INNOVATIONS IN AUDIO • PRODUCT DEVELOPMENT • THE BEST IN DIY AUDIO

www.audioxpress.com

audioxpress

ADVANCING THE EVOLUTION OF AUDIO TECHNOLOGY

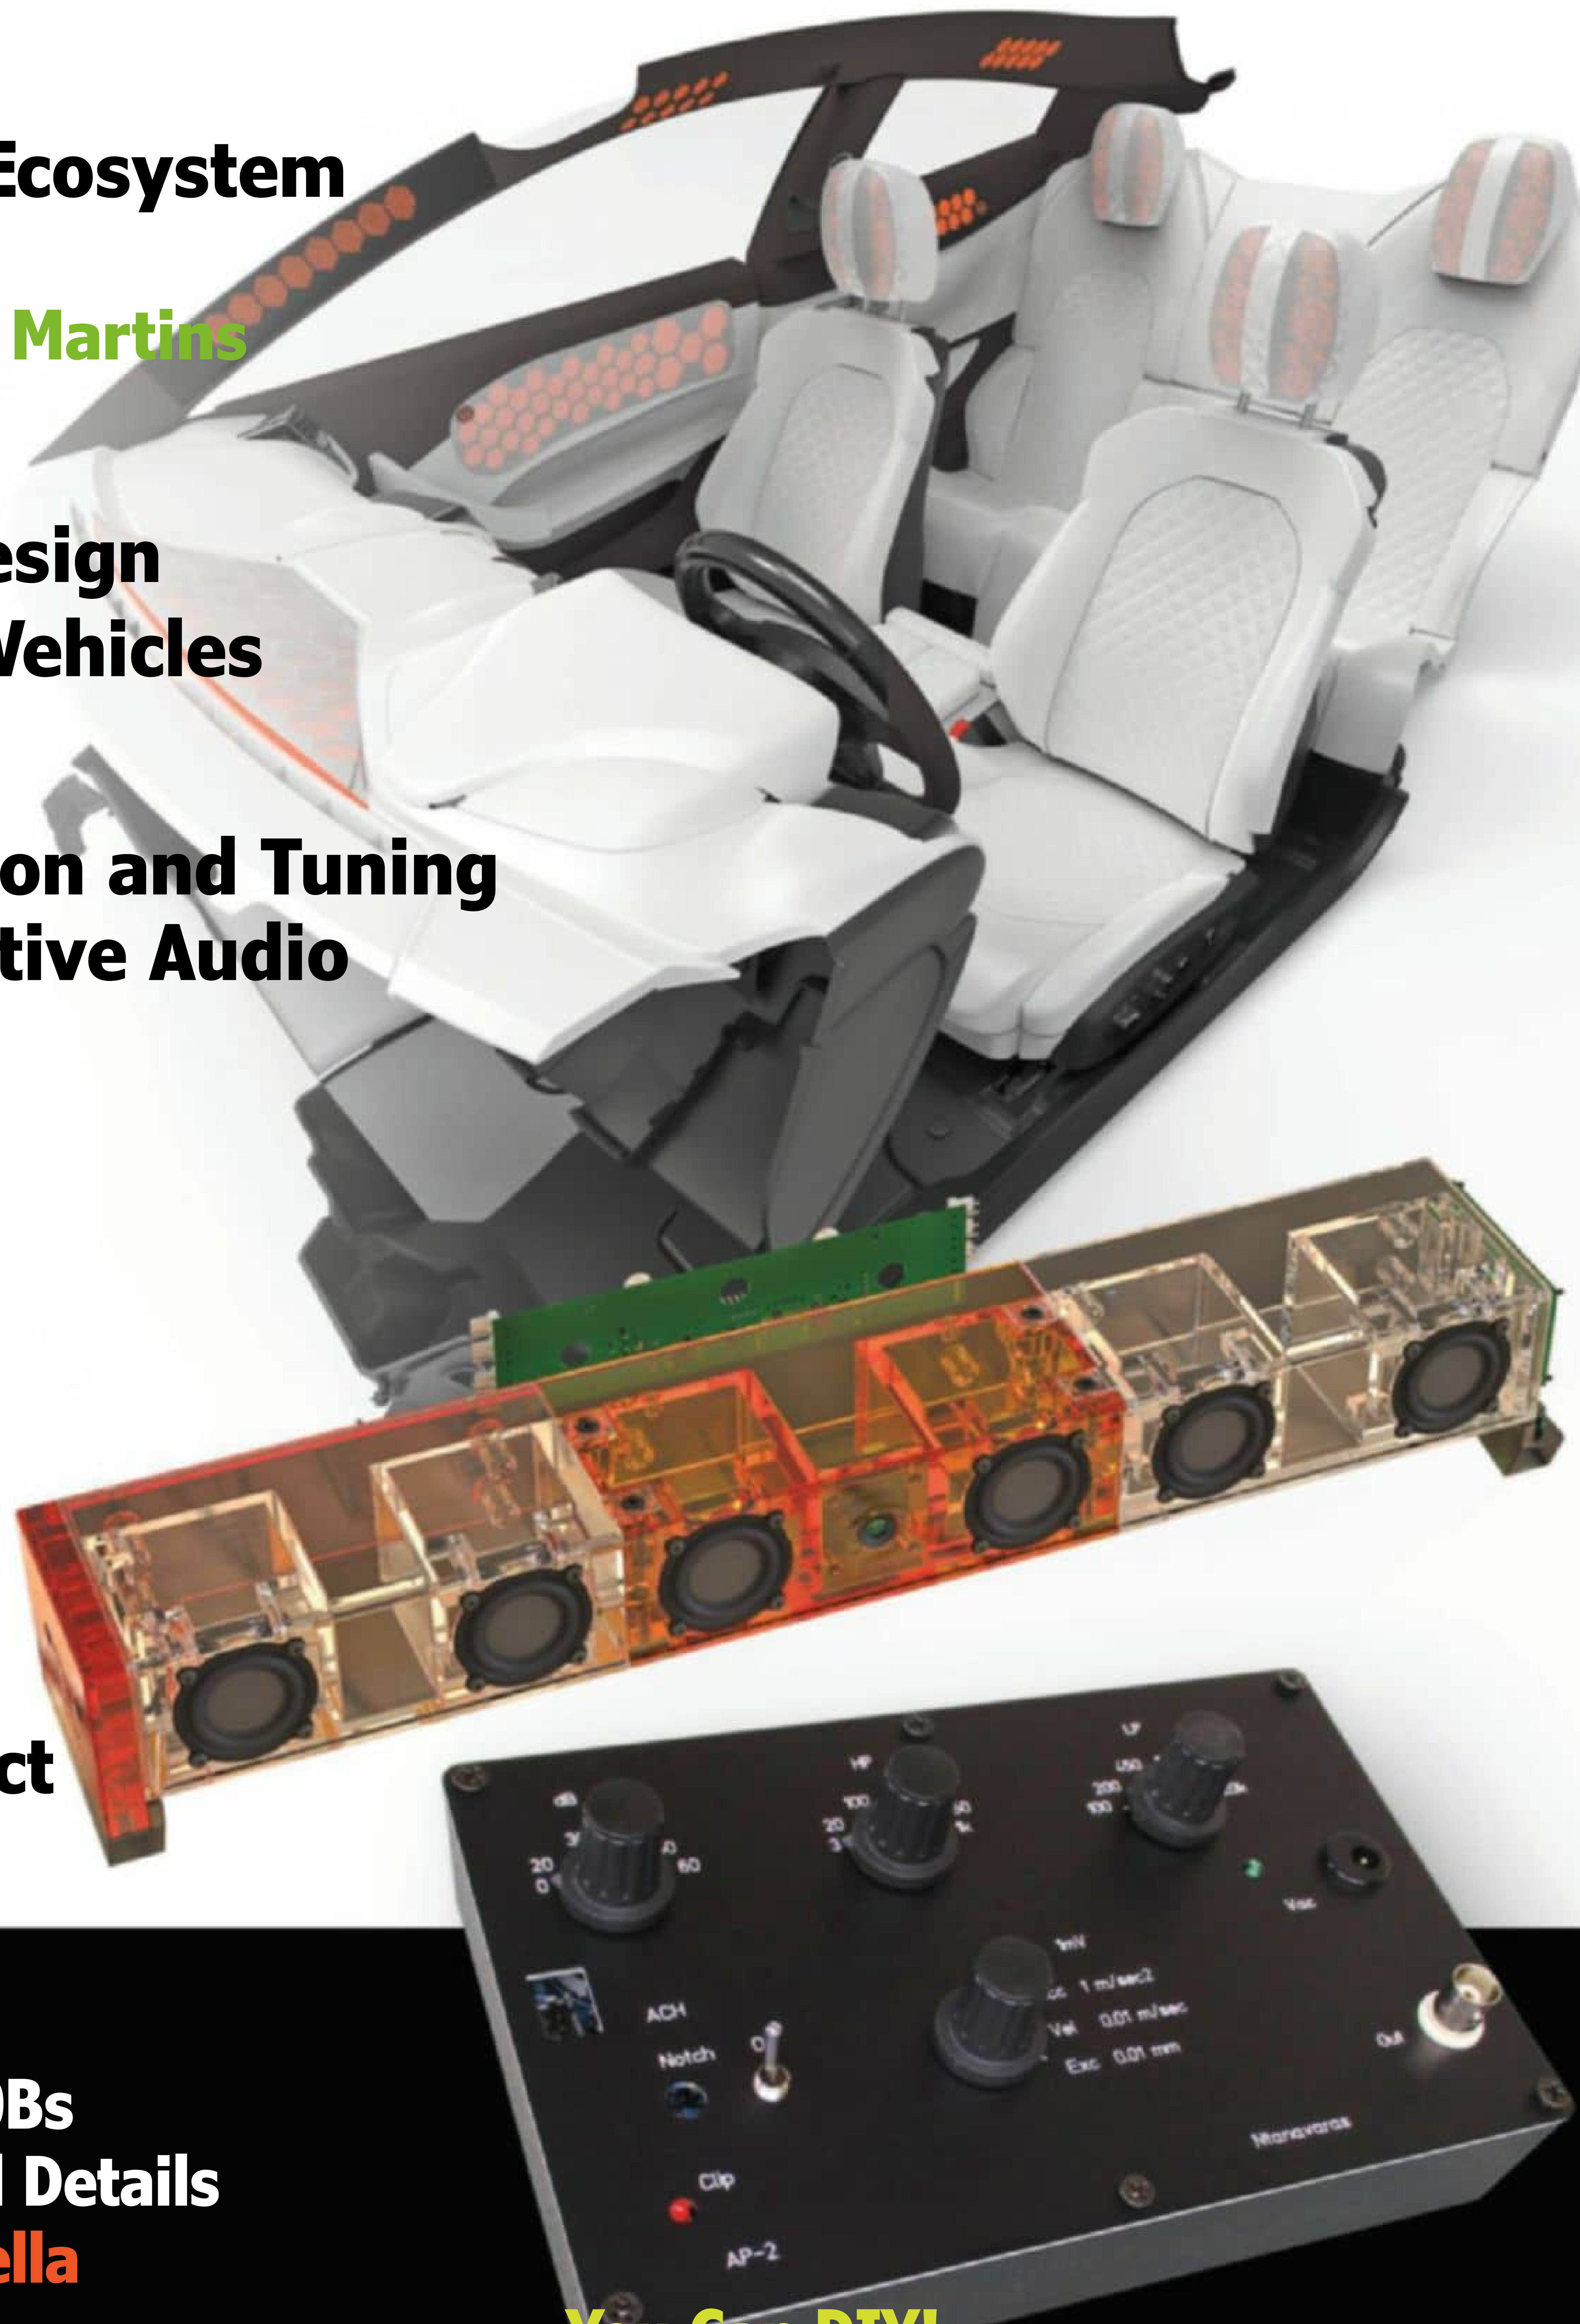

Automotive Audio

The Changing Audio Ecosystem in Electric Vehicles

By Roger Shively and J. Martins

Perspectives on Automotive Sound Design in the Era of Electric Vehicles

By Adam Levenson

DSP Module Application and Tuning Procedure in Automotive Audio

By Philipp Paul Klose

Simulation Techniques

Misconceptions in the Audio Industry

By René Christensen

R&D Stories

Audioscenic: From Vision to Product

By Marcos Simón

You Can DIY!

A (Z)OTL with 300Bs Part 2 — The Final Details

By Thomas Perazella

Sound Control

An Exploration of the Book *Rock and Pop Venues*

By Richard Honeycutt

You Can DIY!

