1," *audioXpress*, May 2010.

[9] "GE Compactrons," Jeff Duntemann's Junkbox, www.junkbox.com/electronics/GE%20Compactron%20List%20Spring%201964.pdf

Introducing Aurora Open-Source DSP

An Interview with Raphael Knoop of Auverdion

Raphael Knoop is the founder of Auverdion, a fledgling company that offers an open-source DSP that he has dubbed Aurora.

By
Shannon Becker
(United States)

Aurora is an idea that evolved when Raphael Knoop sat down with some friends who all love great sound and started talking about their projects. Following two years of development, Knoop and his associates developed Aurora—an advanced eight-channel DSP with awesome specs, unmatched versatility, and an easy-to-control interface. Knoop's fledgling company is called Auverdion (www.auverdion.de).

SHANNON BECKER: Please share some information about your background.

RAPHAEL KNOOP: I was born and grew up in Bielefeld, Germany. After school I studied Electrical Engineering and Communications Engineering at the University of Bremen. Today I am working as a Senior Research & Development Engineer for a big manufacturer of digital music instruments. I am developing DSP algorithm for audio processing and sound synthesis.

BECKER: How did you become interested in audio?

KNOOP: It all started with the organ music by Johann Sebastian Bach when I was a teenager. My father had a CD with a recording from the Cathedral of the Immaculate Conception in Linz and I loved this CD. Fascinated by organ music, I bought CD after CD and my collection was growing. Soon I started thinking about improving my audio system, but I had just finished school and was preparing for university. So there was no money for buying better loudspeakers or electronic devices. But my parents had a garage that I could use as a workshop for building my own loudspeakers and later audio electronics. The list became very long over the years, there have been analog crossovers, preamps, DACs, and with every project as I continued to learn a bit more.

BECKER: Tell us a bit about your company Auverdion, based in Verden, Germany.

KNOOP: Auverdion was founded to make Aurora available for audio enthusiasts. It was a logical consequence when I decided to put Aurora on

the market. Thus, Auverdion is one of these little start-up companies that turns a nice idea into a real product.

BECKER: How did Aurora come about? What compelled you to take on this project?

KNOOP: It is now almost three years ago, when I was looking around for an amplifier module for my new subwoofer. But I could not find anything that thrilled me. Either the amplifier did not have enough power or there was no DSP onboard or the beast was too big, or. Just prior to that search, I had designed a DSP (freeDSP CLASSIC BAL A). That made me think: Let's put it together with a Texas Instruments (TI) TPA3255 Class-D amplifier chip. I did it and showed the module to the community of audio enthusiasts at www.diy-hifi-forum.de. Soon a discussion started, community members came up with questions and ideas, and I learned what the most wanted features are and what the drawbacks are of all commercially available DSPs on the market. With that knowledge, I redesigned the entire DSP module. It is a good feeling to know you are designing something that many other enthusiasts would love and you would make them happy if you share it with them.

BECKER: Can you tell us more about Aurora? How does it work?

KNOOP: Basically Aurora is a member of the freeDSP family (www.freedsp.cc) built around the SigmaDSP ADAU1452 from Analog Devices. The big advantage of the SigmaDSPs is their graphical programming environment SigmaStudio making it possible to build DSP programs even without skills in C/C++ or Assembly programming. Aurora connects to the outside world by an eight-channel 32 bit ADC (AKM5558) and an eight-channel, 32-bit DAC (AKM4458). All analog inputs and output are balanced. Of course S/P-DIF in and out is there, too.

Nowadays, audio streaming from a computer or a media server plays a major role in hi-fi systems. Therefore, a XMOS chip was added to provide an USB Audio Class 2 (UAC2) interface to the DSP. Now you need a way to make settings to your DSP. Therefore, an ESP32 was added to provide WiFi connectivity. And finally there are a lot of connections for things such as the display, the rotary encoder, a volume potentiometer, an IR sensor, the fan, a temperature sensor, and an expansion header for more analog inputs or outputs. In short, it can grow in any direction.

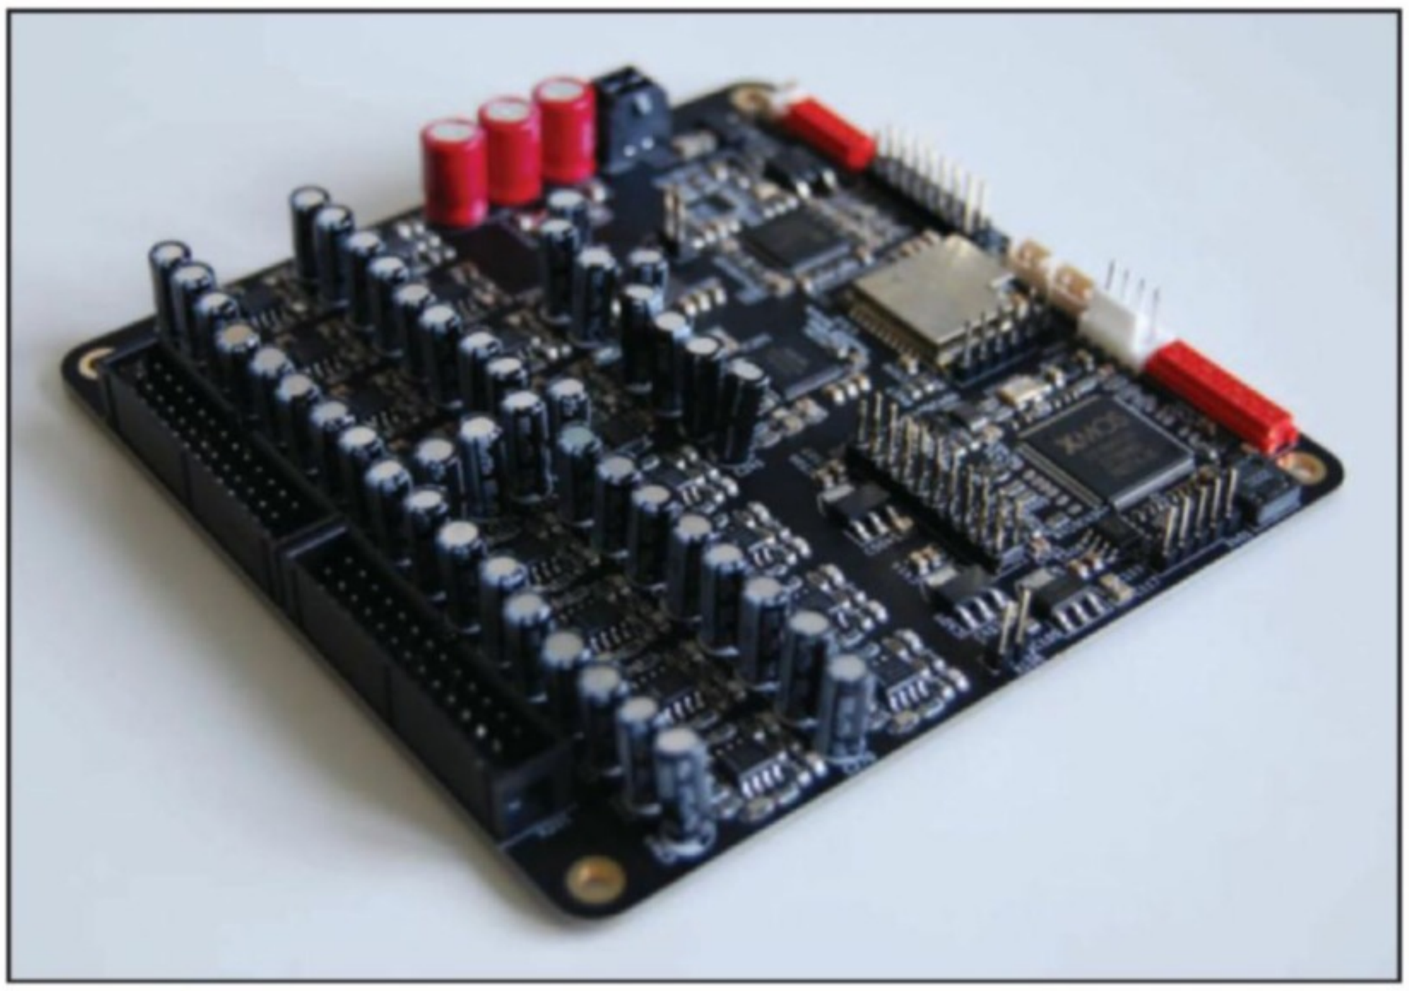

Aurora is an open-source DSP with eight inputs and eight outps, USB Audio Class 2 and wireless control via WiFi and Bluetooth.

Yet the best hardware is useless if the control software forces you to study the manual and the principles of digital signal processing for four years. For Aurora, I wrote the app—auverdionControl (available for macOS, Windows 7 and 10, iOS)—that makes the everyday use of Aurora very pleasant. It communicates with Aurora via WiFi. You can upload DSP-plugins to your Aurora board to change the audio processing. For example, there are plug-ins for running eight audio channels with signal processing in parallel, a 7.1 loudspeaker management for your home cinema, or a four-channel FIR filter plug-in to build a four-way loudspeaker with FIR filtering. But it does not limit the creativity. You can still use SigmaStudio to build in your own ideas.

BECKER: How, if at all, has the design changed from its original concept?

KNOOP: It has radically changed. The original concept was a SigmaDSP ADAU1701 together with the TPA3255 and only one analog stereo input and control via USB. It was designed for a subwoofer. And then it grew almost every month. First, the bigger DSP ADAU1452 came in, analog outputs increased to four, instead of one amplifier channel up to four have been considered and finally the amplification was completely separated to improve the flexibility of Aurora. Another change was the MCU for UAC2. I had used a MCU from Microchip and already built a prototype when end of life for

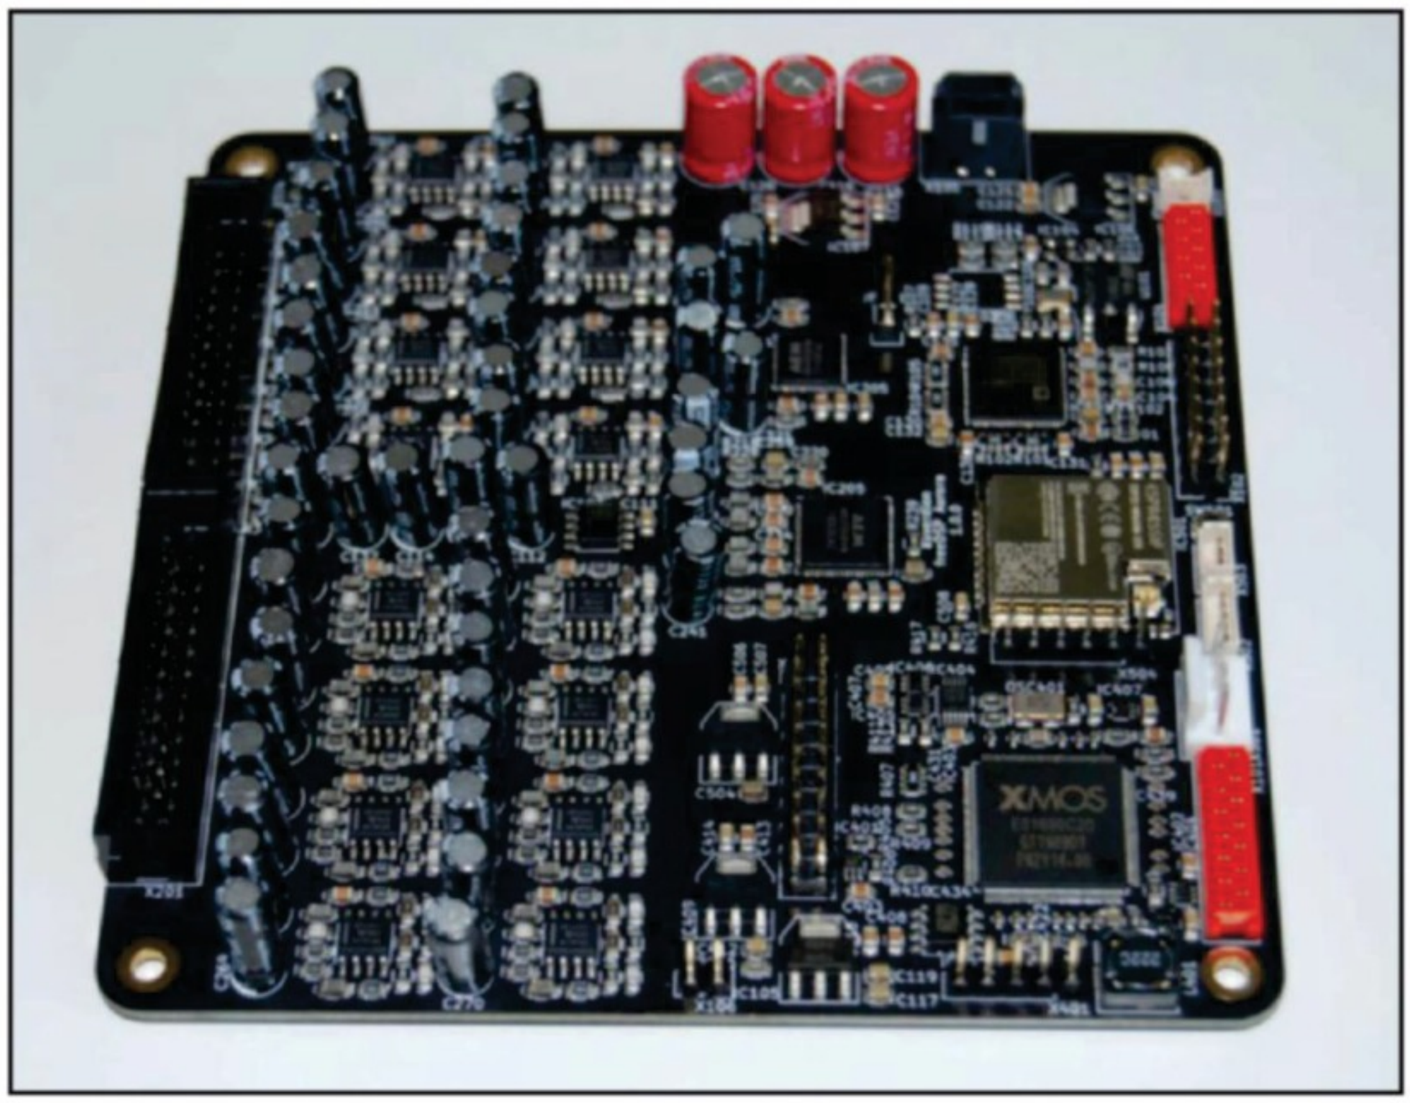

Here is a top view of the Aurora, which has an Analog Devices ADAU1452, 294.912 MHz, 32-bit SigmaDSP; XMOS XE216-512-TQ128 for multichannel bidirectional audio streaming; an ESP32 for WiFi or Bluetooth control; an AKM AK4458 32-bit DAC; and so much more.

this chip was announced. Therefore, I switched over to XMOS. The concept of Aurora how it is now has nothing to do with the first concept from three years ago anymore.

BECKER: Why did you decide to make it an open-source DSP?

KNOOP: I believe in open source. It boosts creativity. If you want to make a product for people and you want them to be creative with your product then you have to give them all freedom to change and hack whatever they want. Borders like closed source hard- and software are a hurdle when trying your ideas and only make you frustrated and looking for another product to buy instead of being creative.

BECKER: What is the idea of the AddOns?

KNOOP: During the design phase of the prototypes I realized that everyone has different requirements for the audio connectors. Some want to have XLRs for the outputs but RCAs for the inputs, another group wants jacks for inputs and outputs but only four of each. The next group is only interested in SPDIF via Toslink and so on.

Therefore, Aurora does not come with any dedicated audio connector like XLR, RCA, or whatever. It provides all the connection options on pin headers and the user select from the AddOns to provide exactly the needed connectors. It even enables users to design a new AddOn if the current options do not cover the requirement.

BECKER: Have you used Aurora for any of your projects?

KNOOP: Currently, I am using Aurora to drive my three-way active loudspeakers for stereo listening. Aurora is driven by a stereo input and does all the crossover and eq'ing. At its outputs, it drives six Hypex UcD amplifiers. I also am planning to use another unit in my home cinema for loudspeaker management.

BECKER: What would you recommend Aurora be best used for?

KNOOP: Actually, you can use it for everything that comes to mind. Thanks to the open-source license there are no limits. Most of my customers used it for either driving multi-way active speakers or for loudspeaker management in their home cinema including room eq'ing. There I see the main application for Aurora.

BECKER: Is Aurora the first product you have promoted through Kickstarter?

KNOOP: Yes, it was my first product. Actually, it was even the first time I promoted a product through Kickstarter. Thus, I had to learn very quickly how to do that on Kickstarter. At the beginning of the campaign, I talked to somebody who works for a marketing agency and he said Kickstarter is so different. Instead of just promoting a product, you

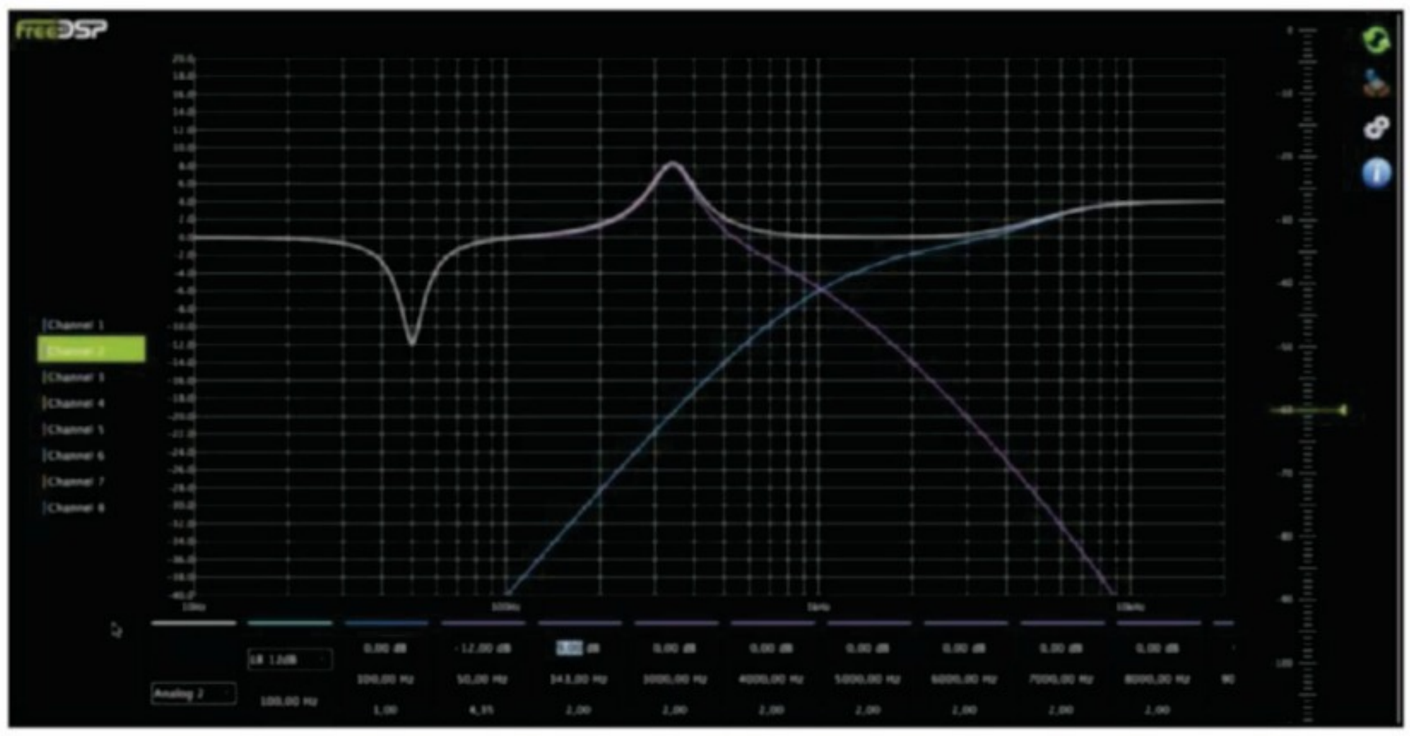

This is a screenshot of the auverdionControl.

can let others become part of your story with their support.

BECKER: Why did you decide to promote it via Kickstarter?

KNOOP: Kickstarter makes it very easy to start a crowdfunding project: Create an account and describe your project. The all-or-nothing principle is fair for both sides, project founder and backers. At the end of the campaign, you either have enough money to realize your project or you don't have enough backers and nobody has to pay any money. If the campaign is successful, it is a win-win for both sides.

BECKER: How did the Kickstarter process work for you?

KNOOP: While creating a project page on Kickstarter was quite simple the real challenge was to keep the campaign running. You need to promote your project wherever possible. I had support from the DIY community all over the world. They spread the news about the Kickstarter campaign everywhere. I even saw a notice about the campaign in a German computer magazine not related to audio but related to the maker community.

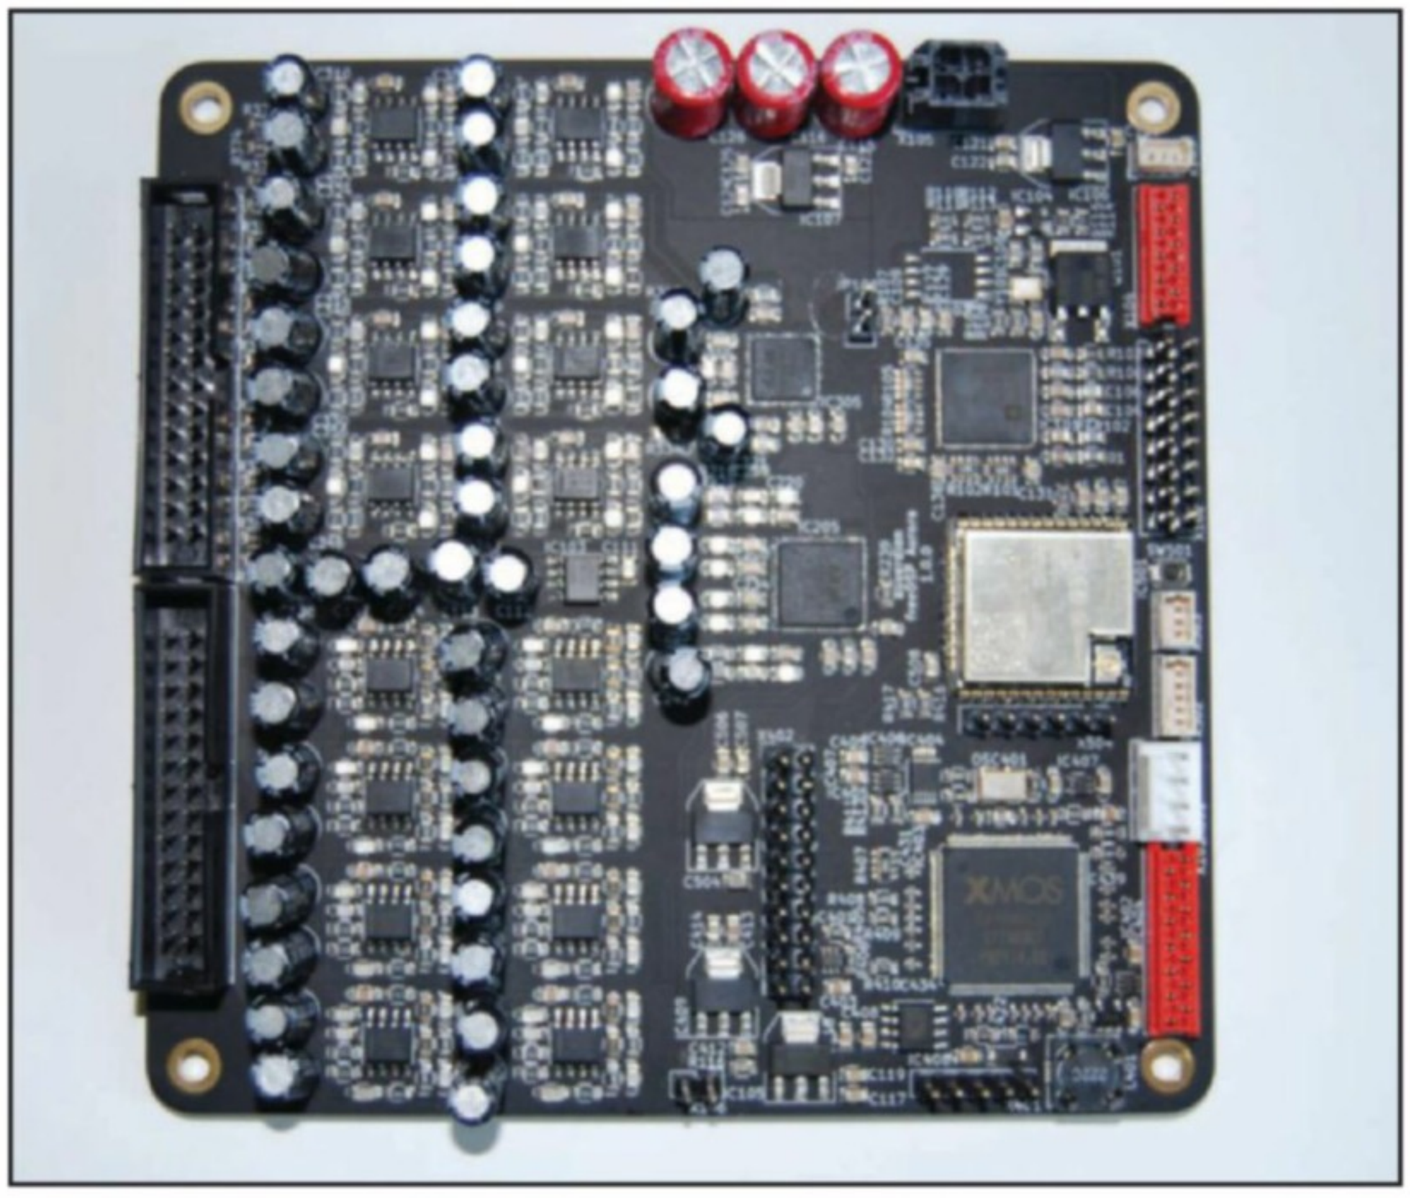

Aurora's 110 mm × 110 mm board also provides suppor for a display, rotary encoder, volume potentiometer, temperature sensor, PWM controlled fan, and an IR sensor.

BECKER: What were some of the challenges?

KNOOP: The biggest challenges were the German and European laws. For a newcomer, it looks like a big thicket of laws and rules all related to putting electronics on the market and sometimes it looks inconsistent. But there are agencies who can help you with that. Yes, they do cost money and that makes small production runs like Aurora more expensive but in the end it spares your nerves and gives you confidence that your business remains legal. Another challenge have been the German authorities. Very often you have to explain what crowdfunding is and what you are going to do.

BECKER: How long did it take from concept to shipping to customers?

KNOOP: From the final concept to shipping it took eight months. In total, it has been about two years with different concepts, designs, and prototypes.

BECKER: Are you working on any new audio innovations?

KNOOP: Yes I am. Mainly new software features for Aurora and auverdionControl will come in the future. One idea is to couple two or more Auroras together to get more power for your loudspeaker management in your home cinema. Integrating Aurora smoothly into your smart home system is another idea. The communication protocol between Aurora and auverdionControl is based on HTTP requests. That makes it easy to connect Aurora to a smart home software (e.g., openHAB).

BECKER: Do you have any advice for *audioXpress* readers who are interested in launching their own audio products?

KNOOP: Make sure that you have a community that backs you. It is not enough to have a good product. Especially audio products, which are a niche market on Kickstarter, will fail to get enough attention and pledges if there is no community. And you should listen to them. The community will tell you what features are needed to make it successful. Distillate the best ideas from it.

And as soon as your campaign is running stay in contact with your backers. They are a part of your story, they make your product possible. Take care of them!

A High-Voltage Delay for Tube Amplifiers

The Sequel

In this article, Jan Didden offers a modern update on his original high-voltage delay design, which he built to complement tube amplifiers.

By
Jan Didden
(Technical Editor, *audioXpress*)

In 2014, *audioXpress* published my design for a delayed on-switch for a tube amplifier's high voltage. Tube heaters need some time to heat up, before a tube starts conducting. Depending on the tube type, this can range from a dozen seconds to a minute. In itself, this is no problem—you just wait a little bit before your music starts playing.

But normally the high-voltage supply for that tube comes on at the same time as the heater voltage comes on—often supplied from the same transformer. That means that the tube sits there for up to a minute with high voltage applied and no conduction. This may shorten tube life—it depends on the tube type and the actual high voltage. The jury is still out on how significant it is, but the effect is known.

Another issue is that the high voltage initially, when the tubes do not yet draw current, can be quite a bit higher than the voltage under load, putting additional stress on the tubes and possibly the high-voltage capacitors. So we should not take chances with limiting the life of our precious tubes, and switch on the high voltage only after the cathodes have been properly heated up.

Why the Sequel?

The original unit worked fine, but it had two disadvantages: It needed a separate 6.3 V heater winding for power, and was limited to a selection of just two delay periods. The heater winding had to be separate from the other heater windings used in the amp because they float on the high voltage. That often meant getting an additional transformer for it, which meant extra cost and extra space required in the chassis. The dual delay selection sometimes meant that one was possibly too short and the other too long.

So I set about to redo this unit by making it self-sufficient—there's no additional outside power required, and you can set the delay between 20 seconds and 254 seconds, which should cover all requirements.

The Circuit

The circuit is solid-state switch-based for noiseless and long-life operation and low power dissipation. It is also easy to integrate even in existing tube amps, and switches on the high voltage from the zero mains crossing to avoid loud thumps

in the speaker. (There can still be an "on" thump but that is then from something else in the amp, like asymmetrical circuitry.) And, of course, no impact on amplifier audio performance!

The topology is shown in **Figure 1**. By inserting the switch in the return line to the power transformer I leave everything after the last reservoir cap unchanged. So there is no impact on the power supply quality and no messing up any carefully laid out ground circuitry. It can be used either with a bridge-type rectifier or a center-tapped secondary with a double-phase rectifier as shown.

The actual switch is connected between Gnd and CT. The unit is powered by a very small current stolen from the high voltage. The complete circuit is shown in **Figure 2**.

The control circuit is simple in its working but getting it to work through a wide range of voltages and under its own power makes it a bit non-intuitive. Let's start with the switch Q3 on the far left, an N-channel high-voltage MOSFET. The drain is connected to the transformer return line as shown in Figure 1. The source is connected to ground, and if this field-effect transistor (FET) switch is closed, the loop from secondary and rectifiers to the cap is closed. The cap will then charge and high voltage is provided to the amp.

On the center-right is the microcontroller U1. It's a small, 8-pin DIL chip that I programmed to turn the FET on after the set delay. Through its pin 7, the controller keeps the FET gate shorted through Q2, keeping the switch open during the delay. (Capacitor C2 keeps the switch open during the short time the controller needs to get its act together). Through R8 and R9, a sample of the secondary AC is provided to U1, so it can see when the AC goes through a zero crossing. After the delay, and at the zero crossing

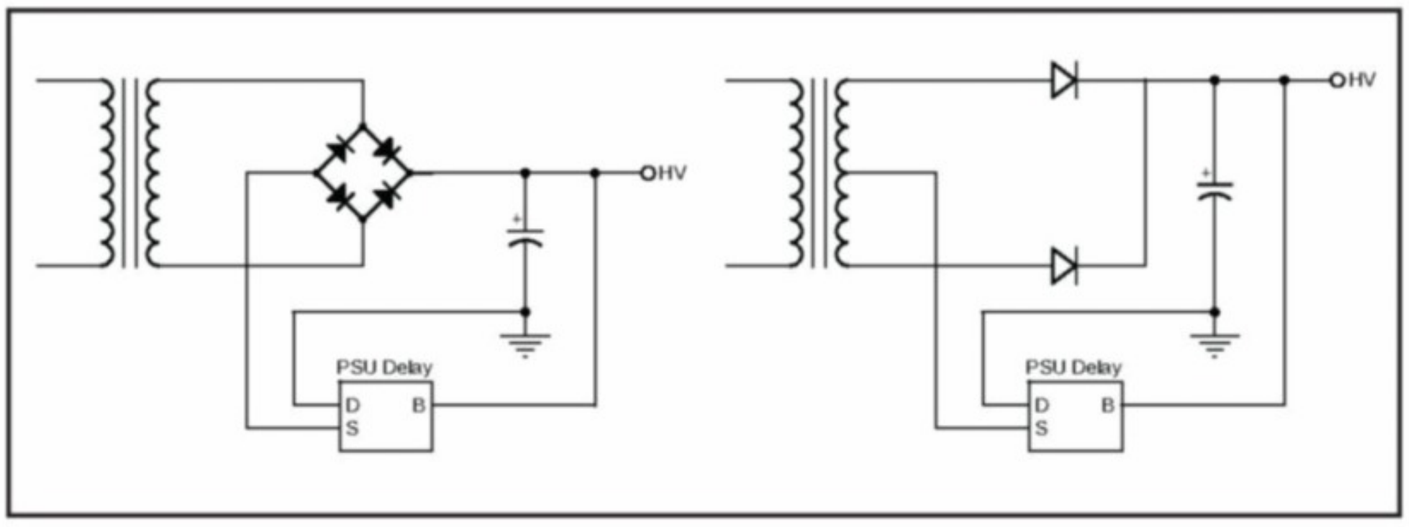

Figure 1: Connecting the high-voltage delay to an amplifier

point, the controller releases Q2 by pulling down its GP0 pin 7, R2 drives the gate voltage up, the FET switches on and the power supply capacitor charges up. Music flows from the speakers!

The controller also drives a LED to show delay or active state: It blinks in a 2 second rhythm during the delay period, then comes on steady (or goes out, see later text) after the high voltage is on.

Q1, driven by Q4, provides a small current from the secondary voltage to power the controller and the LED. I specified a very low current LED, total supply current is just 1.8 mA maximum, and the dissipation in Q1 is generally less than 1 W. Nevertheless, with very high voltages Q1 can become hot to the touch. If you want to play it safe, mount Q1 below the PCB to your amplifier chassis (see Construction).

The Controller

You may wonder what business a digital (there—I said it!) controller has in audio. Well, none of course, that's why it isn't in the audio. And if you are afraid of "digital hash" or whatever, be assured: After the initial delay and switch-on of the high voltage, the

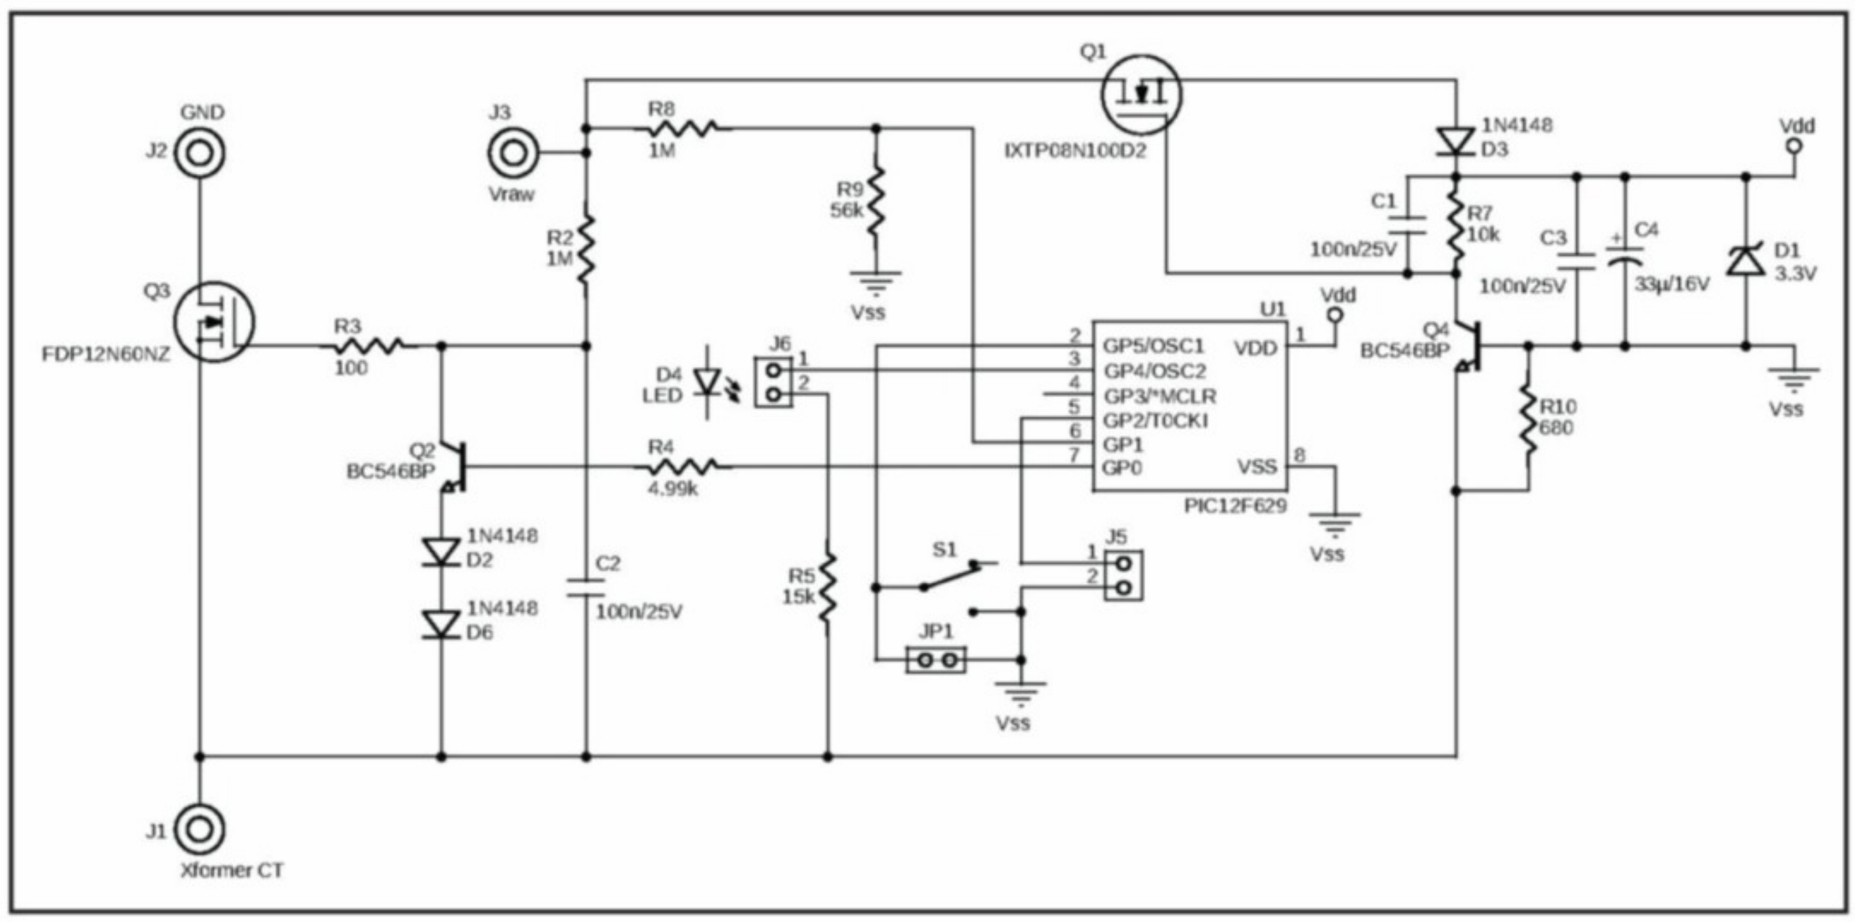

Figure 2: The schematic shows the complete circuit for the high-voltage delay unit. J3/Vraw is the top of the first rectifier capacitor.

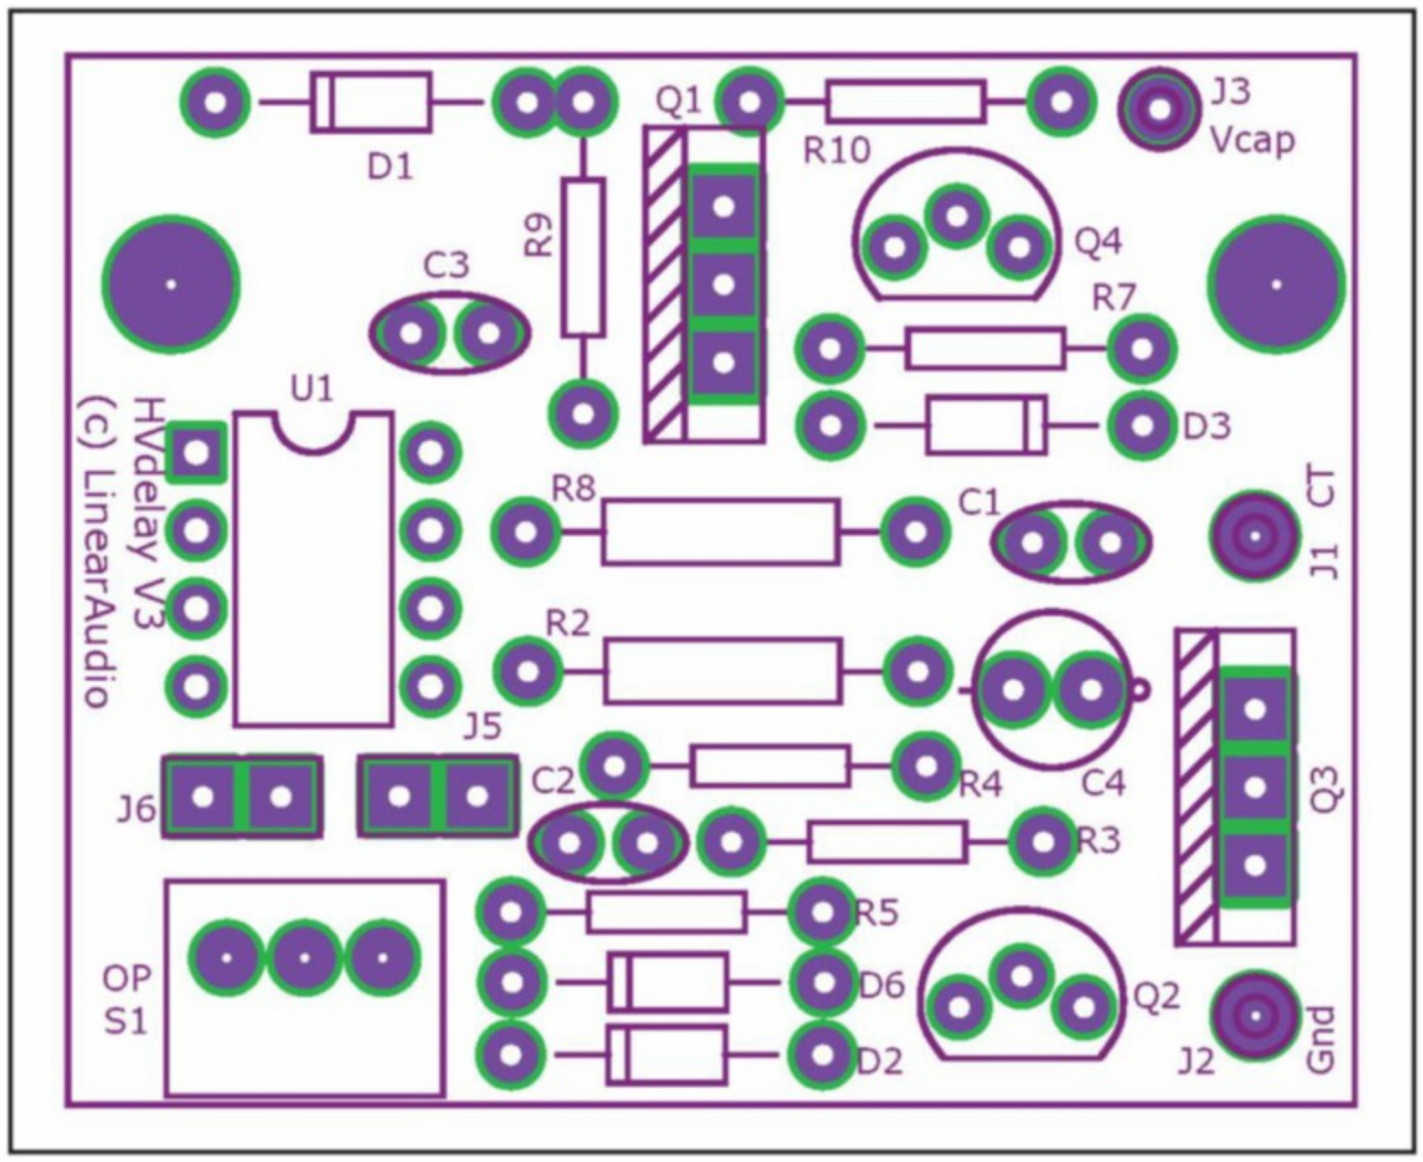

Figure 3: PCB parts location and stuffing guide

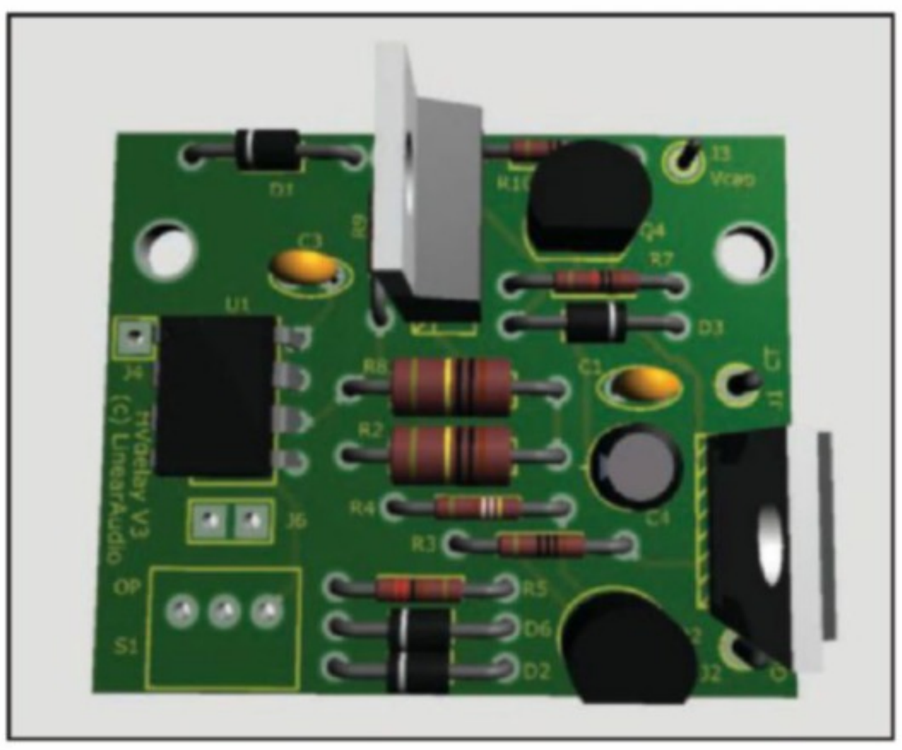

Photo 1: Top of the board in 3D view

About the Author

Jan Didden has written for *audioXpress* since the 1970s. He is retired following a career with the Netherlands Air Force and NATO. He worked in logistics, air defense, and information technology. Retirement has provided him with the time to finish all the audio projects that have piled up for decades. He writes about them on his website linearaudio.nl. Jan is also the publisher and managing editor of the twice-yearly bookzine *Linear Audio.*

Resources

J. Didden, "Extending Tube Life with a High-Voltage Delay," *Linear Audio*, May 2019, https://linearaudio.nl/extending-tube-life-hv-delay.