The AP-2 Amplifier With Analog Signal Processor for the ACH-01 Accelerometer

By George Ntanavaras

$9.00US $10.00CAN

06

7 25274 84697 6

JUNE 2023

June 2023 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, suggestions, and contributed articles to:
audioXpress, Editorial Department
P.O. Box 417
Chase City, VA 23924, US
E-mail: editor@audioxpress.com

The Sound of Silence

Automotive audio is a fascinating topic that *audioXpress* has been increasingly exploring. We have received an enthusiastic response from both our readership and the industry, but our strongest motivation has been the unprecedented level of activity and innovation taking place. Looking back 10 years, when we envisaged the new format and the roadmap for this publication, I confess that this was not something I anticipated. Yes, Tesla started selling its Model S in 2012 and we should have foreseen how fully electric vehicles would impact the audio experience as the familiar noise of the combustion engine was suddenly removed.

Back then, car audio was a separate world and not something with which we envisaged getting involved on a deep level—in particular because there were plenty of industry publications focused on the segment. I believe many audio industry professionals felt the same, until they had the chance to explore sound systems in the new electric vehicle (EV) environment and it became obvious that nothing would stay the same. A completely new approach was needed for audio systems, starting with the need to define the sound of the car itself—for drivers and passengers, and for the outside world.

Everything started to change with the first hybrids—I remember riding in a Toyota Prius and admiring the initial silence when the car was set in motion, until the road noise became more prevalent than we were accustomed to having in other vehicles—because there is no such thing as silence in a moving vehicle. It's just a different type of noise that quickly becomes extremely loud when the traditional engine noise doesn't mask it. I don't remember anyone bragging about the possibility of solving that problem by installing a Burmester sound system for their Prius.

My first experience with ultra-high-end audio in a car was precisely with a Burmester system in the backseat of a Mercedes-Benz that drove me to the airport. I was in absolute awe of the experience, the comfort, the unusual silence inside the cabin, and the magistral sound of the music playing—making it almost impossible to detect the direction from which it was coming. In a way, a truly immersive experience even though the source was in plain stereo, showing me that automotive sound could be a completely different discipline from what we would consider a reference speaker listening environment, in the studio or at home.

Likewise, John Whitecar, Principal Engineer at Tesla from 2015 to 2017, who is interviewed in this issue of *audioXpress* together with his colleague Steve Ernst, now both at DSP Concepts, was working on fairly standard automotive entertainment systems during that period. Basically, because the automotive audio systems that were still being designed in 2012 and a few years after that, were basically entertainment audio systems "fitted" to the car. Where radio receivers and antennas, together with the need to support Bluetooth and connected smartphones, were priorities for audio "infotainment" teams. Such systems were not designed with the car and to enhance the car, like what we see being envisioned today.

And that's what has set in motion a new era for audio innovation, allowing for many new players to step into this completely new environment. As the saying goes, "necessity is the mother of invention" and EVs, being inherently silent, needed completely new audio technologies. As an example, new speaker topologies that have never received much attention for conventional applications and were simply considered to be a "different" way to achieve the same thing, are now an essential way to meet the two fundamental requirements of reduced weight and lower power, mandatory in new vehicles. Noise reduction or cancellation systems were always more effective when combined with active acoustics, and efforts in that direction are showing up everywhere, because EVs need it. And yes, immersive audio systems, which could be considered an optional luxury add-on, are fundamentally the right approach for sound systems in these new vehicles.

Because sound needs to be a part of the car driving experience, it doesn't hurt that audio systems also offer a great entertainment experience when listening to music or any other content. And that's even before we get to the point where that entertainment becomes basically the center of the whole autonomous car experience.