J. Didden, "Tube Amplifier High-Voltage Delay," *audioXpress*, February 2014, https://www.audioxpress.com/article/how-to-extend-the-useful-life-of-those-precious-tubes-use-a-tube-amplifier-high-voltage-delay

controller goes to sleep and switches itself off until the next time you cycle the mains voltage.

If you are new to controllers, or are unfamiliar with it, it may seem like a bit of magic, but it is simple engineering. The controller program is just a bunch of instructions that are executed by the chip after it wakes up at power-on, something like (in order of execution):

- Immediately set pin 7 high to keep the FET switched off.
- Read the voltage at pin 5 to determine whether the manual delay jumper at JP1 is placed. If not, get the saved delay time from memory and start it. If yes, start the delay and keep on checking the jumper. If it is removed, terminate the delay and save the delay for next time. Flash the LED during the delay.
- After the delay, wait until the secondary voltage at pin 6 goes through zero.
- Set pin 7 low to allow the FET to switch on.
- If the jumper at J5 is set, turn off the LED. If not set, turn the LED on to indicate 'active' state.
- Turn yourself off and go to sleep. (Poor controller—so much power, so few challenges ...).

Customization

The way I have set up the program is that the delay is 30 seconds out of the box. You can set it to your need as follows. Place jumper JP1 (or switch S1, see later text) before switching on the amp. Determine the delay you want to set, and divide that number by two. Switch on the amp and the LED will start flashing in a 2 second rhythm, so you count off the number you have previously determined. When you get to zero, pull off the jumper (or flip the switch).

Important note: Be careful when handling a jumper with the amp on, since high voltage is present on the board. You are responsible for your safety—not me, not the publisher. If you anticipate experimenting, you could replace the jumper with one of those small PCB toggle switches, they should fit the 0.1" pitch, but use a switch with a plastic or isolated lever. Jumpers are available with a little extra "handling strip" so you can remove them without touching the jumper itself. Better yet, use pliers with insulated handles.

You can also change the way the LED works. If you already have some on/off indicators on your amp, you may want to solder the LED directly to the board and use it only for checking/setting the delay. In that case, place a jumper at J5—the LED

will then blink during the delay only and after that go to sleep, along with the controller.

The Specs

There's only a few specifications. The absolute maximum voltage (that is the maximum at your supply cap, when switching on the amp, before it is warmed up!) is 600 V. This is where the professional manufacturers would say: "exceeding this maximum spec may permanently damage the device."

Maximum high voltage supply current that can be switched has no practical limit. If you have an amp that takes more than a couple of amps, you probably operate a transmitter and even then, no sweat! The delay time default is 30 seconds and can be changed to any value, in 2 second steps, between 20 and 254 seconds.

Construction

I have laid out a small PCB for this. **Figure 3** shows the stuffing guide. **Photo 1** and **Photo 2** give a 3D view of the completed small board; maybe I should call it a boardlet. Be sure to use a standard DIP08 socket for the controller in case you need to trouble shoot or replace it!

The dissipation in FET switch Q3 is so small that you can just solder it in an upright position to the top of the board as shown in Photo 1. Q1 may get warm to the touch at very high voltages, and if you want you can mount Q1 below the PCB to your amplifier chassis with a standoff as shown in Photo 2. Since the tab of Q3 is at Gnd, isolation is not strictly required, but mounting it on the chassis could impact the ground curents in your amp and is a reason to use isolation hardware. Then connect the board as shown in Figure 1, GND and CT to J2 and J1, and a wire from Vraw to the top of the (first) power supply capacitor, and you're good to go.

The Bill of Materials (BOM) for this project is shown in **Table 1**. The LED that comes with the diyaudio.com kit is a very low current, high brightness, red clear lens one. Feel free to replace with your favorite color, but do not replace the LED series resistor (R5) because increasing the LED current will significantly increase the dissipation in Q1. If the brightness is too high for your liking, you can place a resistor in series with one of the LED leads to dim it.

Availability

Our friends at diyaudio.com agreed to make a kit available for this project, consisting of the PCB, the programmed controller chip, the two power MOSFETs and the low current LED. You can order it at https://diyaudiostore.com/collections/kits. All other parts are available and/or in your parts bin.

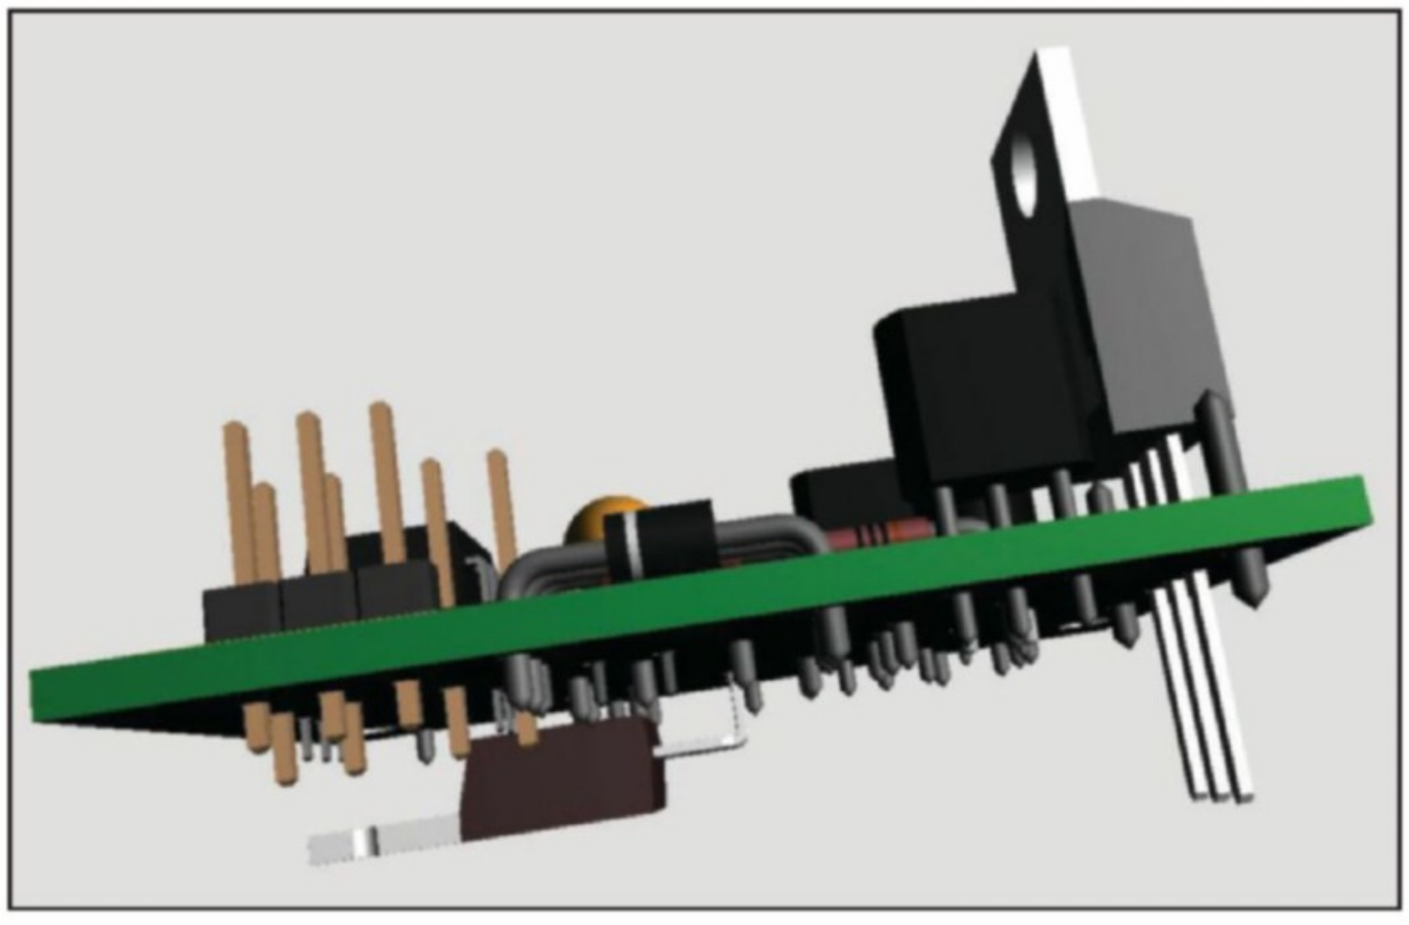

Photo 2: Mounting Q1 on the chassis in 3D view

Quantity	References	Value	Comment
3	C1-C3	100 n/25 V	0.1" pitch ceramic
1	C4	33 µ /16 V	0.1" pitch
2	R2,R8	1 M	1 W
1	R3	100	1/4 W
1	R4	4.99 kΩ	1/4 W
1	R5	15 kΩ	1/4 W
1	R7	10 kΩ	1/4 W
1	R9	56 kΩ	1/4 W
1	R10	680	1/4 W
1	U1	PIC12F629	programmed controller (part of kit) Use DIP socket!
1	Q1	IXTP08N100D2	part of kit
2	Q2,Q4	BC546BP	or equivalent small signal TO-92
1	Q3	FDP12N60NZ	part of kit (see text)
1	D1	Zener, 3.3 V	400 mW
3	D2-D3,D6	1N4148	or equivalent small signal
1	D4	LED	low current, Mouser 604-WP710A10SEC/J3 (part of kit)
1	J1	Xformer CT	off-board connection
1	J2	GND	off-board connection
1	J3	Vraw	off-board connection
2	J5-J6	2-pin hdr 0.1"	LED header
1	JP1	JUMPER	3-pin header (see text)
1	S1	SPDT	optional (see text)

Table 1: Here is the Bill of Materials (BOM) for the High-Voltage Delay project.

The Tenacious Bass-8 Powered Subwoofer

In last month's October issue of *audioXpress*, Thomas Zarbo, Jr. detailed his Bantams microspeaker system. This month, he shares the perfect accompaniment to that system, a subwoofer with a passive radiator that he has dubbed the Tenacious Bass-8.

By
Thomas Zarbo, Jr.
(United States)

Photo 1: The Tenacious Bass–8 powered subwoofer is pictured with the Bantams microspeaker system that was detailed in the October 2019 issue of *audioXpress* sitting on top of it.

Since I was already blazing new territory in designing my first truly tiny speaker with the Bantams speaker system detailed in the last issue of *audioXpress*, I decided that for my next project I would try to design a subwoofer with a slot vent, which I had never previously attempted. As you will see, however, an issue with my design revealed itself part-way through construction of the enclosure, necessitating a partial rework of the project. Thankfully, with some good advice from a fellow DIYer and just a few extra hours of work, I was able to turn lemons into lemonade—and the end result was the solid performing Tenacious Bass-8 subwoofer (see **Photo 1**).

Photo 2: Front and rear view of the Tang Band W8-740P 8" subwoofer. Note the massive bumped back plate, thick mounting flange, and large foam roll surround.

The Design

With my goal in mind, I set out to find a suitable 8" driver and eventually settled on the Tang Band W8-740P 8" subwoofer. The name of this project, the "Tenacious Bass–8" is a nod to the special Tang Band driver at its heart.

As you can see in **Photo 2**, it's an aggressive-looking monster of a driver with its huge vented, and bumped back-plate, as well as extensive venting in the frame. It has a 2" voice coil employing an aluminum former, and sports a full 12 mm of excursion. After some modeling, it appeared that this driver would work well in a small slot-vented enclosure with an F3 just under 30 Hz. (F3 is the

point at which the system response is 3 dB lower than reference—or where a person is likely to begin to notice a reduction in output)

I designed a slot vent with openings between the main cabinet and a thick base on the exterior, and a built-up plate amp sub-enclosure/vent combo unit on the inside (see **Photo 3**).

The plate amp I chose was the Yung SD200 (no boost) subwoofer plate amplifier. It is a Class-D style amplifier that has the ability to play at its full-rated output over long periods of time. It is fairly compact as well—with a depth of only 2.25"—which makes a minimal intrusion on cabinet volume.

After the cabinet was mostly completed, I temporarily clamped the vent/bottom base assembly to the box and played some bass test tones. To my dismay, I found that while the setup did produce a great deal of deep bass, I was also treated to wind buffeting coming from the vent that was comparable to cracking open a car window at speed. In talking the issue over with a friend who also uses BassBox Pro to model enclosures, we eventually discovered that my box modeling program had become corrupted after a recent update, and was spitting out faulty information.

The graphs told me all was well, but reality told me something far different—this vent was vastly undersized, and the vent air speed was excessive, causing non-musical turbulence, or "chuffing" as the air exited the vent. In reality, a little vent turbulence is acceptable, as there are a multitude of other sounds being reproduced in conjunction with the low bass—masking any moderate vent noise that may occur. My vent had apparently taken "chuffing" to another level however. This design needed an overhaul.

Discussing my problem with a Parts Express "Tech Talk" forum member convinced me to convert this project to a passive radiator design, which after some modeling appeared to be the perfect solution. As it turns out, utilizing the Dayton RSS265-PR 10" Passive Radiator in the same size box yielded a very similar F3 with no chance for vent noise (see **Figure 1**). A reworked slot vent would have made the cabinet substantially bigger and I wanted a more compact form factor for this design.

Modifying the Cabinet

Rather than start all over, I ended up slicing off the entire top portion of the cabinet on my table saw to cut about an inch down from the top to access the inside and remove the vent, which I did with a vibrating multi-tool (see **Photo 4**). While the "lid" was off, I added an additional 3/4" thick piece of 10–7/16" square MDF to the inside on the bottom—

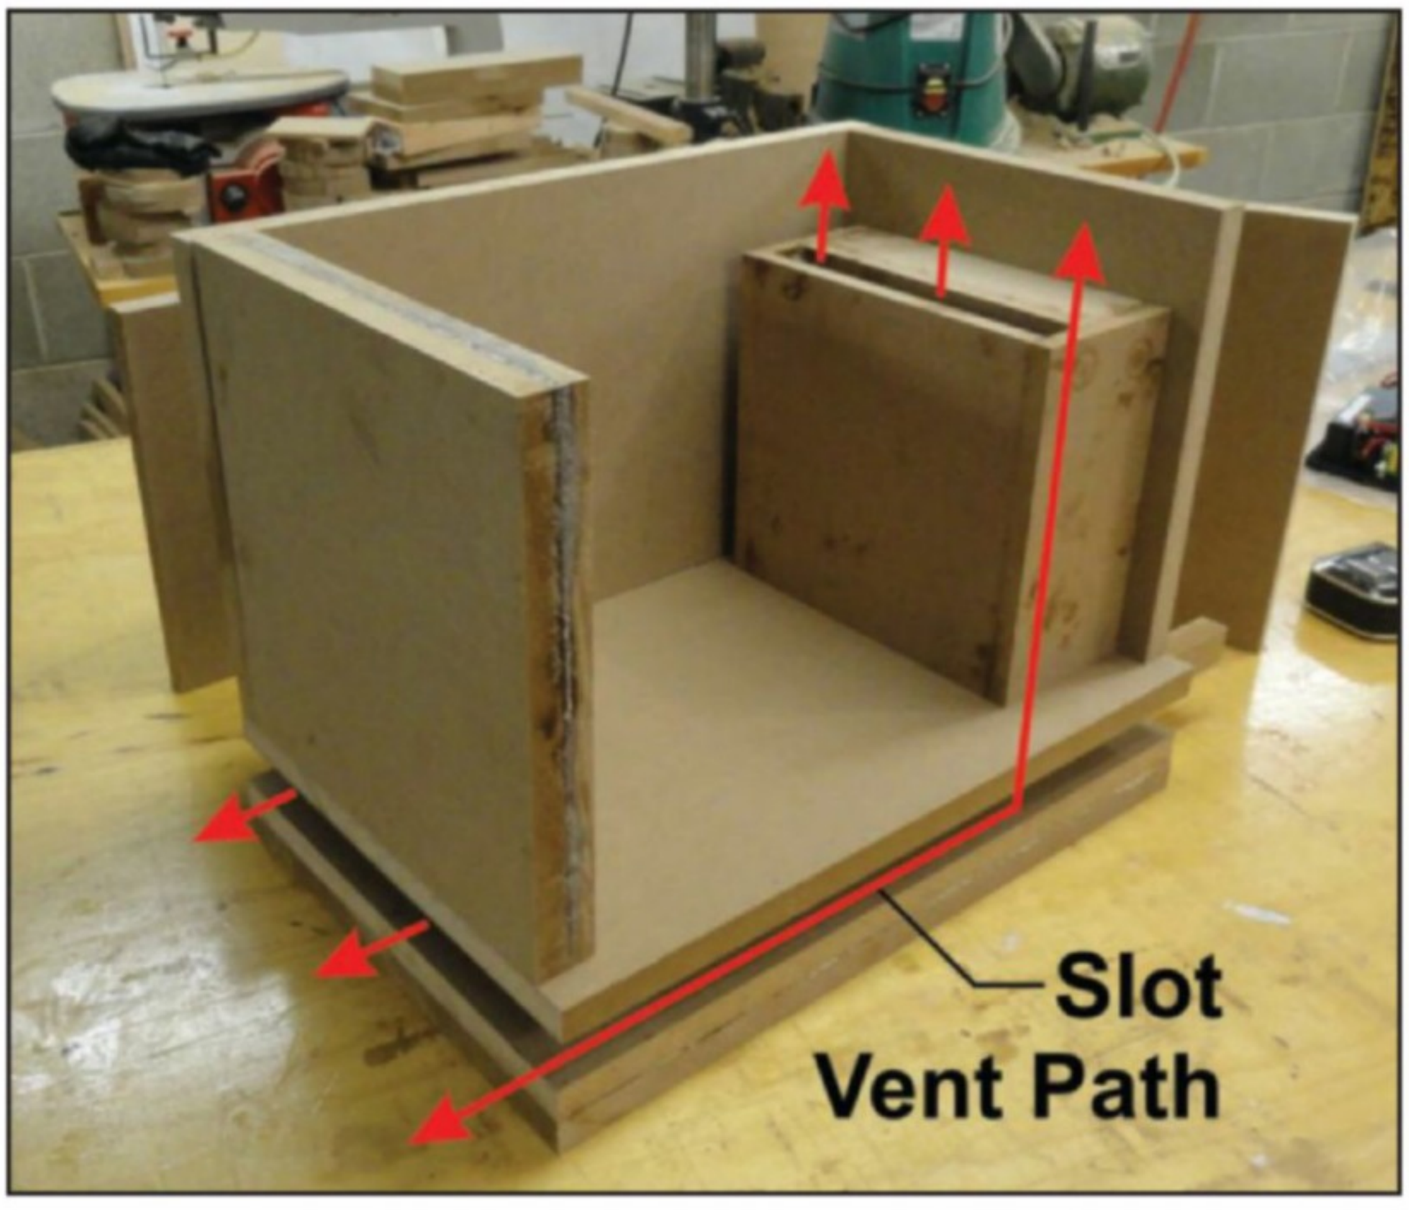

Photo 3: A side view of original slot vent arrangement. Inside the cabinet, the vent structure also serves as part of the housing for the Yung plate amplifier.

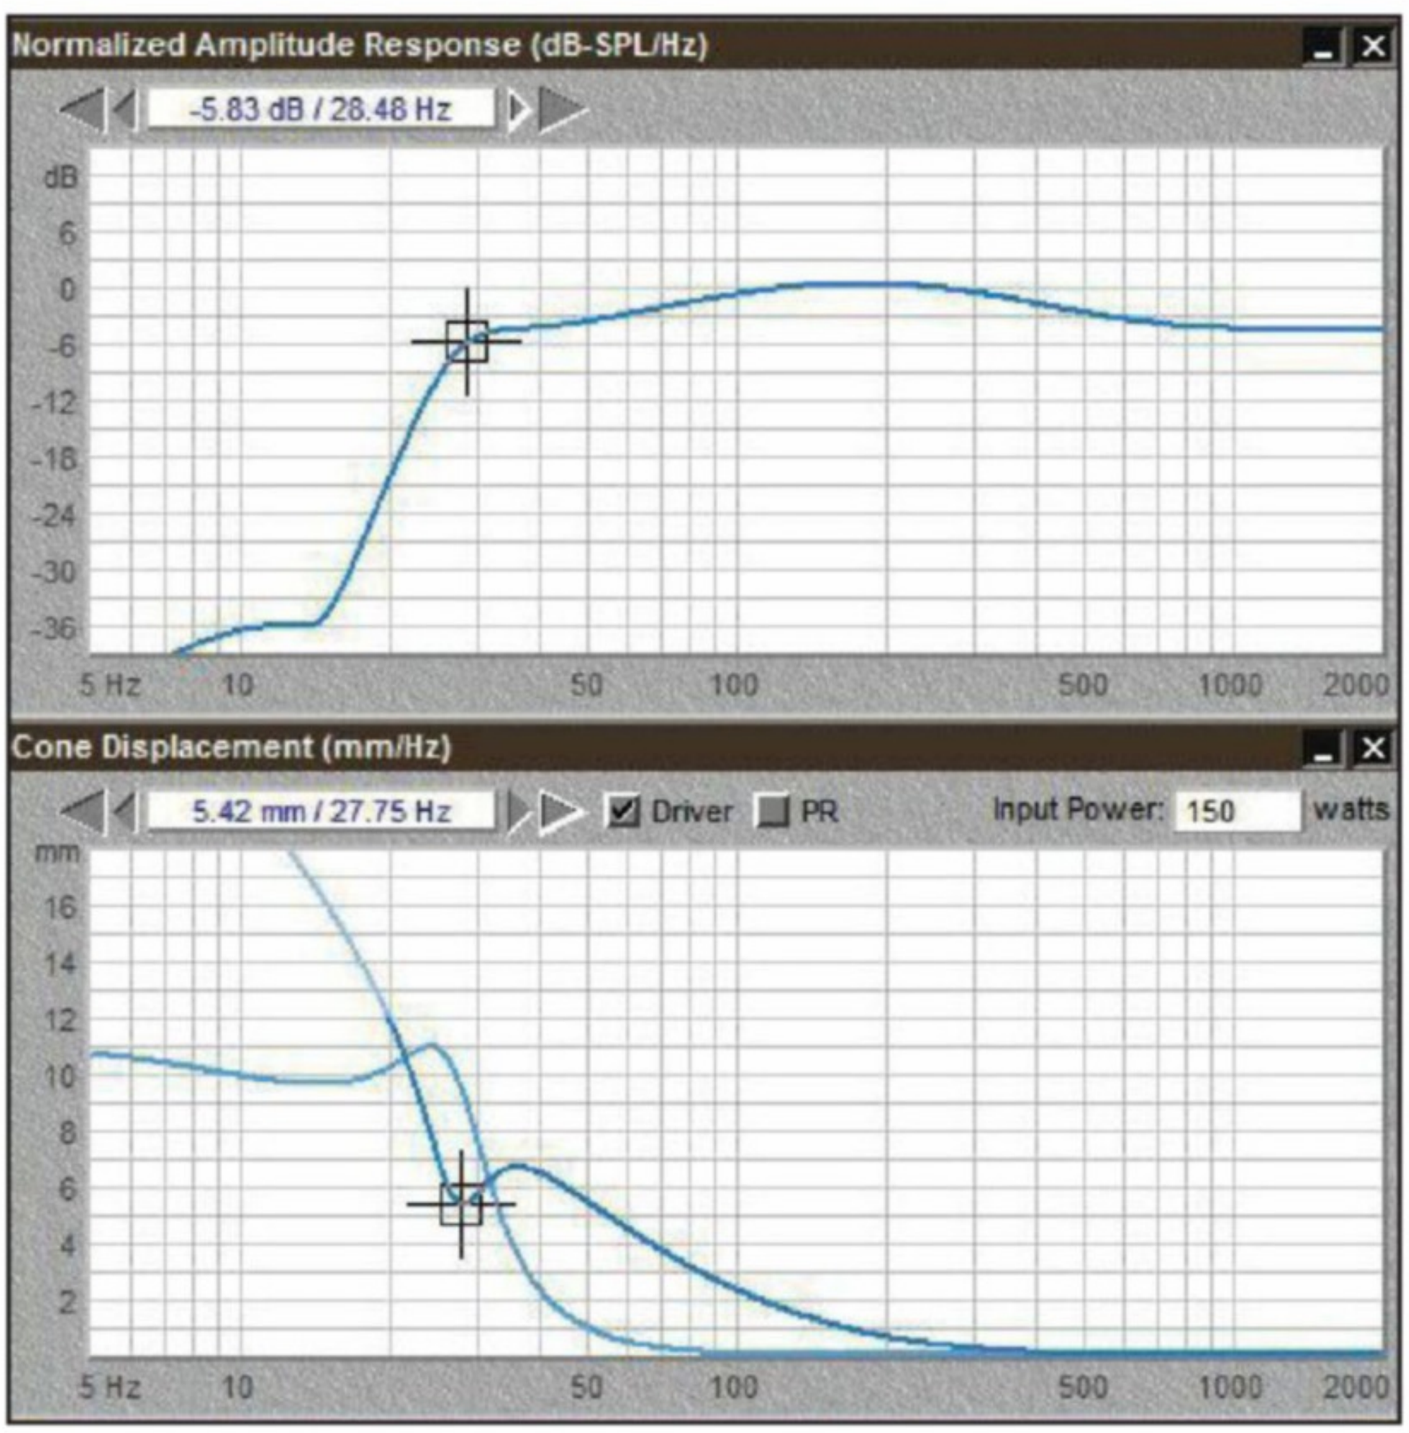

Figure 1: The F3 of the system is around 28 Hz when the lower "reference" frequency of 50 Hz is considered. Woofer excursion doesn't breach Xmax until 20 Hz with 150 W. Note that the driver excursion minimum nearly coincides with passive radiator maximum.

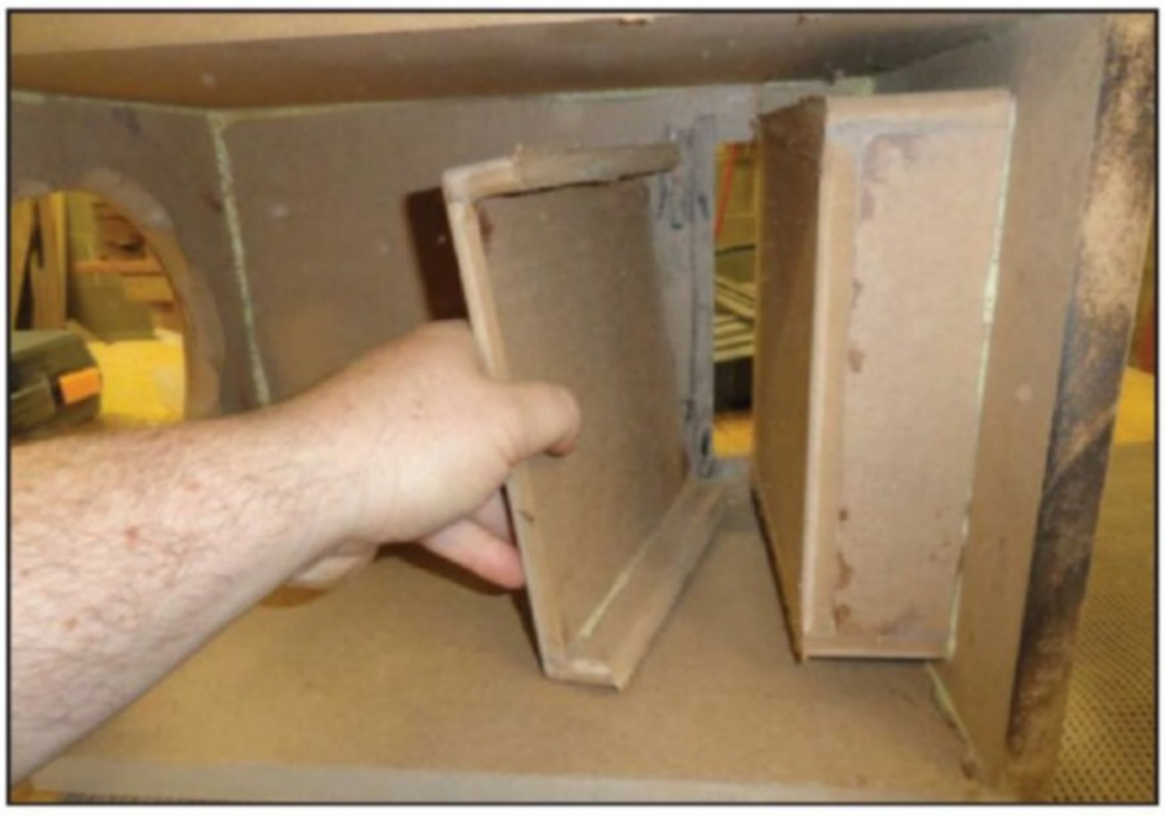

Photo 4: Removing the inside portion of the slot vent from the plate amp sub-enclosure. Note the rectangular exit on the bottom of the cabinet that will need to be filled in.

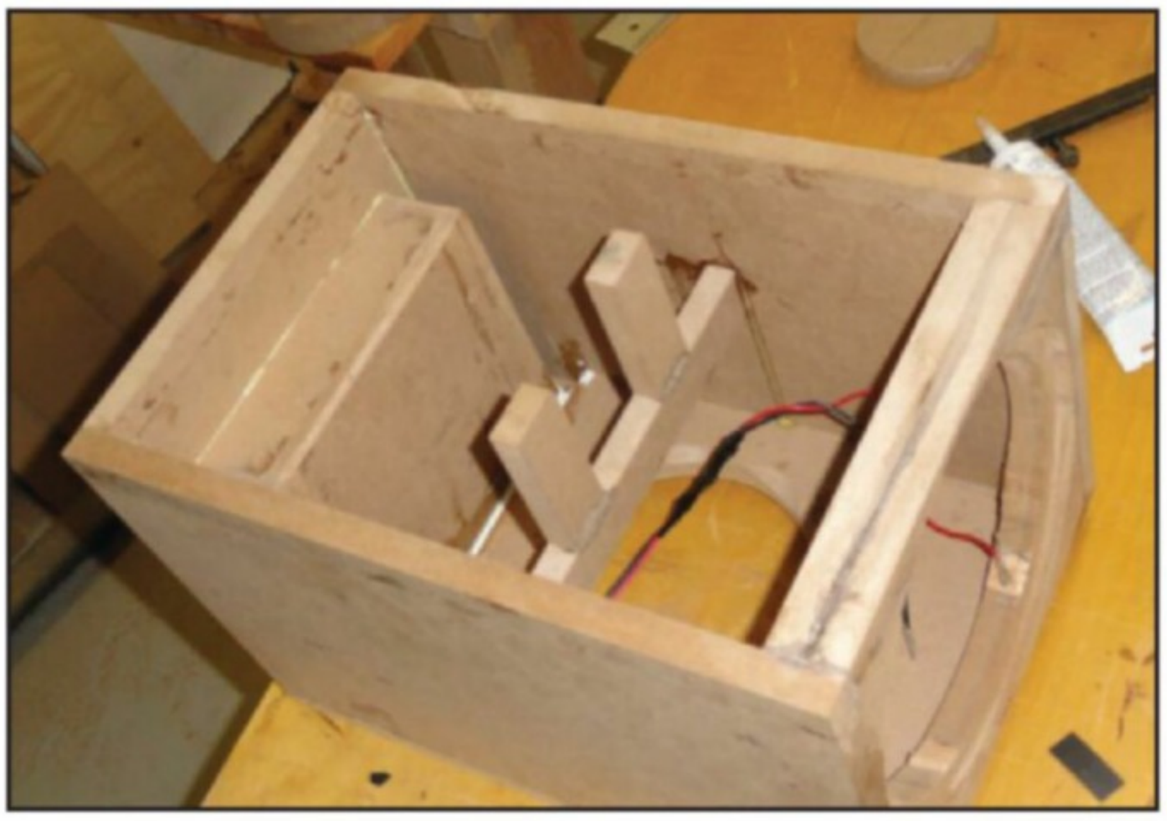

Photo 5: The sole brace structure in the cabinet consists of a 1.5" piece of 3/4" MDF across the width of the cabinet, along with two 1.5" wide 3/4" MDF braces that join the top of the cabinet to the horizontal brace.

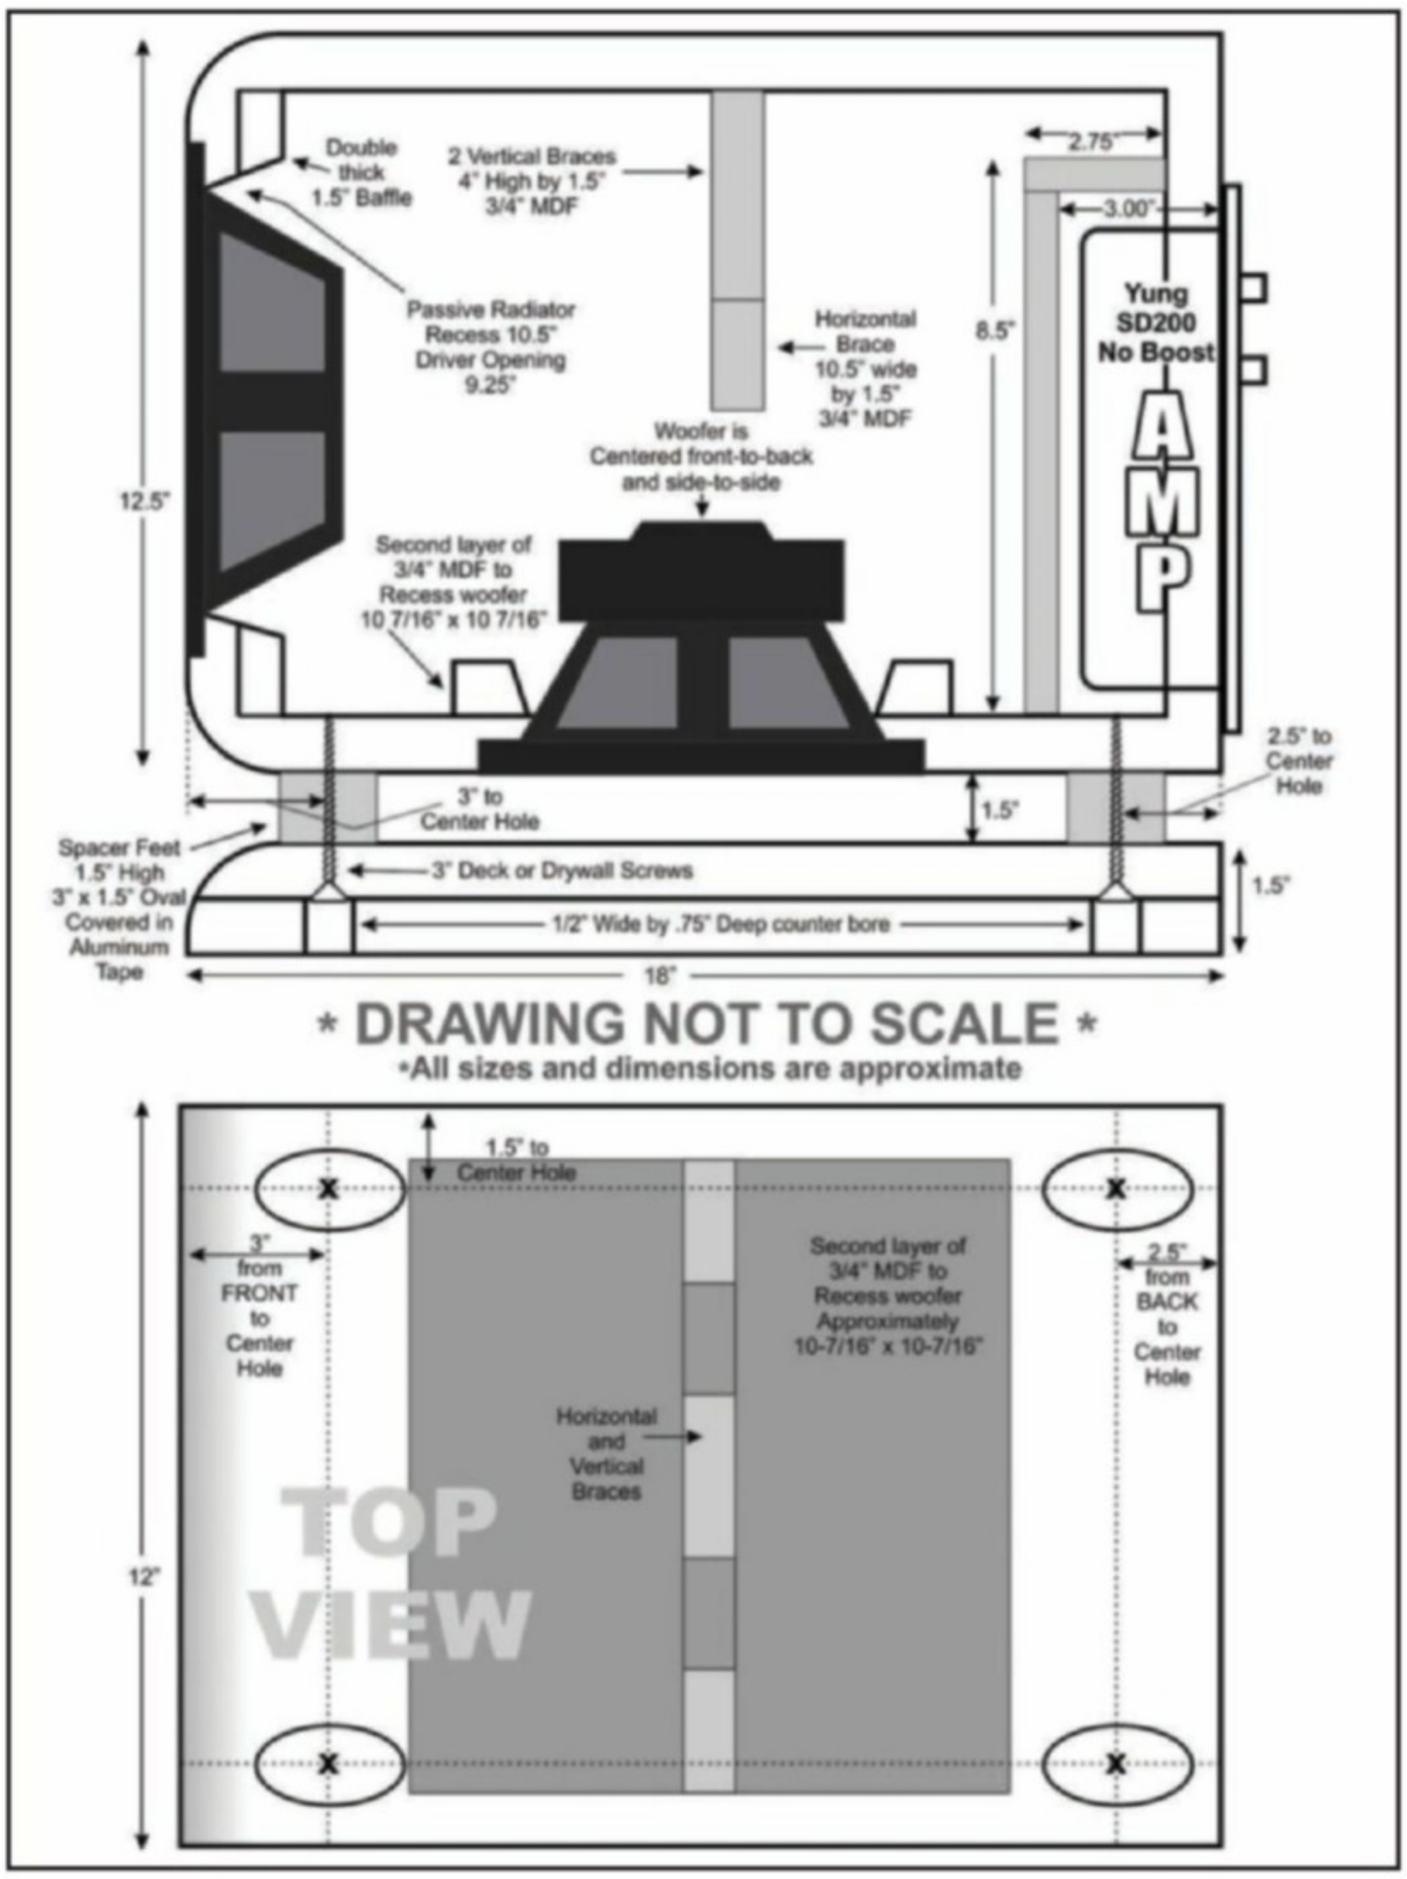

Figure 2: These are the dimensions for the Tenacious Bass–8 subwoofer.

the woofer's new location—to allow it to be recessed flush with the cabinet. Before attaching the new top, I glued in some additional 1" tall pieces of 3/4" thick MDF to the cabinet to raise the height so that the large 1.25" round-over I wanted to create wouldn't intrude on the 10" passive radiator that would now be installed on the cabinet front. Other than that, the cabinet would be roughly identical in shape to the original compact design.

As a rule, it's usually a bad idea to mount larger woofers facing down (or up), as their weight can cause them to sag over time. To see if a particular driver is suitable for mounting this way, Parts Express has a section in its "Technical Resources" area at the bottom of the website which, when clicked on, reveals another link—"How to tell if a woofer can be mounted up- or down-firing." That link supplies the calculation based on the drivers Thiele-Small (T-S) parameters to determine if that driver is suitable for "down-firing"—which as I suspected with its extremely taunt suspension—was no problem for the Tang Band.

Realizing that the pressures inside the cabinet will be fairly intense, I created a small brace system that links the top and sides of the cabinet together (see **Photo 5**). This simple system links the largest unsupported panels together, and makes the cabinet largely inert. The reworked cabinet including the base ended up being 18" deep, 12" wide, and 15.5" high (see **Figure 2**). The large 1.25" round-overs were machined on the assembled cabinet, using my home-built router table in several careful passes. I realize that not everyone has such a large router bit in their toolbox, or would want that much sharp

steel spinning inches from their hands, but I really wanted this particular curve on this cabinet as I feel it adds a bit of style to the design. Also, the larger round-over made it easier to bend the veneer over the curve, since I decided to have the length of the grain run perpendicular to the curve. Pre-made MDF round-overs are available for purchase online in lieu of machining your own.

The Sub-Enclosure

Although the Yung plate amp is designed to be air-tight, I made a sub-enclosure for it out of 1/2" MDF, hoping to reduce the strain it will face over the long-term. The large round-over is strictly for aesthetics, straight corners will work fine as well. As previously noted, the cabinet has a double-thick 1.5" baffle area for the active driver, as well as the passive radiator. I wanted to recess these into the cabinet, but the flanges for both are so deep, that using a single layer of 3/4" MDF would have left precious little material for the mounting screws to bite into (see **Photo 6**).

Next, I set to looking for a snazzy way to lift the cabinet off of the base now that there would be no slot vent to serve that purpose. I remembered

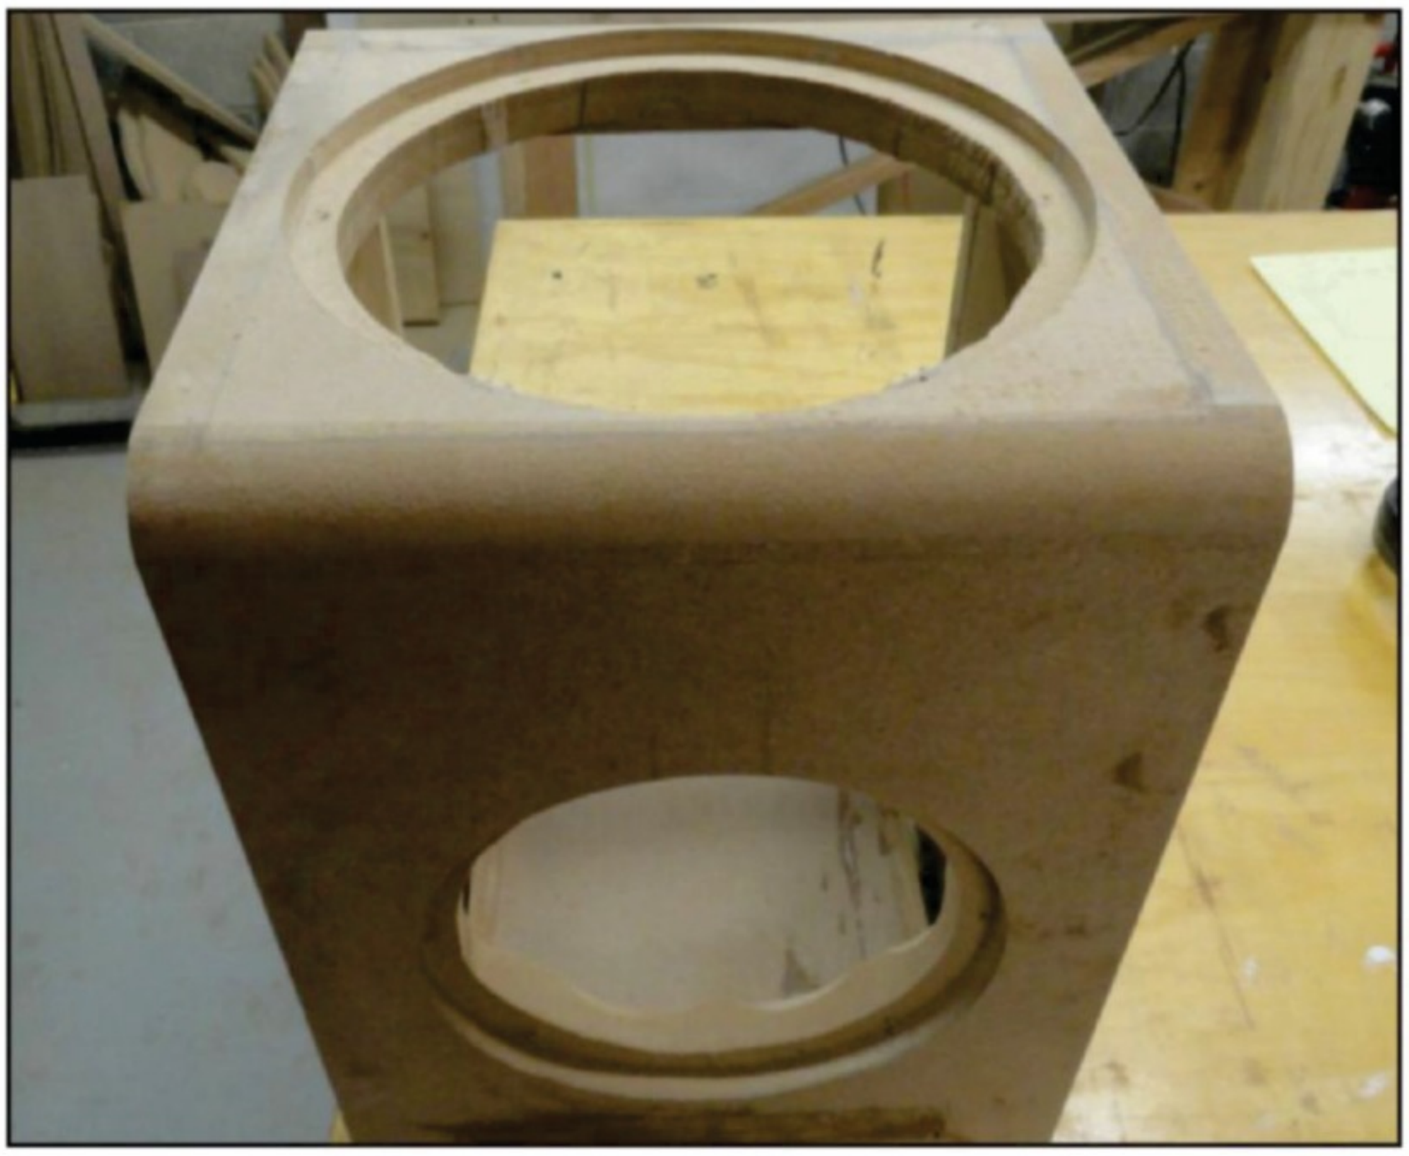

Photo 6: Tenacious Bass–8 cabinet has a large round-over, driver recesses, and scalloped areas for the active driver to "breathe." Also visible is the patch on the bottom, filling in the area where the slot vent originally exited the main cabinet.

that a friend of mine from the Tech Talk forum had creatively used stainless steel sign standoffs for a similar use on a set of speakers he had built. I loved how they looked, but after pricing them out online I discovered they were somewhat expensive. I ended up making my own version out of a few layers of plywood, cut on the bandsaw, and sanded into an oval shape that I thought matched the longer depth of the cabinet better than round feet would have. I covered these with aluminum flashing tape

Photo 7: The four feet that raise the cabinet off of the base are constructed of two layers of 3/4" plywood, cut on the bandsaw and sanded into an oval shape. These were covered with aluminum tape, which gives the approximate look of stainless steel or brushed aluminum at a fraction of the cost.

Photo 8: The Dayton RSS265-PR 10" Passive Radiator has a strong cast-aluminum frame and 24 mm (nearly a full inch!) of maximum one-way mechanical movement (Xmech). Four of the five included 75 gram discs are used. The washers are secured with the included lock-nut.

purchased from the hardware store. They are screwed through the base, into the main cabinet with four long screws (see **Photo 7**).

The Dayton RSS265-PR 10" Passive Radiator comes with five 75 gram weights for various tunings. The webpage for this product conveniently lists individual T-S specifications for this radiator, depending on how many weights are used for easy and accurate modeling with software. I ended up using four of the included washers for a final tuning of around 26 Hz (see **Photo 8**).

The cabinet parts were covered with paper-backed walnut veneer using the "iron-on method" that I prefer. I machined the passive radiator opening and recess in the enclosure before applying the veneer, as I wanted to test the concept before finishing—but it would be much easier to apply the veneer first, then machine the opening afterward, as getting the veneer to lay flat with such a large opening proved difficult. For a durable finish, I applied 10 coats of rub-on polyurethane for a nice, even shine (see **Photo 9**).

Before installing the passive radiator, I painted the driver recess black as seen in Photo 8 so there wouldn't be any distracting, light-colored bare

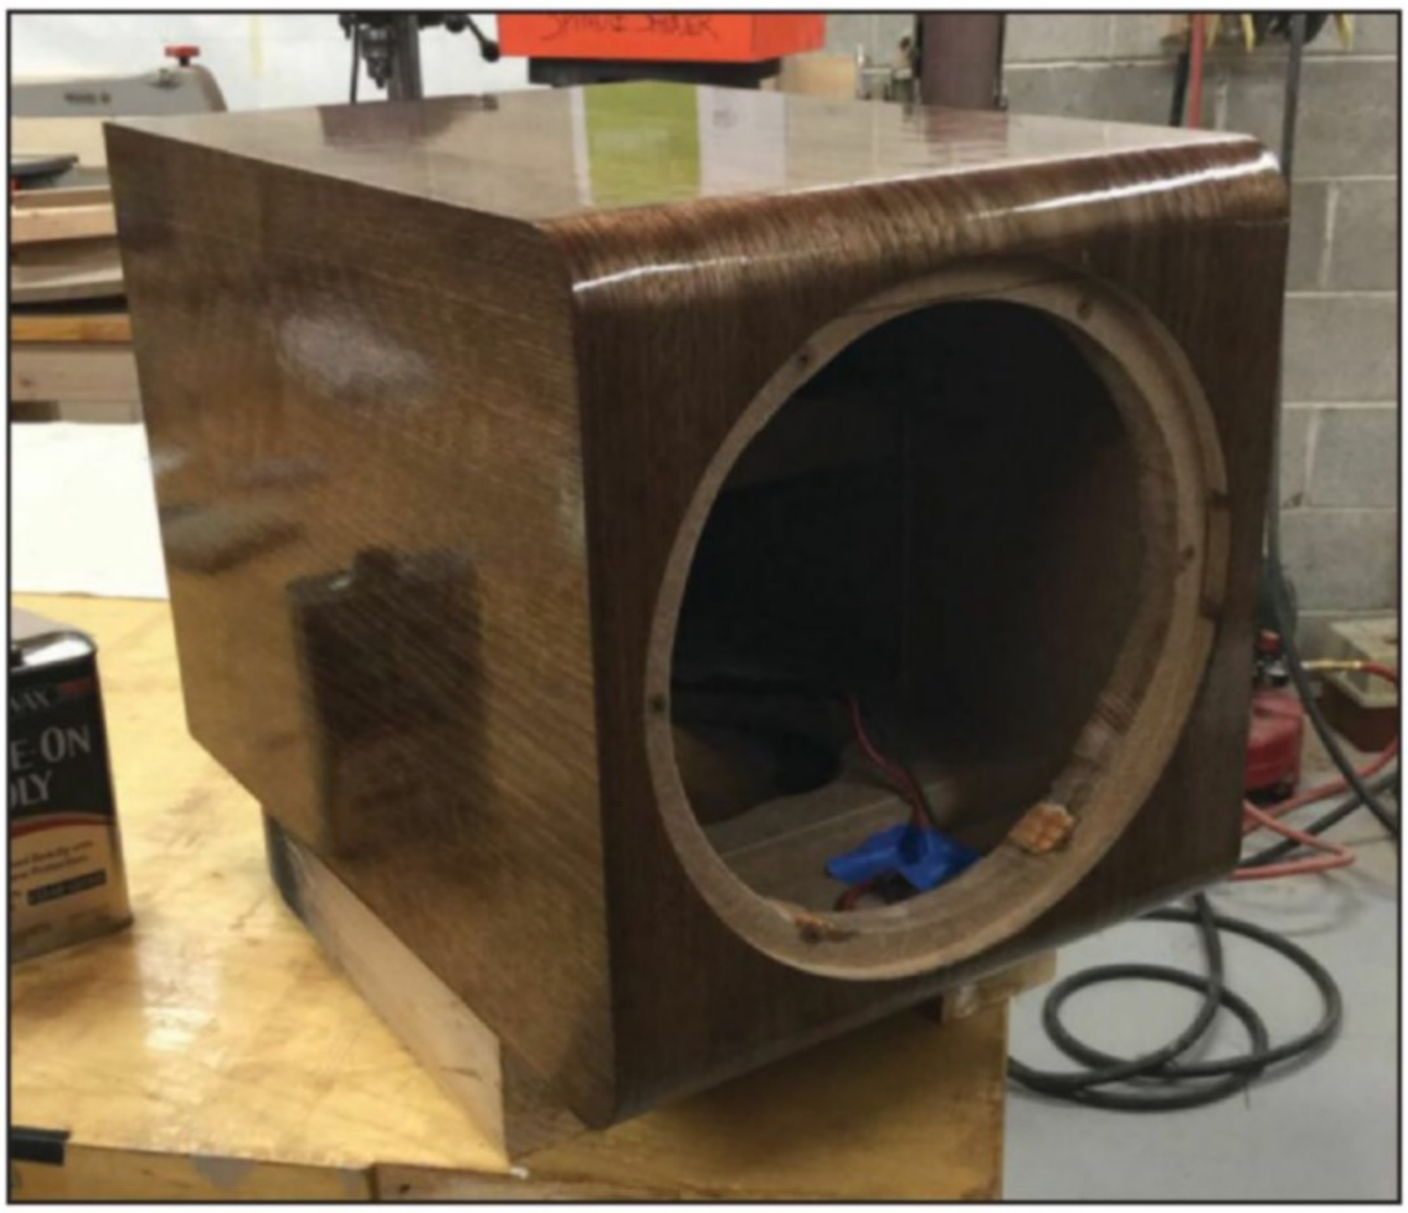

Photo 9: The finished cabinet is nearly ready for assembly. The opening for the passive radiator takes up almost the entire front of the enclosure.

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

Spiral Air™
Interconnects

Coax
Is Cool
Again!

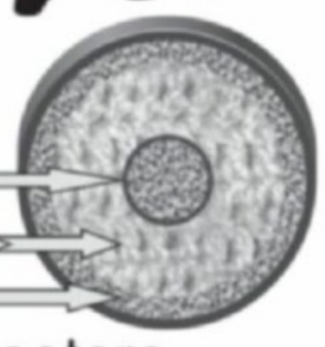

Cardas Golden Strand
Air Filled Cotton Dielectric
Unique Spiral Litz Shield
ETI Tellurium Copper Connectors

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

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

Photo 10: The Bantams mini-speaker system are shown with Tenacious Bass–8 powered subwoofer. This combination is small enough to fit nearly anywhere, but provides legitimate full-range sound at satisfying listening levels.

About the Author

Thomas Zarbo, Jr., is a preschool teacher living in southern New Jersey, and he has been building speakers since he was a teenager. Thomas' interest in speaker building began when he and his father took apart an old stereo console unit and transplanted the speaker drivers in some home-made enclosures. He said he was instantly hooked on the hobby. His love of DIY things "audio" has grown to also include various smaller amp and radio projects. Thomas said he is thankful for the kind folks on the "Tech Talk" forum who have helped him over the years gain a good understanding of how the various aspects of speaker design and construction fit together. Thomas added that he was grateful for their continued help and guidance, as well as to Parts Express for hosting the forum, and *audioXpress* for promoting DIY speaker building.

Sources

Dayton RSS265-PR 10" Passive Radiator
Parts Express | www.parts-express.com

Yung SD200 200 W Subwoofer Plate Amplifier (no boost)
Parts Express | www.parts-express.com

Tang Band W8-740P 8" subwoofer
Tang Band Speaker Co., Ltd. | tb-speaker.com

MDF visible between the cabinet recess and the passive radiator frame. A Tech Talk forum member shared this useful tip with me several years ago, and it's little details like this that can help make a DIY project look every bit the equal of a store-bought product.

The Yung plate amp easily mounts with 10 #6 black oxide screws. I extended the wires on the amp through its wood housing and sealed the wire exit with silicone sealant. The amplifier has three knobs for adjusting the volume level, low-pass crossover frequency, as well as phase. Upon firing up the subwoofer, I adjusted the volume until it sounded proportionate with the music coming from the Bantams, then dialed the crossover frequency knob until I could no longer hear any localizable output. I then readjusted the volume a bit more to re-blend the sound to my liking.

Performance

So, how does this subwoofer perform? It shakes the walls in the house! You can feel the bass this little monster puts out. I wouldn't place it in the super-loud category, but it has "loud" covered all day long. Although it's not listed in the manual or product webpage, it appears the Yung amplifiers employ a second-order (12 dB per octave) subsonic filter fixed at around 30 Hz, so there isn't much 20 Hz content, but it plays strong down to nearly 25 Hz or so.

Although the Bantams and the Tenacious Bass–8 subwoofer were not designed from the start to be a complete 2.1 system—they do complement one another very nicely. On first glance, it might seem that this larger subwoofer is somewhat overkill for such a small speaker pair, but the Bantams really deliver a surprisingly clean and loud listening experience for their size. Their lack of output in the lowest octave is powerfully compensated for by the low-tuned Tang Band 8" subwoofer, and the pair just seems to mesh seamlessly with one another—I only hear the Bantams while listening to music.

Together, they make up a truly full-range system that my wife and I have been enjoying for several years now. They are my wife's pick for the best-sounding speaker system I have designed to-date—and she has quite a few from which to choose (see **Photo 10**).

For more information on these projects, visit the Parts Express website and click on "Project Gallery" at the bottom of the page. A search for "Tenacious Bass" or "Bantams" will take you to the respective write-ups for each project. Additional information is also available at the Parts Express Tech Talk forum if you perform a similar search.

Build your dream system

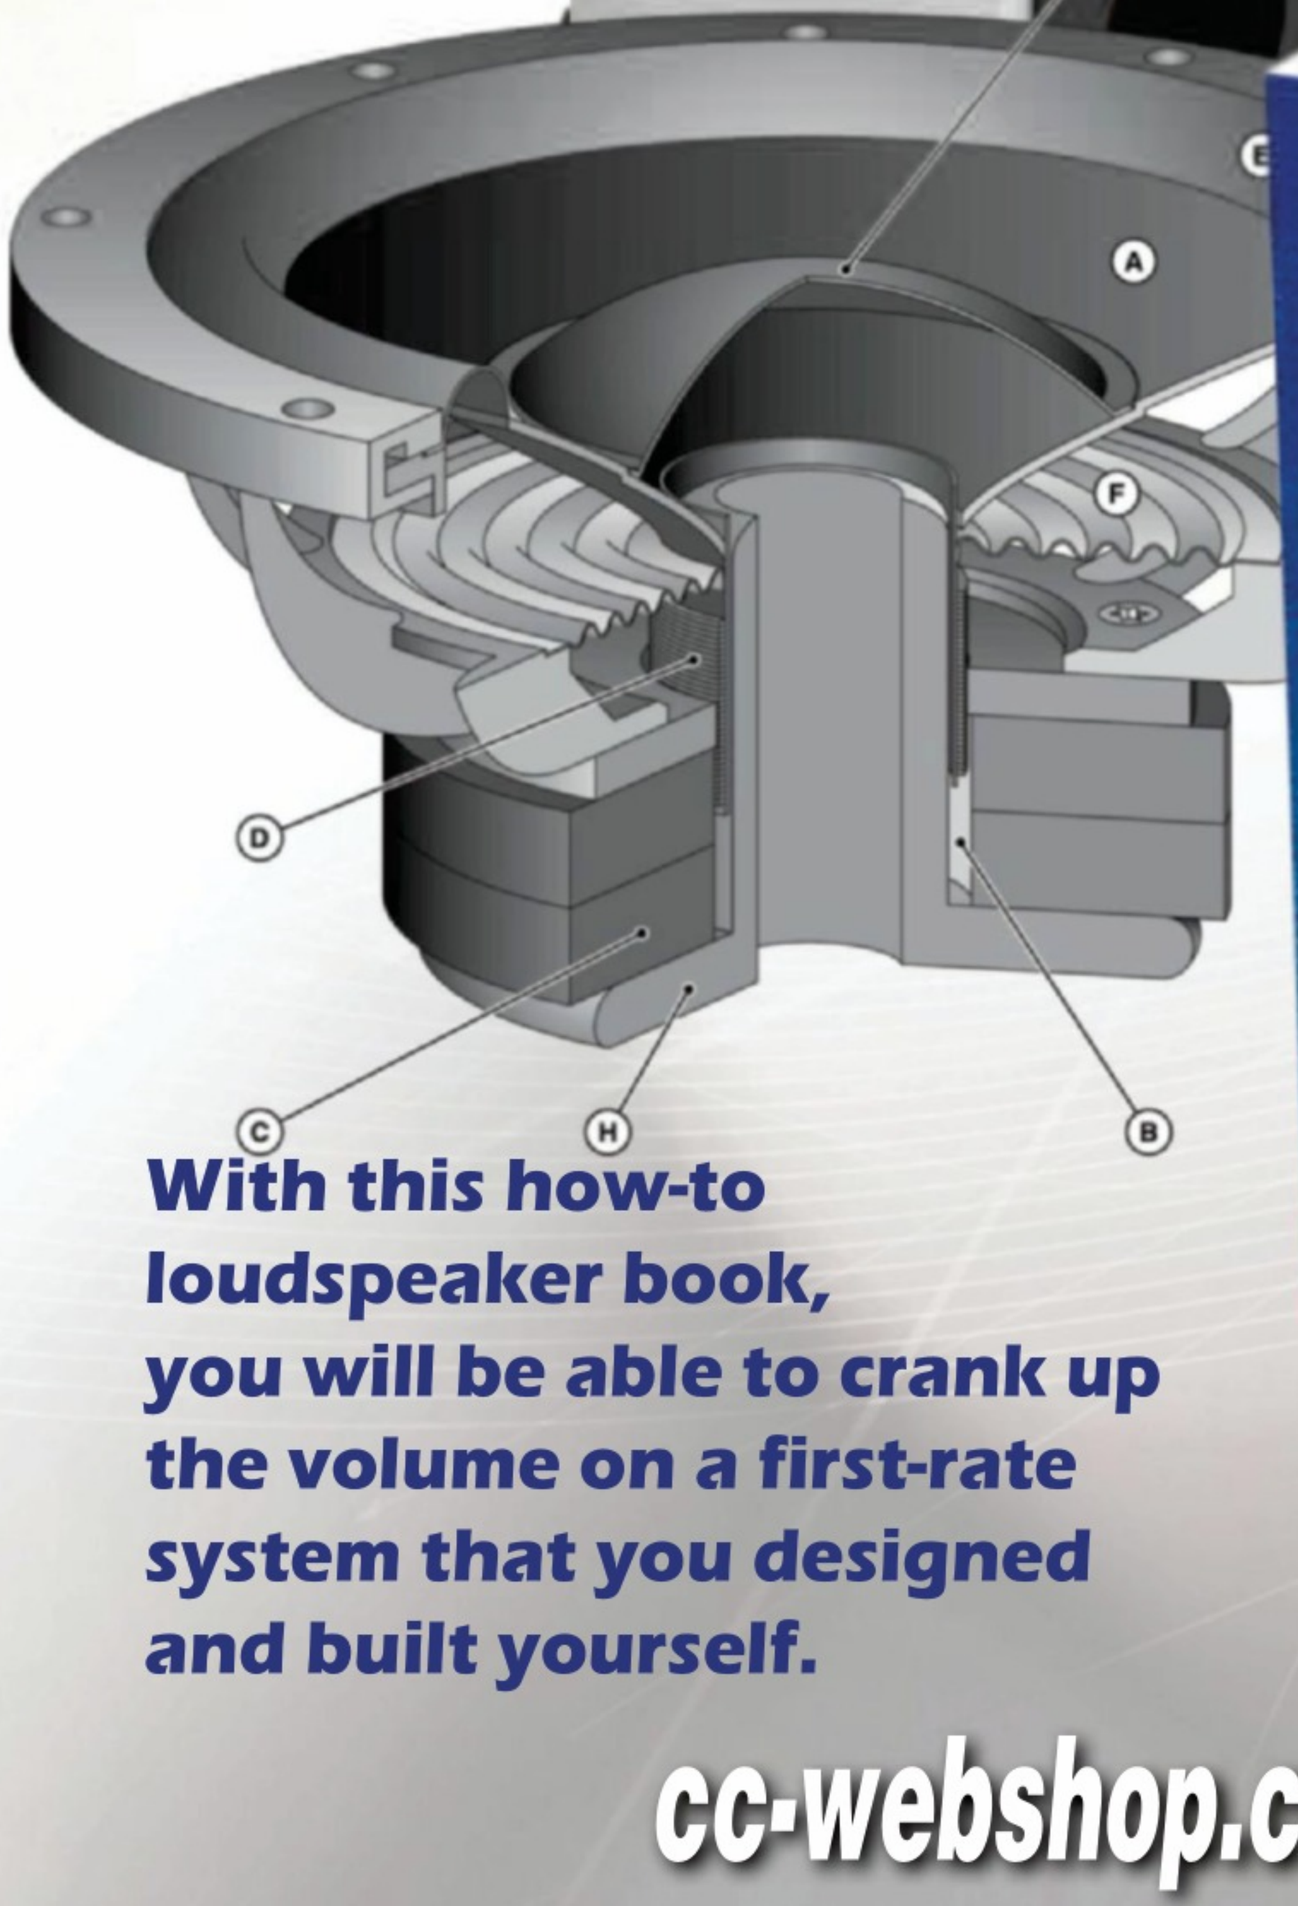

With this how-to loudspeaker book, you will be able to crank up the volume on a first-rate system that you designed and built yourself.

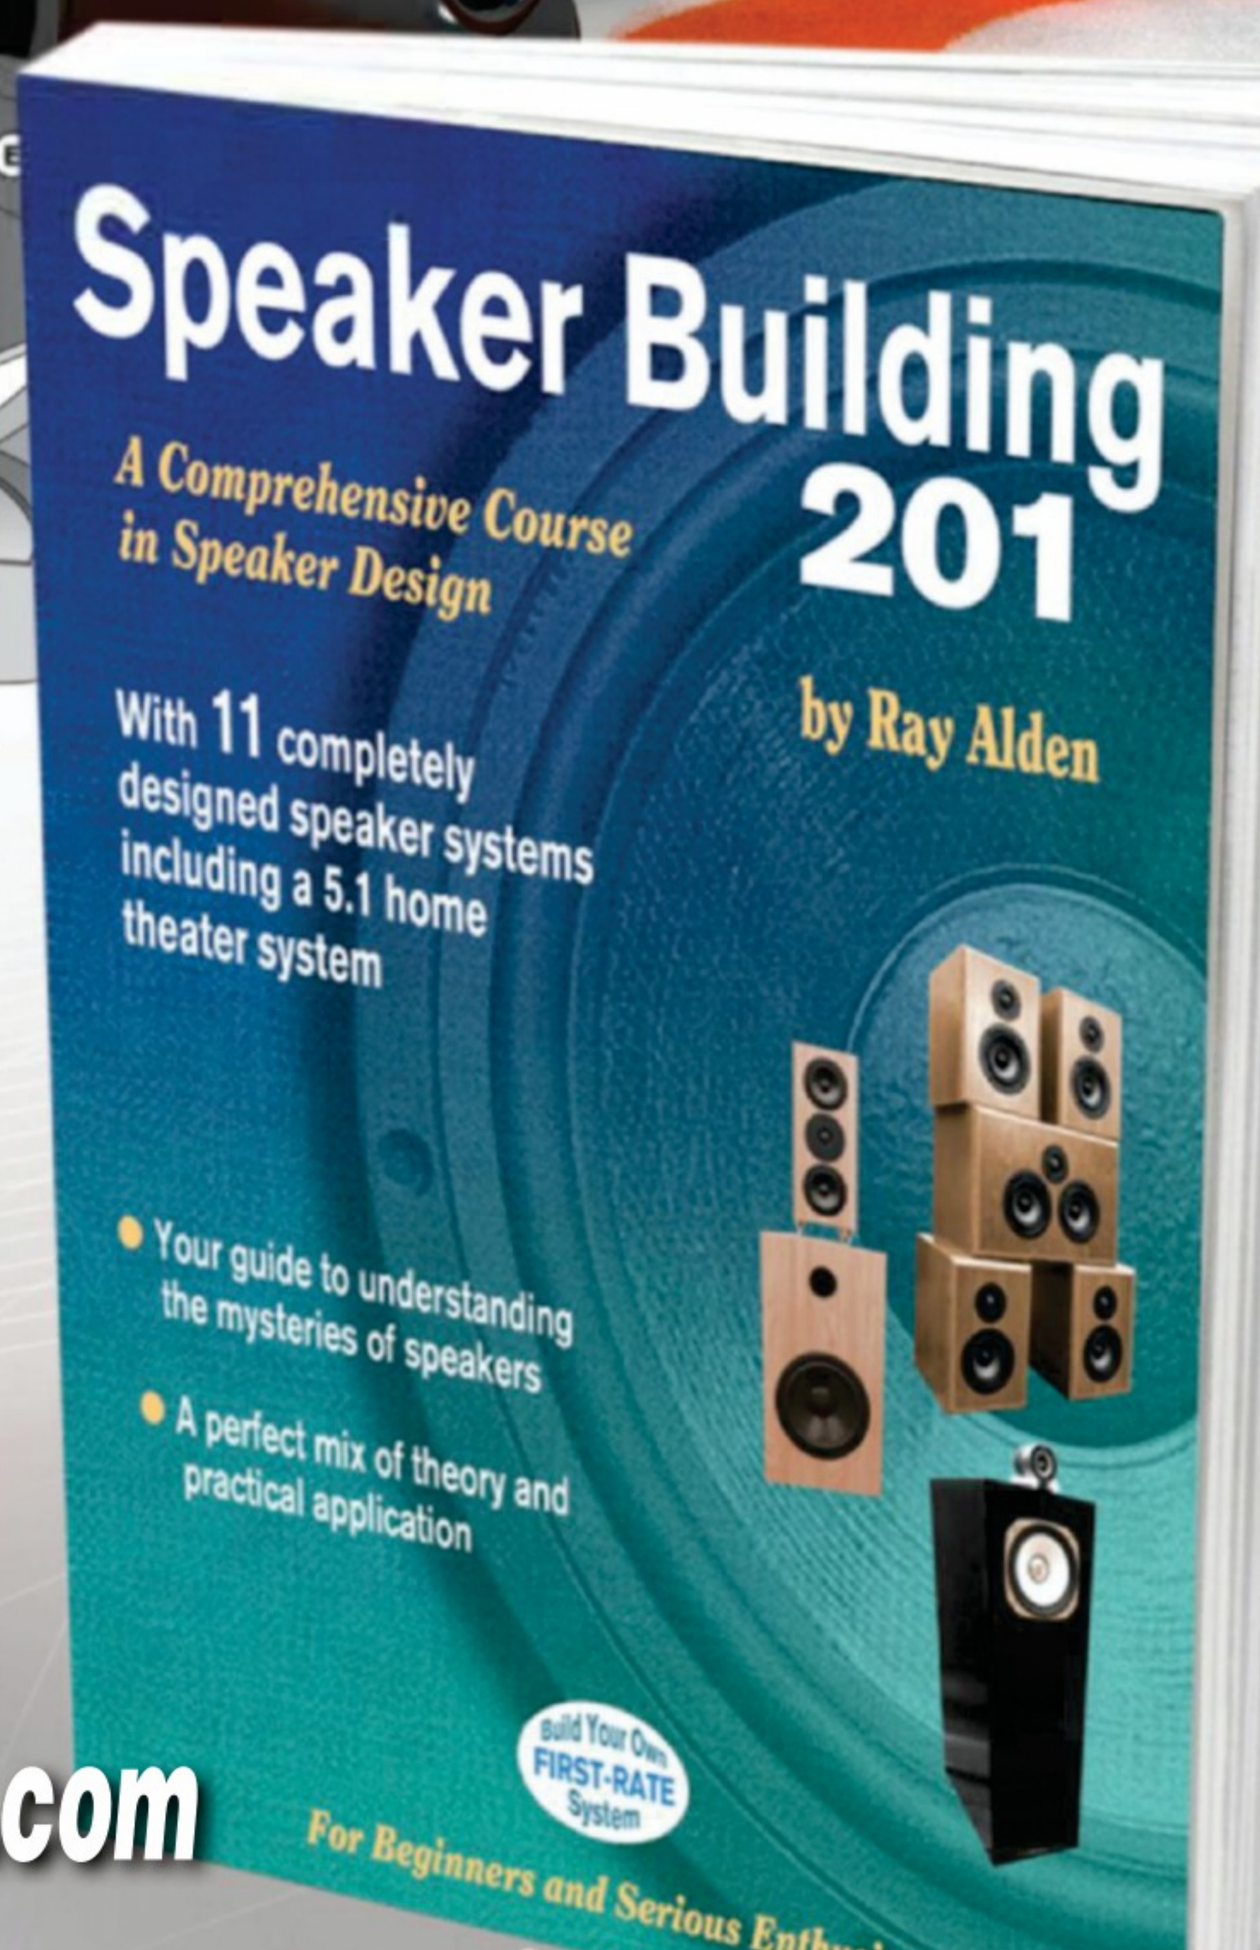

cc-webshop.com

The Hammond Growl

In this article, Richard Honeycutt discusses the Hammond organ's distinctive "growl" and attempts to decipher the origins of its unique sound.

By
Richard Honeycutt
(United States)

Photo 1: The high-frequency driver-horn assembly of a Leslie 122 speaker/amplifier.

When Laurens Hammond invented the Hammond organ, his intention was to create an electromagnetic instrument with the sound of a pipe organ. Because a Hammond organ cost much less than a pipe organ, many churches and theaters bought them. As discussed in the Hollow-State article printed in the October 2019 issue of *audioXpress*, features were added to the instrument over the years, including the unique speaker introduced by Donald Leslie in 1941.

The Growl Gains Popularity

The Hammond organ was discovered by musicians of other genres, notable among whom are Booker T. and the MGs, who's 1962 hit "Green Onions" hooked listeners from the very beginning with its Hammond organ vamp. In the late 1960s, organists such as Keith Emerson of Emerson, Lake and Palmer; and Steve Winwood of the Spencer Davis group established a sound called the "Hammond growl" that has been an important element in a multitude of pop, rock, blues, and jazz songs ever since.

With the advent of synthesizers as substitutes for Hammond organs, keyboardists started trying to recreate the Hammond growl through the new (and much lighter-weight) instruments. Although some manufacturers and players of synthesizers have come close to the original sound, many keyboard players maintain that the pot of gold is still out of reach: The only way to match the Hammond growl is to use a real Hammond organ with a

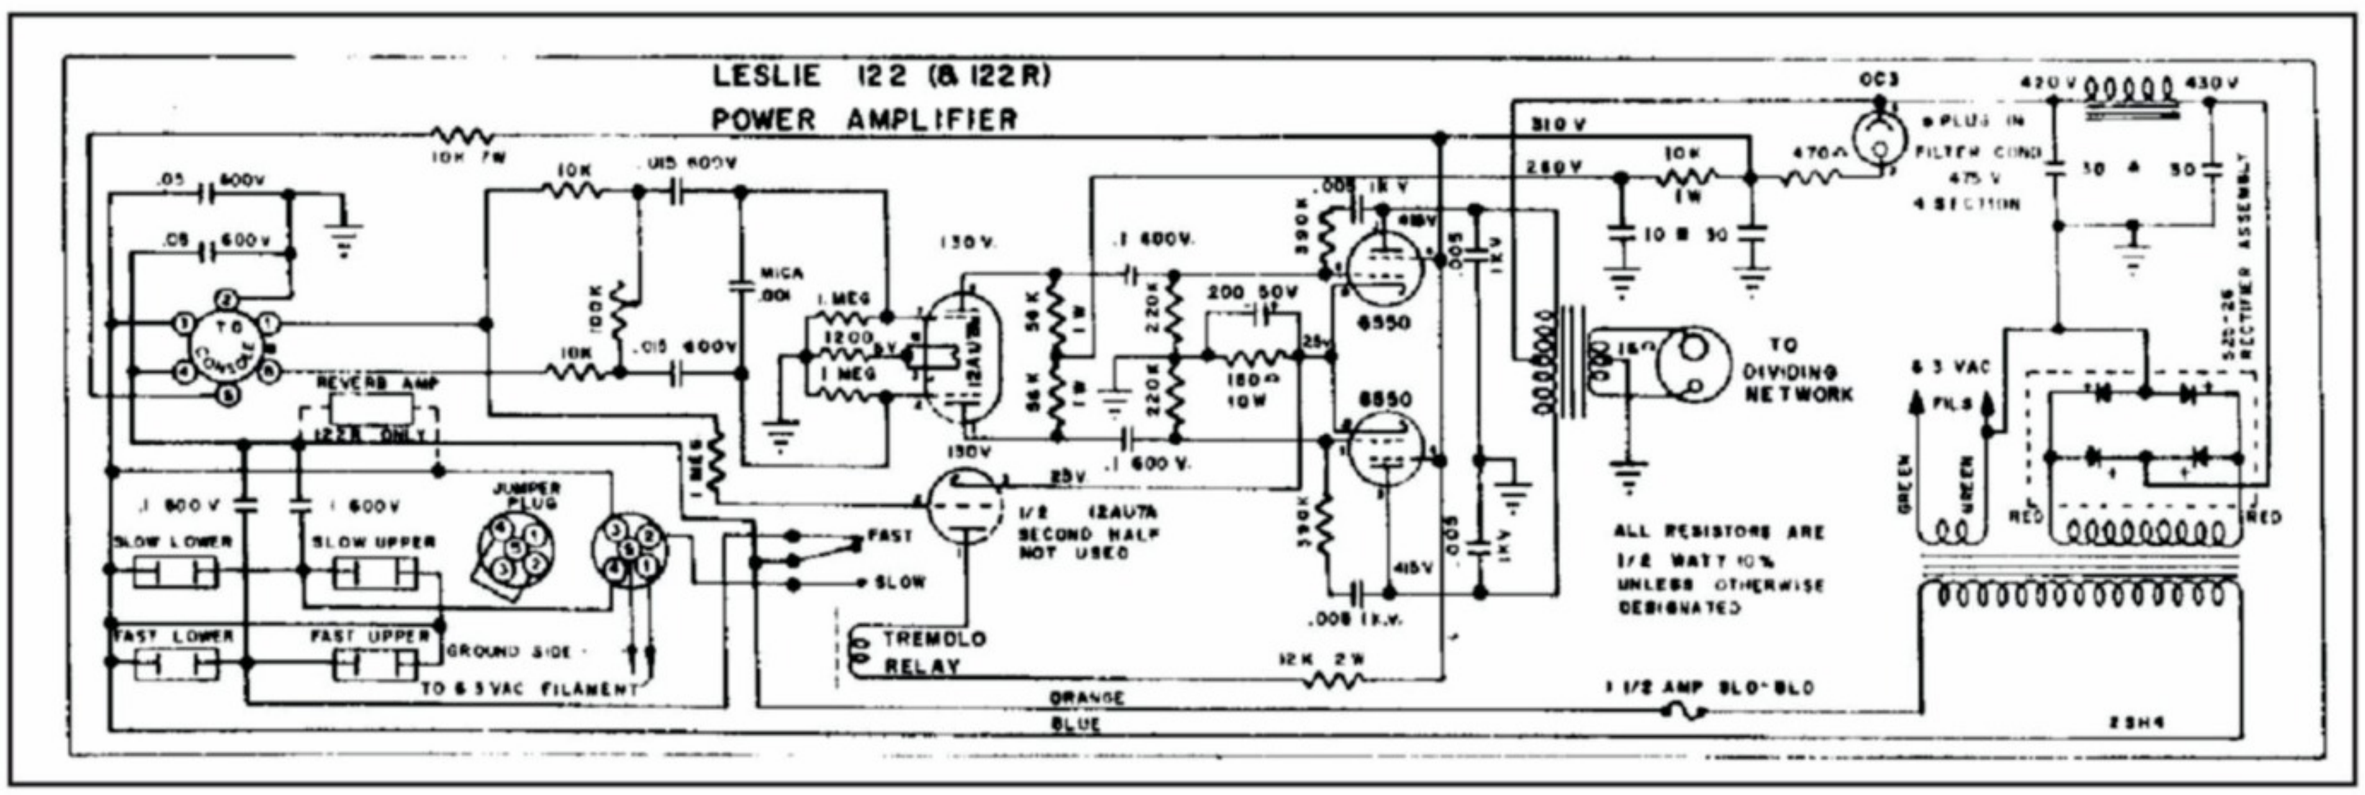

Figure 1: An original Leslie 122 Power Amp schematic

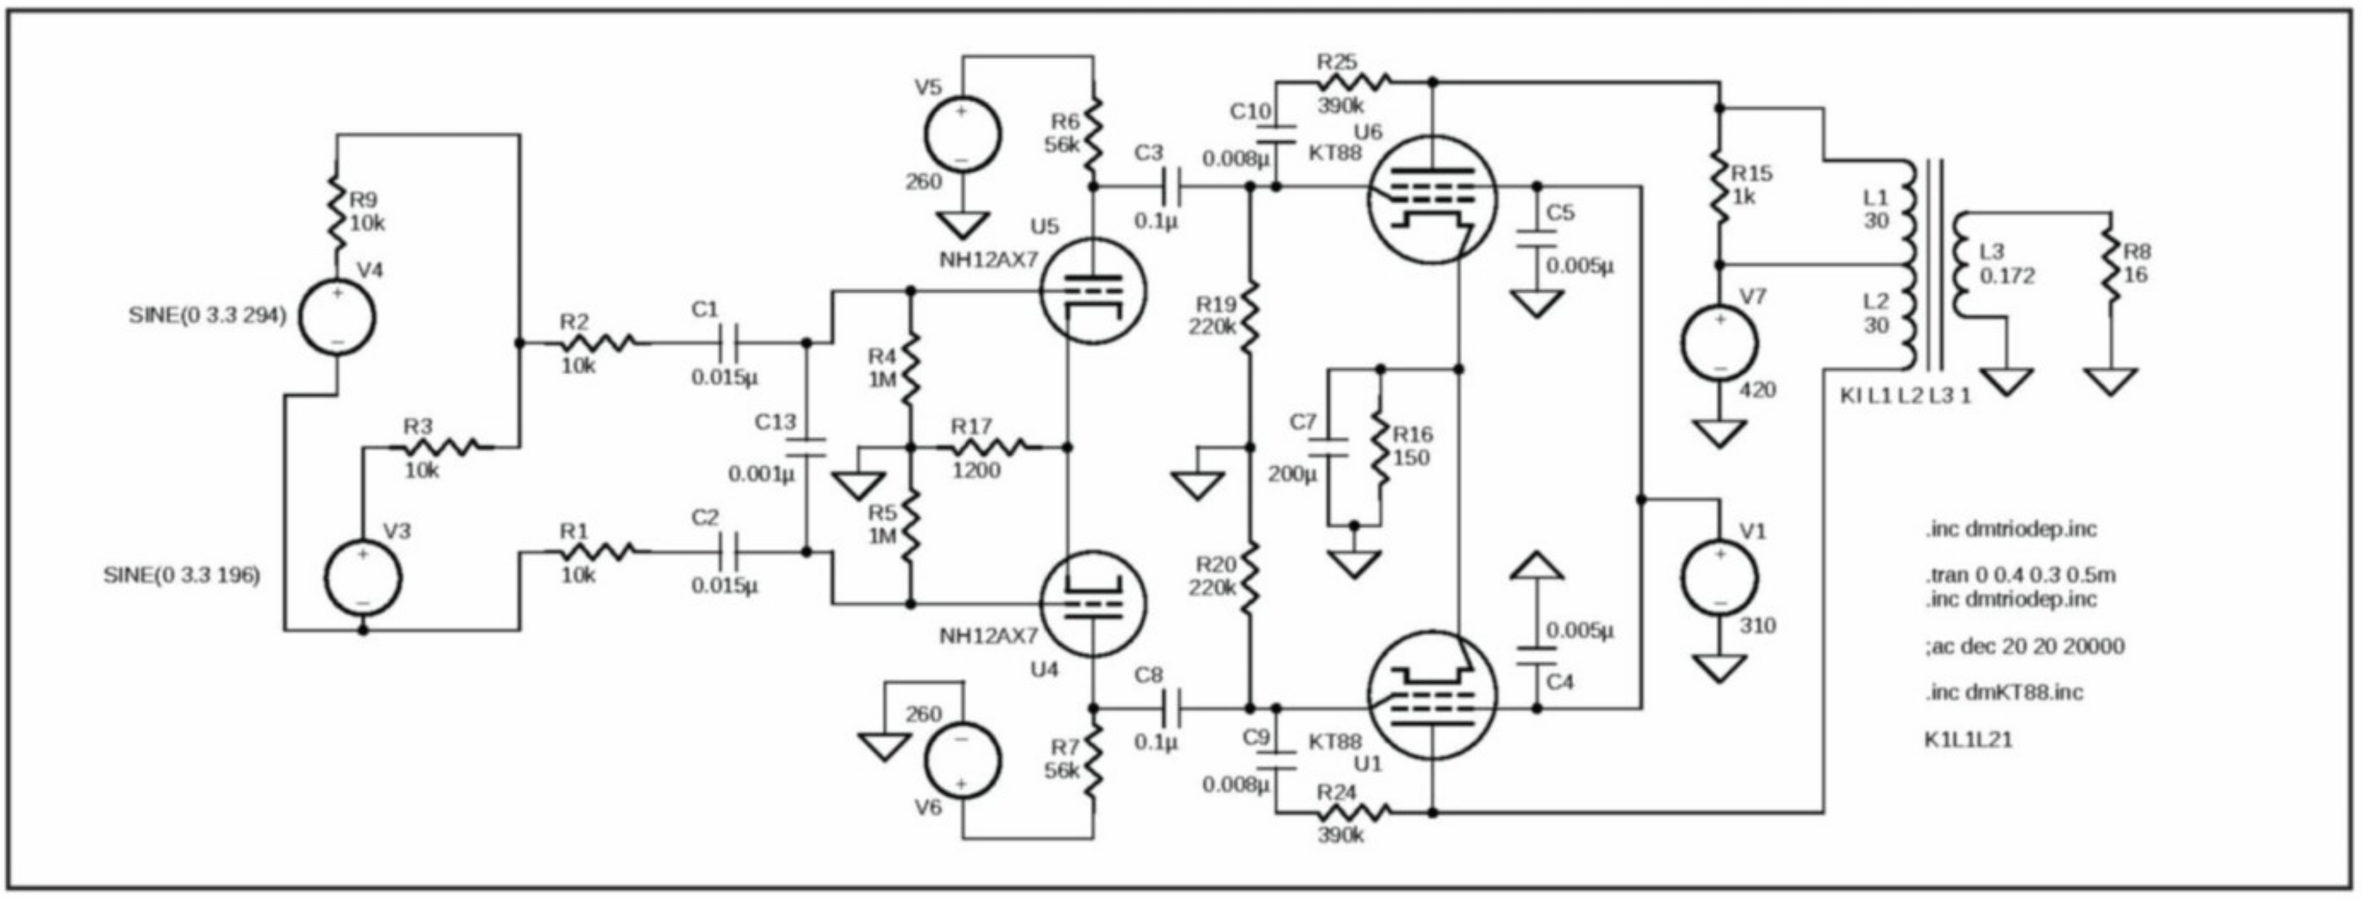

Figure 2: An LTspice model of the Leslie 122 power amplifier

mechanical tone generator and a Leslie speaker with the original amplifier and suitable speaker drivers.

The Growl's Characteristics

All this brings us to the central question for this article: What is involved in the growl? First, let's discuss a few characteristics of the Hammond organ and the Leslie speaker—specifically the Leslie 122 speaker and its power amplifier. The Hammond organ produces pitches up to just below 6 kHz. The specific pitches produced when a given key is pressed depend upon the drawbar settings, as these add subharmonics and harmonics to the fundamental pitch.

The Leslie 122 speaker consists of a 15" woofer passively crossed over to a compression driver (originally a Jensen V21) mounted to a pair of horns (see **Photo 1**) that can be made to rotate, creating a rhythmical Doppler-shift vibrato effect. The Jensen driver design was inspired by the Western Electric 555 driver, but the Jensen used a phenolic diaphragm and had a high-frequency response limit of about 6 kHz, perfectly matching the Hammond's needs.

The Leslie's 15" woofer fed into a rotating baffle that provided vibrato similar to that of the high-frequency section. Never having measured a Leslie speaker, I cannot attest to its low-frequency limit, but since the Hammond can produce output as low as 32 Hz, one would expect that the 122's low-frequency limit is close to that frequency. The low-frequency enclosure has a small vent where the horizontal baffle does not completely seal to the sides and back, so the exact shape of the low-frequency curve depends upon the speaker model used: the original Leslie or an aftermarket Electro-Voice, Altec-Lansing, JBL Eminence, etc.

Figure 1 shows the original schematic of a Leslie 122's power amplifier. **Figure 2** shows an LTspice model of the Leslie 122's power amplifier. Not having exact specifications for the output transformer, I approximated; I knew the model would not be precise. Nevertheless, I could investigate the frequency response and distortion of the amp. **Figure 3** shows the frequency response, which is flat within 3 dB to just under 6 kHz. Adding the amp's 3 dB drop in output at 6 kHz to that of the horn, the system output is down about 6 dB at 6 kHz.

As you may know, amplifier distortion can be split into two kinds: harmonic distortion (HD) and intermodulation distortion (IMD).

- Harmonic distortion results from the creation of multiples (harmonics) of the input frequencies. The aural effect is a change in timbre: more shrill, blanketed, etc., depending upon how much of which harmonics are added.
- Intermodulation distortion (IMD) involves the creation of new frequencies at the sum and difference frequencies of the input tones. Thus, if we fed a 60-Hz tone and a 1-kHz tone

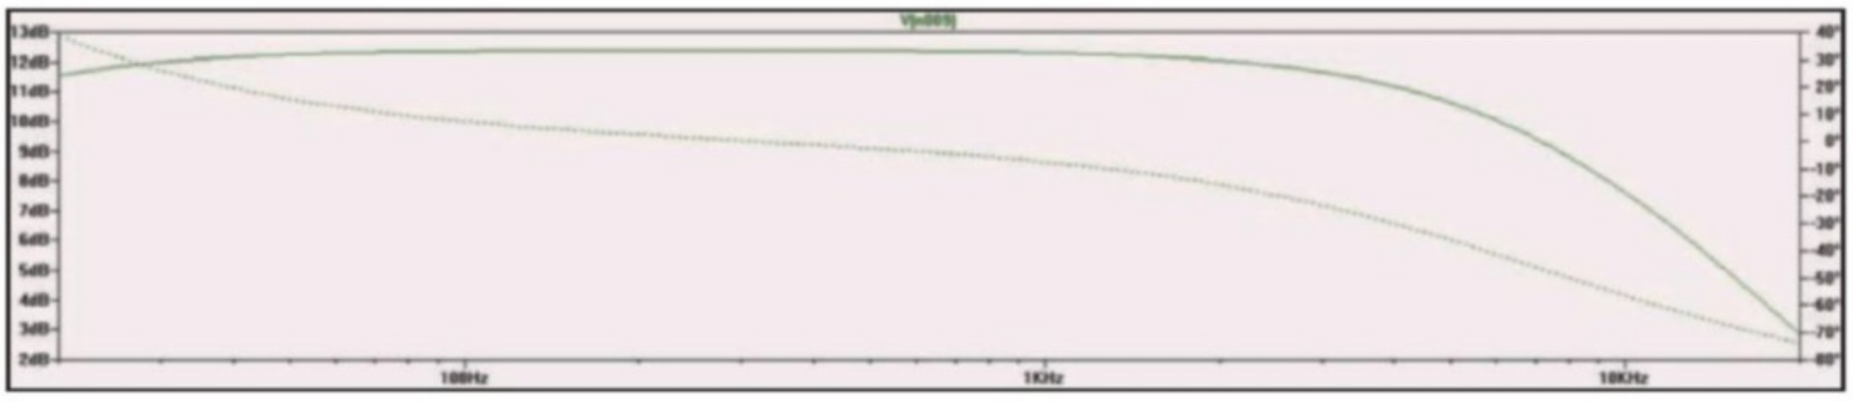

Figure 3: Frequency response of the Leslie 122.

to an amplifier producing IM distortion, the output will contain components at 60 Hz, 940 Hz (difference), 1,000 Hz, and 1,060 Hz (sum). IMD can produce a warbling sound, or in our present case, a growl. The exact nature of IMD's sound depends upon the frequencies involved, the % IMD, and the % HD. (Where there's IMD, there's also HD, and vice-versa.)

The high-frequency limitation of the amp and the speaker help keep the distortion from being obnoxious, since excessive high-frequency distortion products cause a razzy sound that many people find unpleasant.

In order to investigate the distortion produced by a Leslie 122 amplifier under realistic conditions, I fed the modeled amp two inputs, each 2.33 V_{RMS}, and looked at the simulated output wave in the time and frequency domains. The input voltage was chosen to drive the amp into slight clipping. I wanted to compare the model's output to a recording of Steve Winwood's classic intro riff in "Gimme Some Lovin," by the Spencer Davis group. Winwood's first chord in that riff was a G7, including the notes F3 (174.6 Hz), G3 (196 Hz), D4 (293.66 Hz), and F4 (349.23 Hz). These frequencies may not be spot-on, since the Hammond is not an equal-tempered instrument, and I've given the frequencies that would be produced by a perfectly tuned piano. However, these frequencies are close enough. Examining the output at the sum and the difference frequencies of these four notes would be unnecessarily difficult, so I chose only G3 and D4 for the input frequencies in my model.

Using equal voltages for the two input frequencies represents another approximation. Winwood used the first four drawbars, which comprise the subharmonics

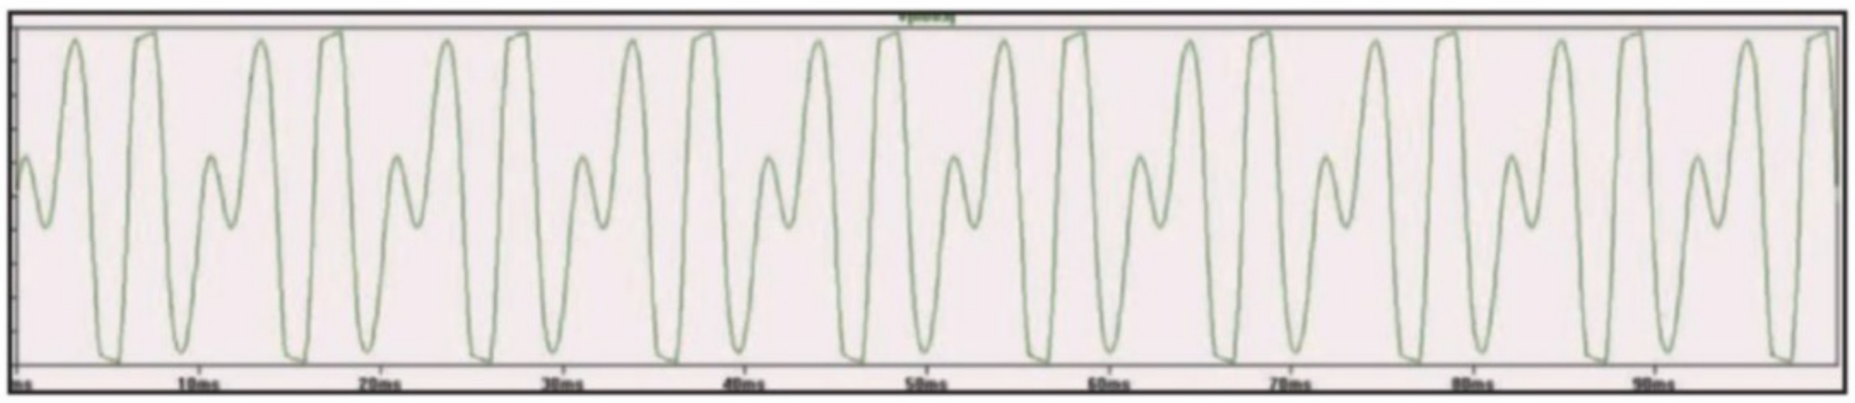

Figure 4: Output wave of the simulated amplifier

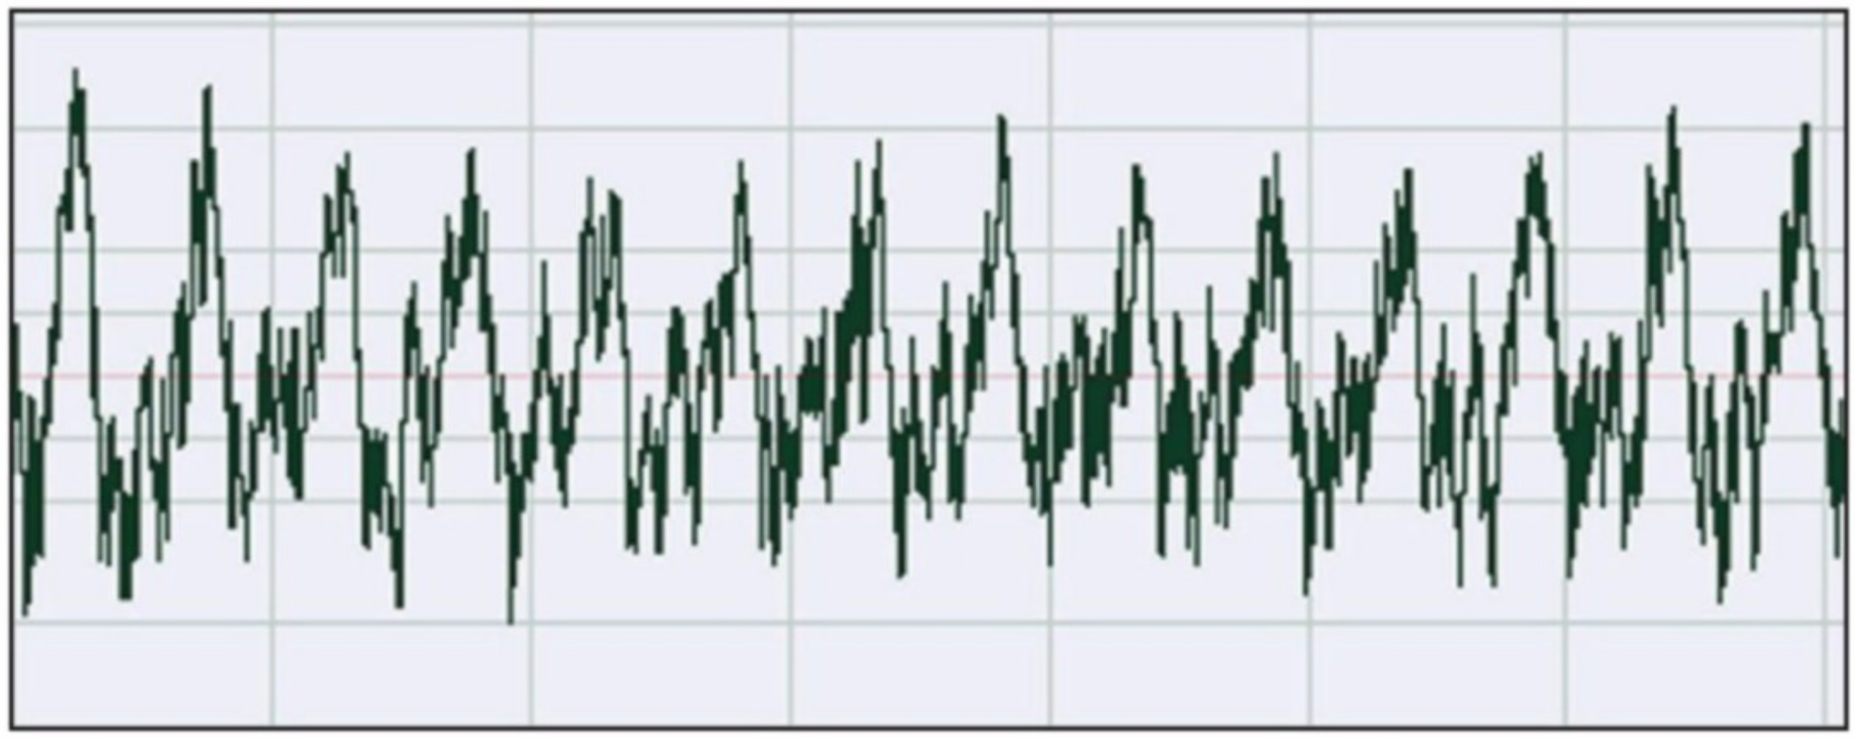

Figure 5: Recorded waveform

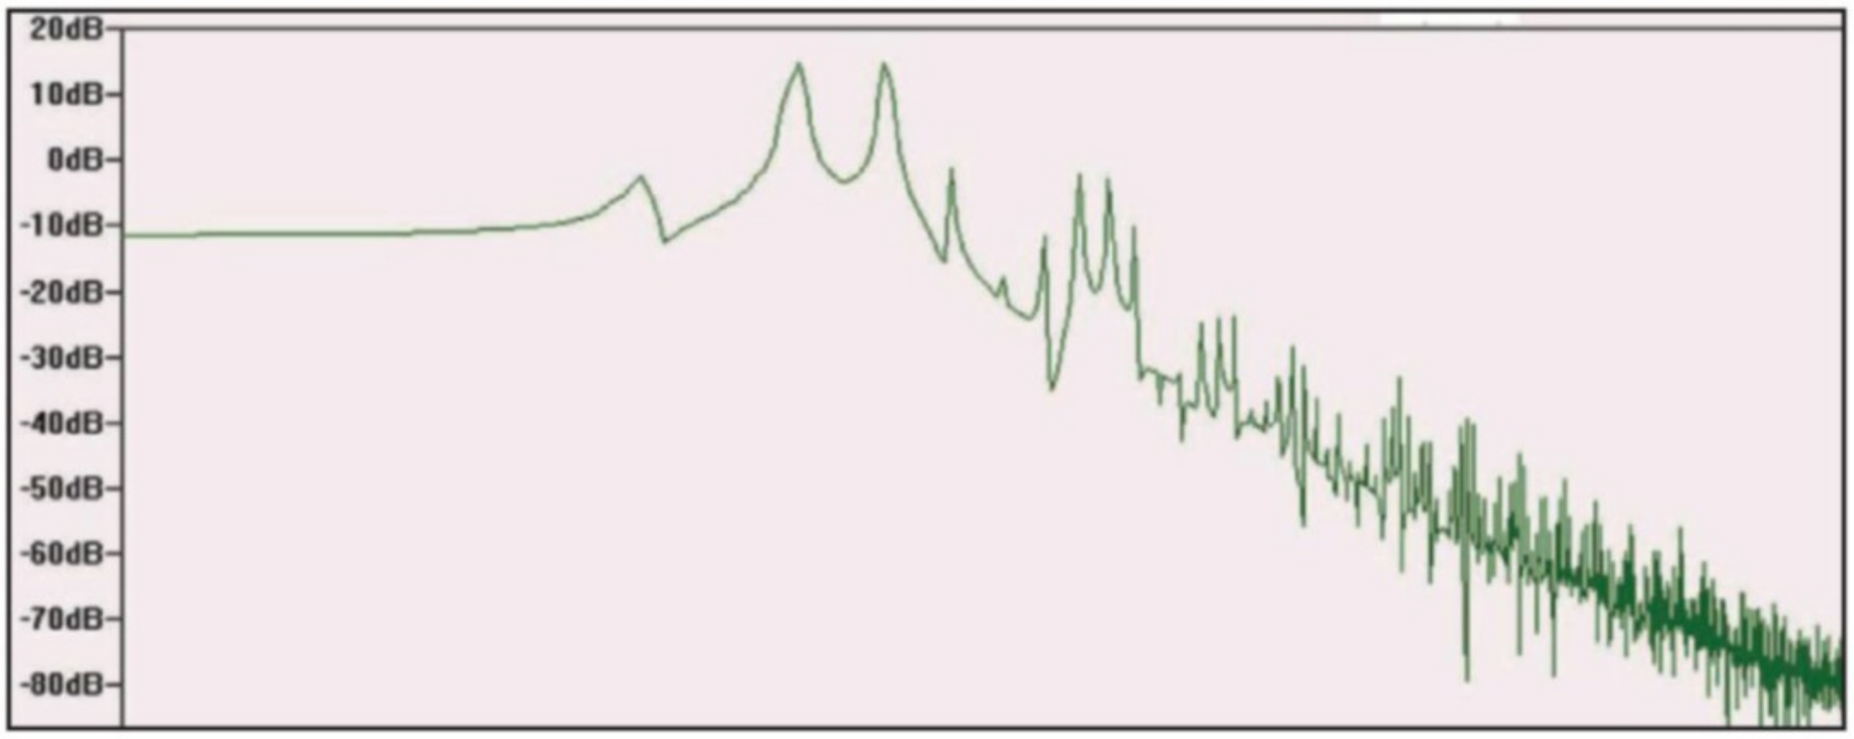

Figure 6: FFT of the modeled waveform

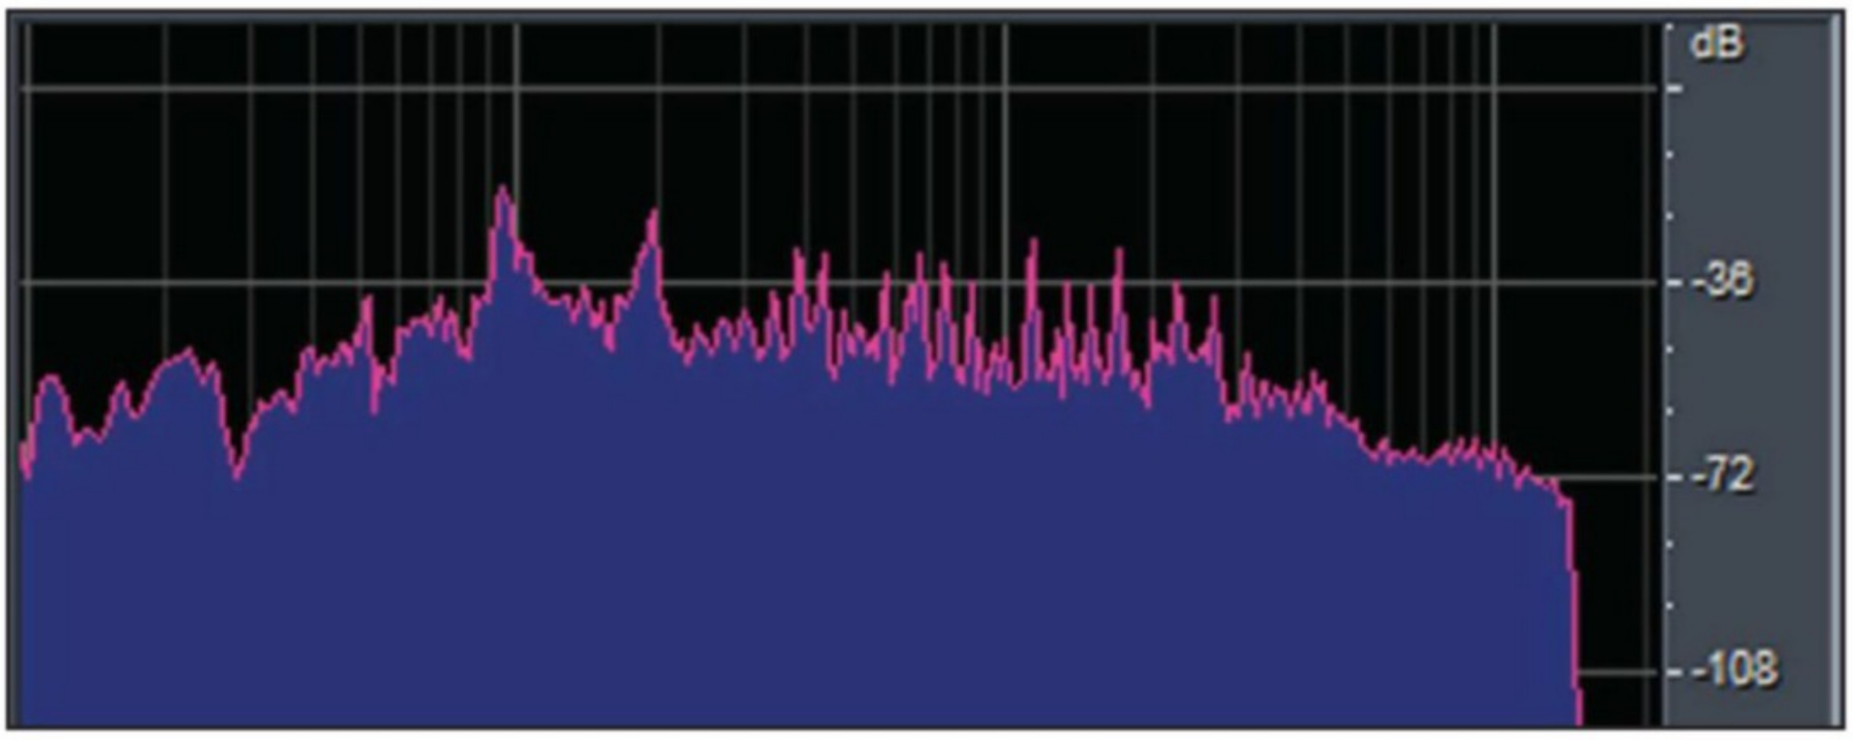

Figure 7: FFT analysis of the recorded waveform

(F2, G2, D3, and F3), the fifth above the subharmonics, the fundamentals, and the second harmonics. The fourth drawbar was not set at maximum, though the lower three were.

Figure 4 shows the output wave produced by the modeled amp. For comparison, **Figure 5** shows the waveform from a recording of "Gimme Some Lovin." Note the difference in the waveform resulting from the greater number of frequency components due to the inclusion of 16 pitches (four keys producing four pitches each). Also, the rest of the band was playing, producing additional frequency components, although the organ was dominant.

Figure 6 and **Figure 7** show the Fast Fourier Transforms (FFTs) of the modeled and recorded waveforms, respectively. Both FFTs have strong components at the difference frequency of 97.66 Hz, but neither shows a visible sum-frequency component at 489.66 Hz. The LTspice model's output shows just over 10% first-order IMD, and 11% second-harmonic distortion (2HD) of G3, plus 5% 2HD of D4.

The FFT of the recorded wave shows higher IMD and HD, but it is difficult to make a meaningful comparison, since the organ was feeding so many more frequencies to the amp. [Note the high peak at 99.55 Hz, whose amplitude exceeds that of the G3 and D4. While this component is at the difference frequency between G3 and D4, its higher amplitude (4 dB higher than the amplitude of G3) is difficult to explain unless the bass guitar was playing that same note.] It is also possible that Winwood was driving his Leslie amp further into clipping than I did in the LTSpice model.

The frequency peaks that appear in both FFTs are shown in **Table 1**. You can see many additional peaks in Figure 7, the FFT of the recorded wave.

Distortion	Frequency (Hz)
Difference	98 Hz
G3	196 Hz
D4	294 Hz
Second harmonic of G3	390.525 Hz
Second harmonic of D4	580.587 Hz
Seventh harmonic of the Difference	688.85 Hz
Fourth harmonic of G3	780.608 Hz
Third harmonic of D4	876.12 Hz

Table 1: Frequency peaks in the FFTs

Unleashing the Growl

Over the years, I have done work for many keyboardists, including replacing Leslie low- and high-frequency drivers with units having higher power ratings, replacing Leslie 122 amplifiers with other amps including Ampeg V4 and SVT heads, and comparing different output tubes with the stock 6550s in the 122 amp. Some of these modifications were simply repairs, and some were attempts to make the organ louder. In no case have I ever found any modification that improved upon the stock Hammond B3/Leslie 122 growl.

Speakers—especially direct radiating speakers—typically produce more HD than IMD, and the distortion percentage increases at lower frequencies. The FFTs show significant distortion well into the fundamental voice range, not mainly at or below 100 Hz, as would be the case if the speakers played a major role in creating the growl.

I am now satisfied that the Hammond growl is a result of overdriving the bandlimited 122 amplifier, producing lots of IMD. Most of the studies of distortion audibility have shown that 1% HD may or may not be noticeable, but 1% IMD is usually objectionable. So the approximately equal percentages of IMD and 2HD produced by a Leslie amp would result in more noticeable IMD effect than HD effect: more growl than timbre modification.

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.

December 2-6, 2019
178th Meeting of the Acoustical Society of America
The Hotel del Coronado, San Diego, CA
acousticalsociety.org

The 178th Meeting of the Acoustical Society of America (ASA) will be held during the first week of December 2019 in San Diego, CA. The Society holds two meetings each year—one in the US and the second one abroad. At each meeting, invited and contributed papers and poster presentations, usually number between 850 and 1,100, reflect the diversity and interests of the membership. Tutorials, lectures, technical and standards committee meetings, and visits to university, government, and industrial laboratories are regular features, along with receptions and a variety of social programs. These meetings offer opportunities for students and young researchers as well as experienced acousticians to share information. Members who are able to attend meetings regularly find the personal contacts a rewarding experience. 2020 will be the International Year of Sound! IYS2020 will be a global initiative to highlight the importance of sound in all aspects of life on earth and will work toward an understanding of sound related issues at the international level. https://sound2020.org.

***

January 2020

January 7-10, 2020
CES Show 2020
Las Vegas Convention and World Trade Center (LVCC) and 10 other locations in Las Vegas, NV
www.ces.tech

CES is where business gets done. That's a valid way to describe the largest consumer electronics trade show, and how it continues to be an essential meeting point for the global technology. Far from the days of a partially consumer-oriented giant exhibition, as the Consumer Technology Association (CTA) evolved into a wider reaching standards and trade organization, the CES show also dramatically evolved into a mainly B2B affair and showcase for channel partners and the media. More than half of the show's visitors spend more time in meetings outside the traditional show floor, which in turn expanded with the evolution of completely new industries and market segments, such as automotive and mobility, 5G networks, and the smart home technology.

Now more than ever, CES 2020 is expected to be the world's gathering place for all who thrive on the business of consumer technologies. It is the global stage where next-generation innovations are introduced to the marketplace. As the largest hands-on event of its kind, CES features all aspects of the industry and the 2020 edition will be a pivotal moment to check out the latest trends on applied technologies in areas such as wireless audio, personal audio, HDMI, automotive sound, wireless charging, voice interfaces, and the smart home.

Now in its 53rd year, the world's leading innovation and most influential technology event on the planet and promoted every year by the CTA, continues to be the launch pad for breakthrough technology and innovation. While the ongoing trade disputes and some market uncertainties will no doubt result in a reduced presence from Chinese companies, the Las Vegas event also promises to be the ideal setting where many new vendors and brands will look to capitalize in a thriving and stable market environment, such as the United States. 2020 promises to be once again a record year in terms of attendance, with more and larger companies committed to exhibit or simply conduct business in their own meeting spaces—spread all over the city.

Last year, the CTA audit concluded that CES 2019 attracted 175,212 in total attendance—35% of which came from outside of the United States. This number compares to 182,279 in CES 2018 and the high of

184,279 at CES 2017. The CTA says this is the result of its deliberate strategy of "managing attendee numbers and attracting the most highly qualified attendees. This strategy is aimed at helping limit hotel price increases and the pressure on travel to and around Las Vegas during CES. While CES attracted government leaders from around the world, many members of the US government were not able to attend due to the government shutdown." Still, the impression we had from previous years is that professionals attending CES and not even registering to enter the exhibition halls is also on the rise, with meeting and conferences areas in practically all hotels in Las Vegas being fully booked.

The CTA also reported that CES 2019 continued to grow in exhibit space, with more than 4,500 attending companies expanding to an incredible 2.9 million net square feet in 2019.

With 11 official venues, CES features more than 25 marketplaces and product categories. Venues are grouped into three geographical areas: Tech East, Tech West, and Tech South.

Product Categories and Marketplaces of interest for *audioXpress* readers include "Audio and Video," "High-Tech Retailing," "Home Cinema," "IoT Infrastructure," "iProducts," "Smart Home," "Wearables," and "Wireless Devices and Services." But don't forget to check out the Design & Source Showcase and the Eureka Park.

At Tech East —located in the Las Vegas Convention and World Trade Center (LVCC), the Westgate Las Vegas (Westgate), and the Renaissance Las Vegas (Renaissance)—there will be innovations in audio, HDMI, automotive electronics, gaming, video, wireless devices, wireless services, digital imaging/photography and all "i" devices.

The Tech West exhibits—at the Sands Expo (Sands), The Venetian, The Palazzo, Wynn Las Vegas and Encore at Wynn (Wynn/Encore)—will feature emerging technologies, encompassing innovations in fitness and health, the Internet of Things (IoT), smart home and other high-growth technologies changing the way we live, work, and play. The Audio Marketplace is included in Tech West at the Venetian.

Finally, Tech South—located at the ARIA, Cosmopolitan, and Vdara hotels—will host a variety of C Space activities, including conference sessions, networking events, exhibits, and meeting rooms geared toward the advertising and marketing community at CES.

CES remains a trade-only event and is not open to the general public, which is one of the reasons why the traditional Venetian Audio Marketplace cannot (and should not) be considered a traditional audio show, increasingly replaced by B2B activities of the most important manufacturers, meeting with suppliers, technology vendors, and global distribution partners.

The "Must Have" reference for loudspeaker engineering professionals.

Home, Car, or Home Theater!

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

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

A $99 value!

Yours today for just $39.95.

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

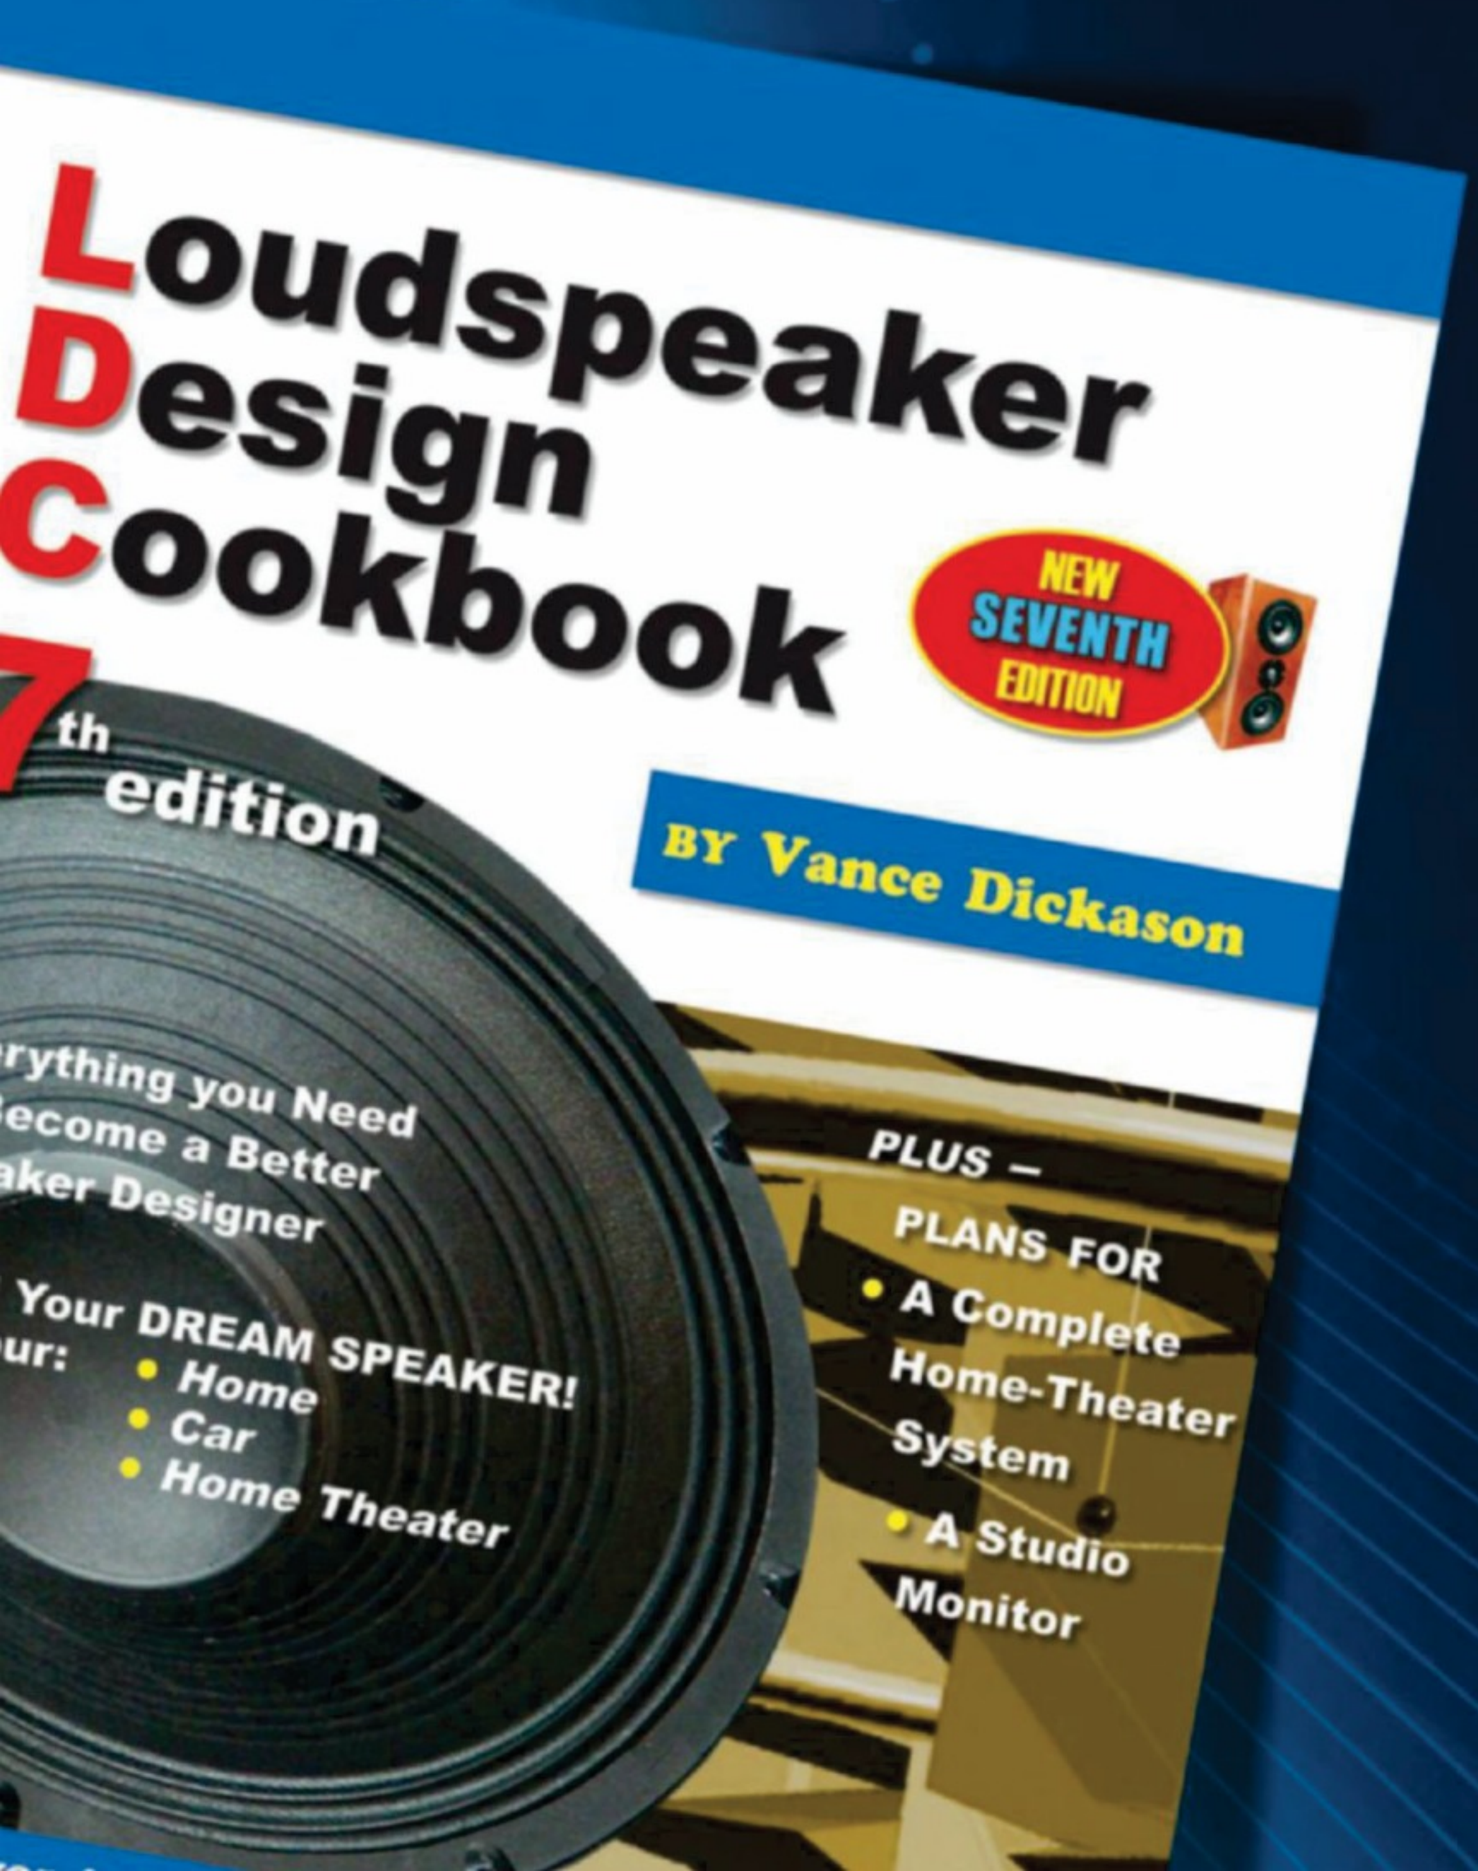

Accurately Measure Speaker Parameters ***Faster Than Ever!***

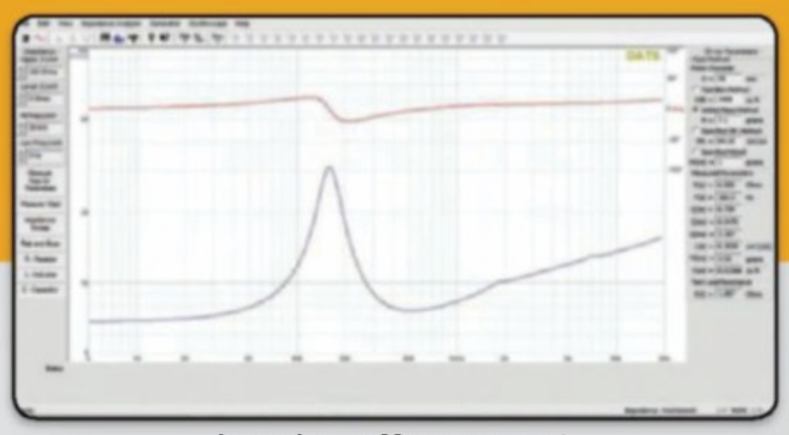

Impedance Measurement

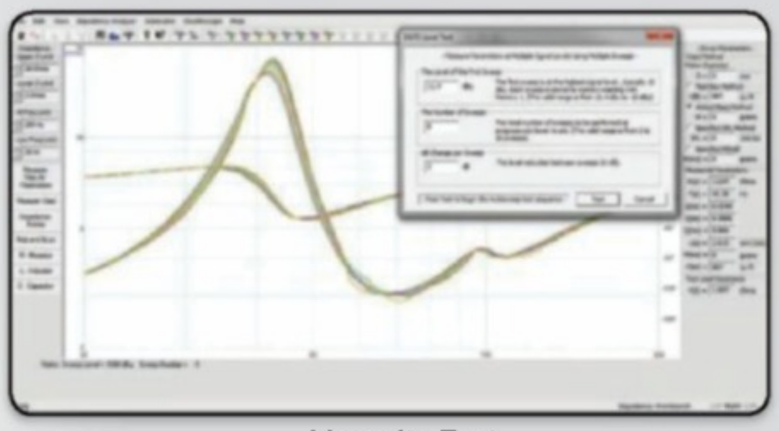

Linearity Test

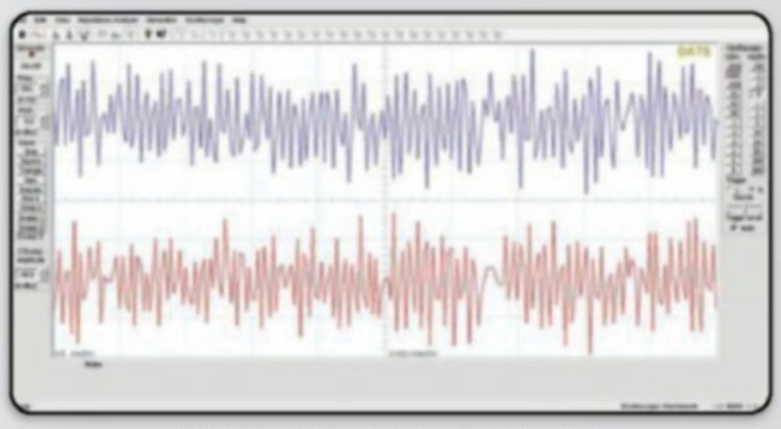

Oscilloscope Wave Generator

DATS V3

DAYTON AUDIO TEST SYSTEM

DAYTON AUDIO

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!

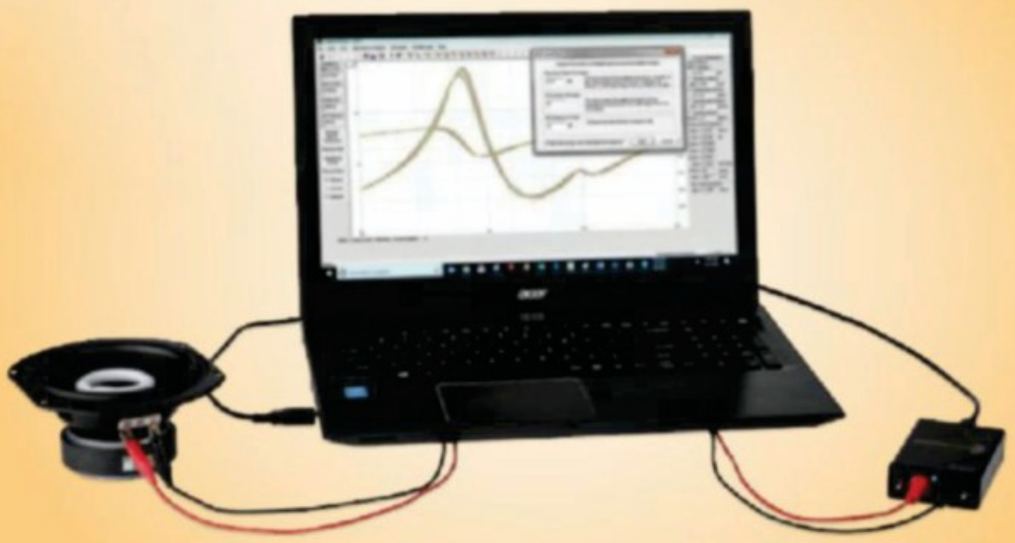

Improvements include:

- 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

AVAILABLE MID-NOVEMBER

PRE-ORDER NOW!

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