J. Martins
Editor-in-Chief

OUR NETWORK

audioxpress | circuit cellar | LIS Loudspeaker Industry Sourcebook | Voice Coil

SUPPORTING COMPANIES

Contents

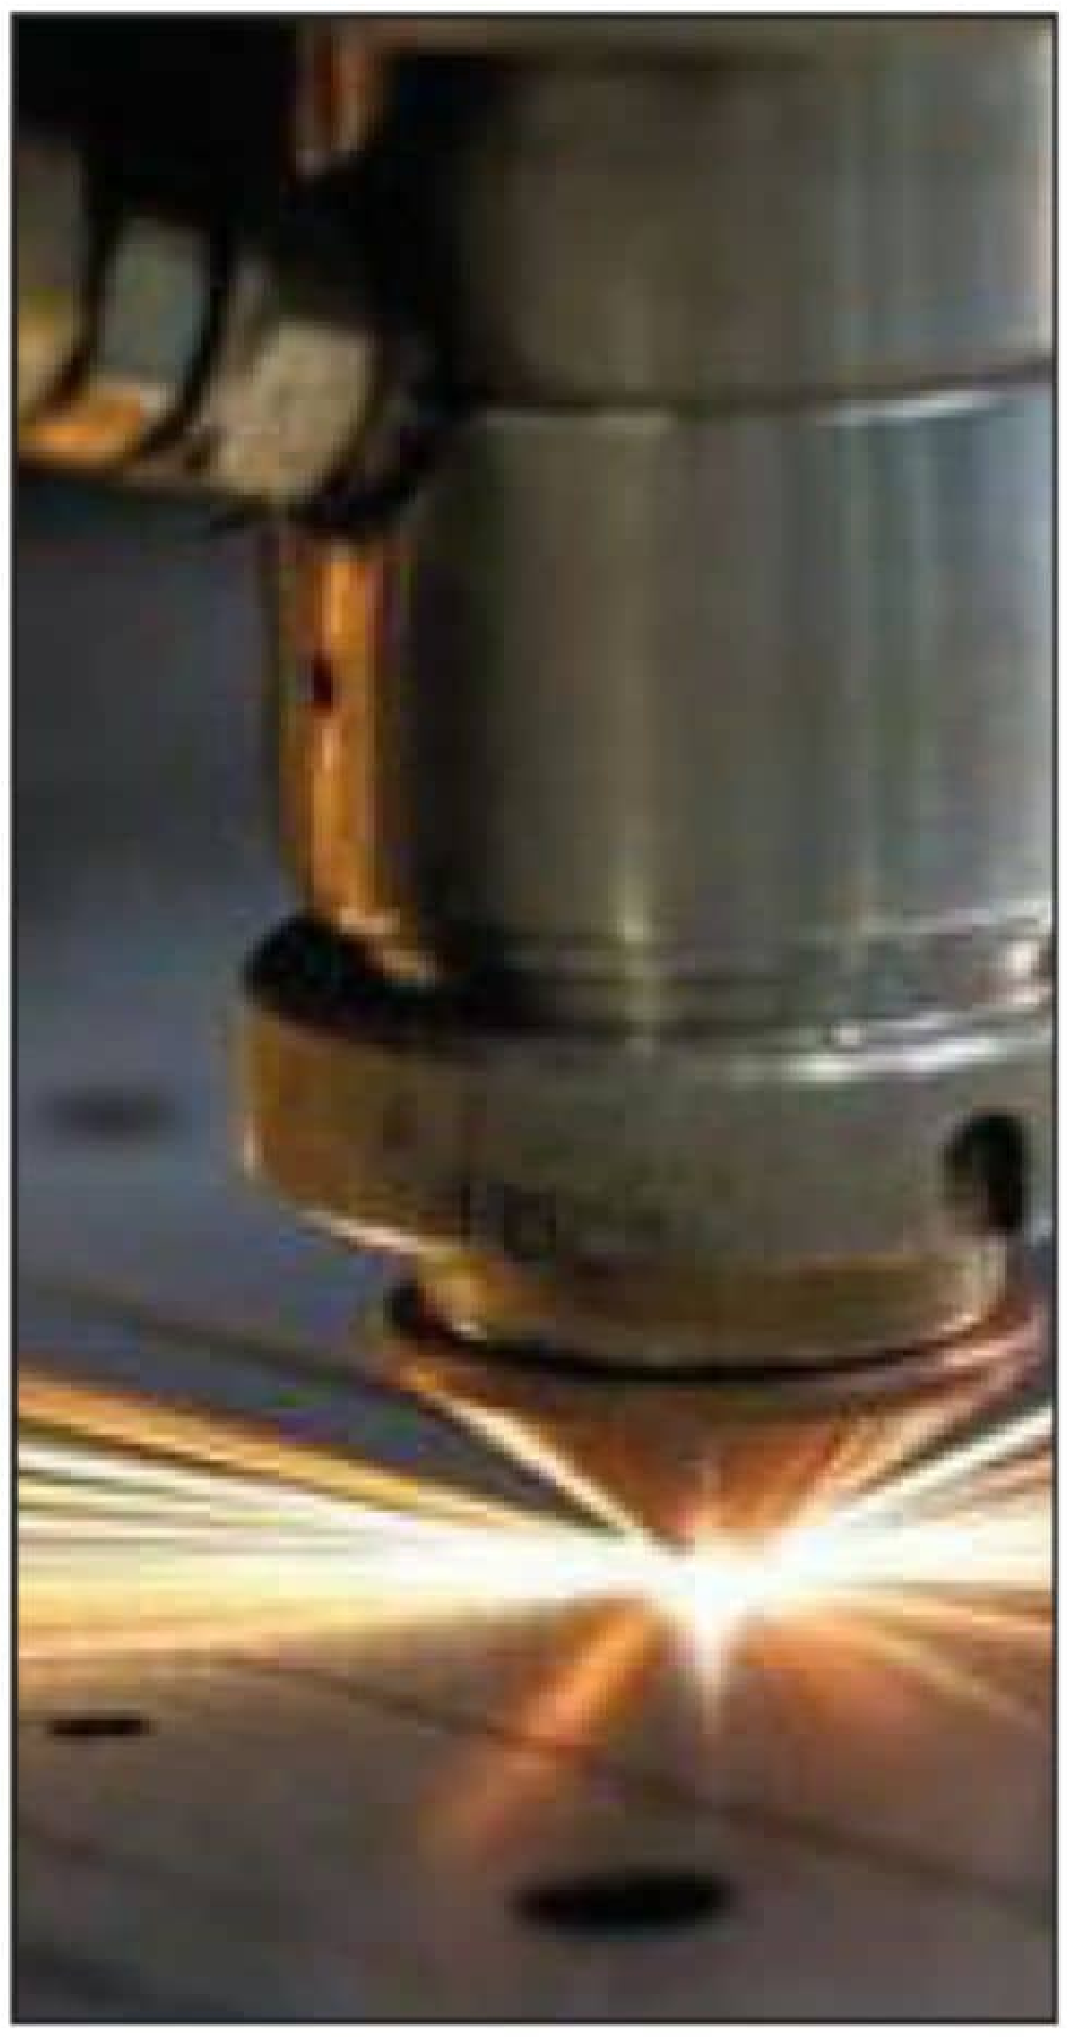

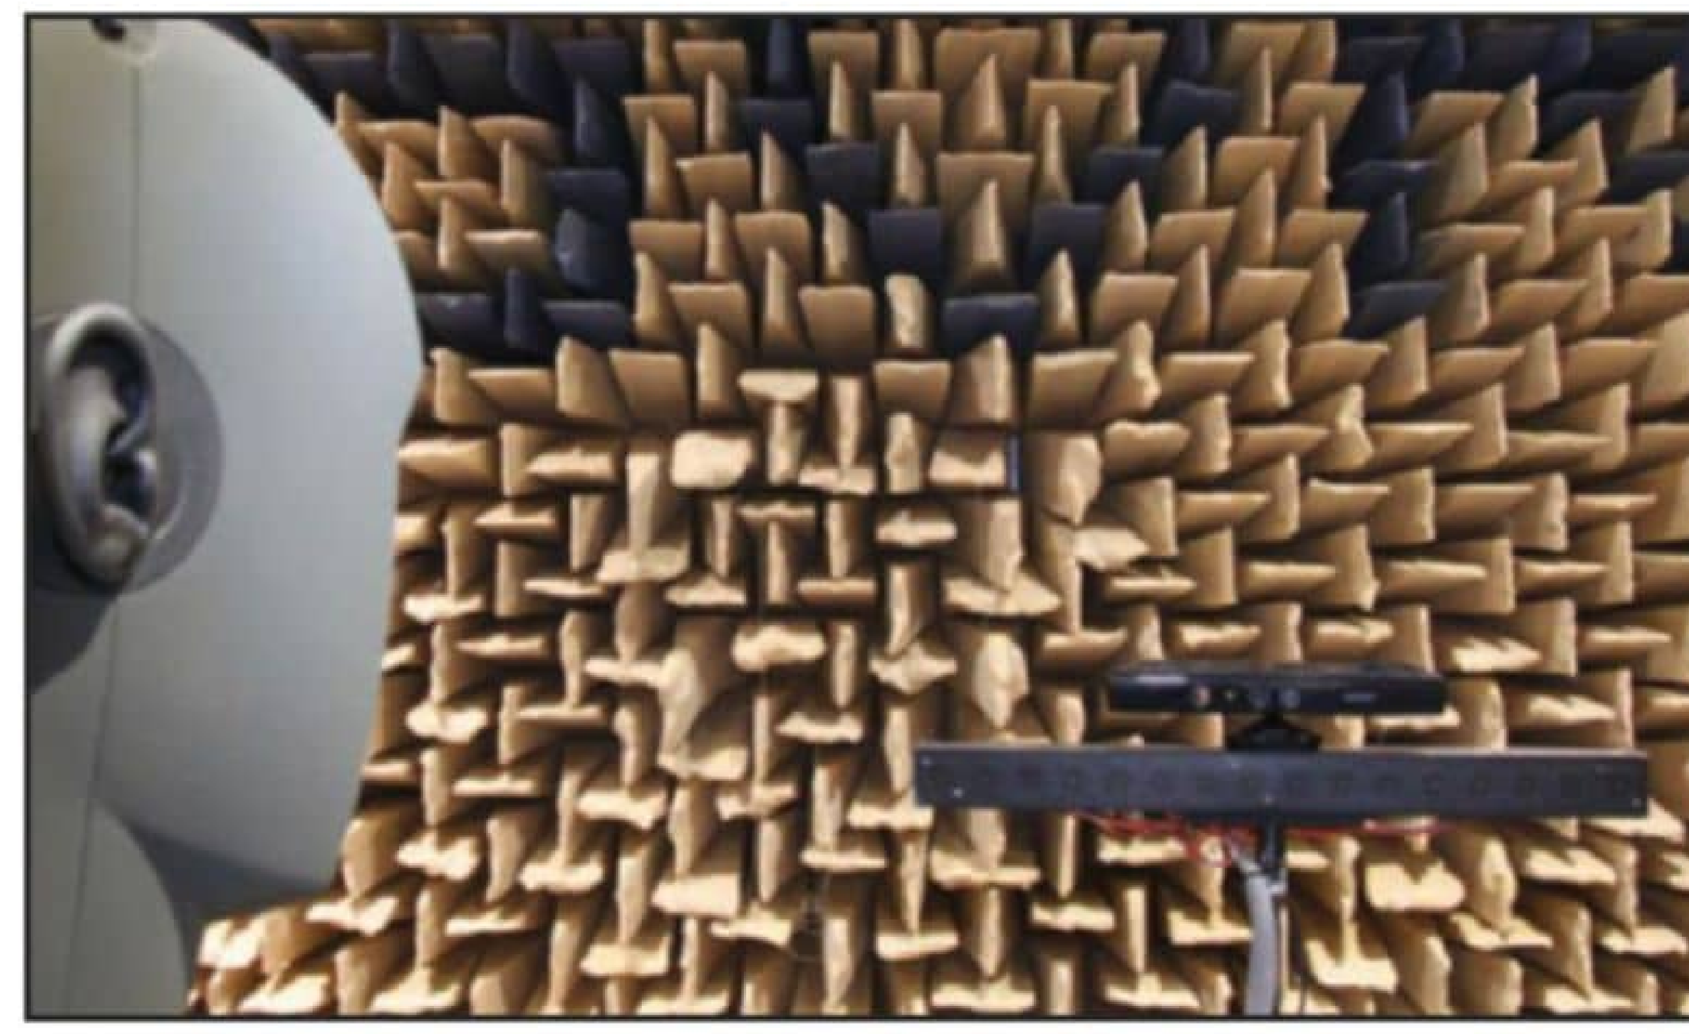

Features

COVER ART

Warwick Acoustics automotive electroacoustic panel technology

****

Audioscenic and Tectonic soundbar design

Volume 54 – No. 6

June 2023

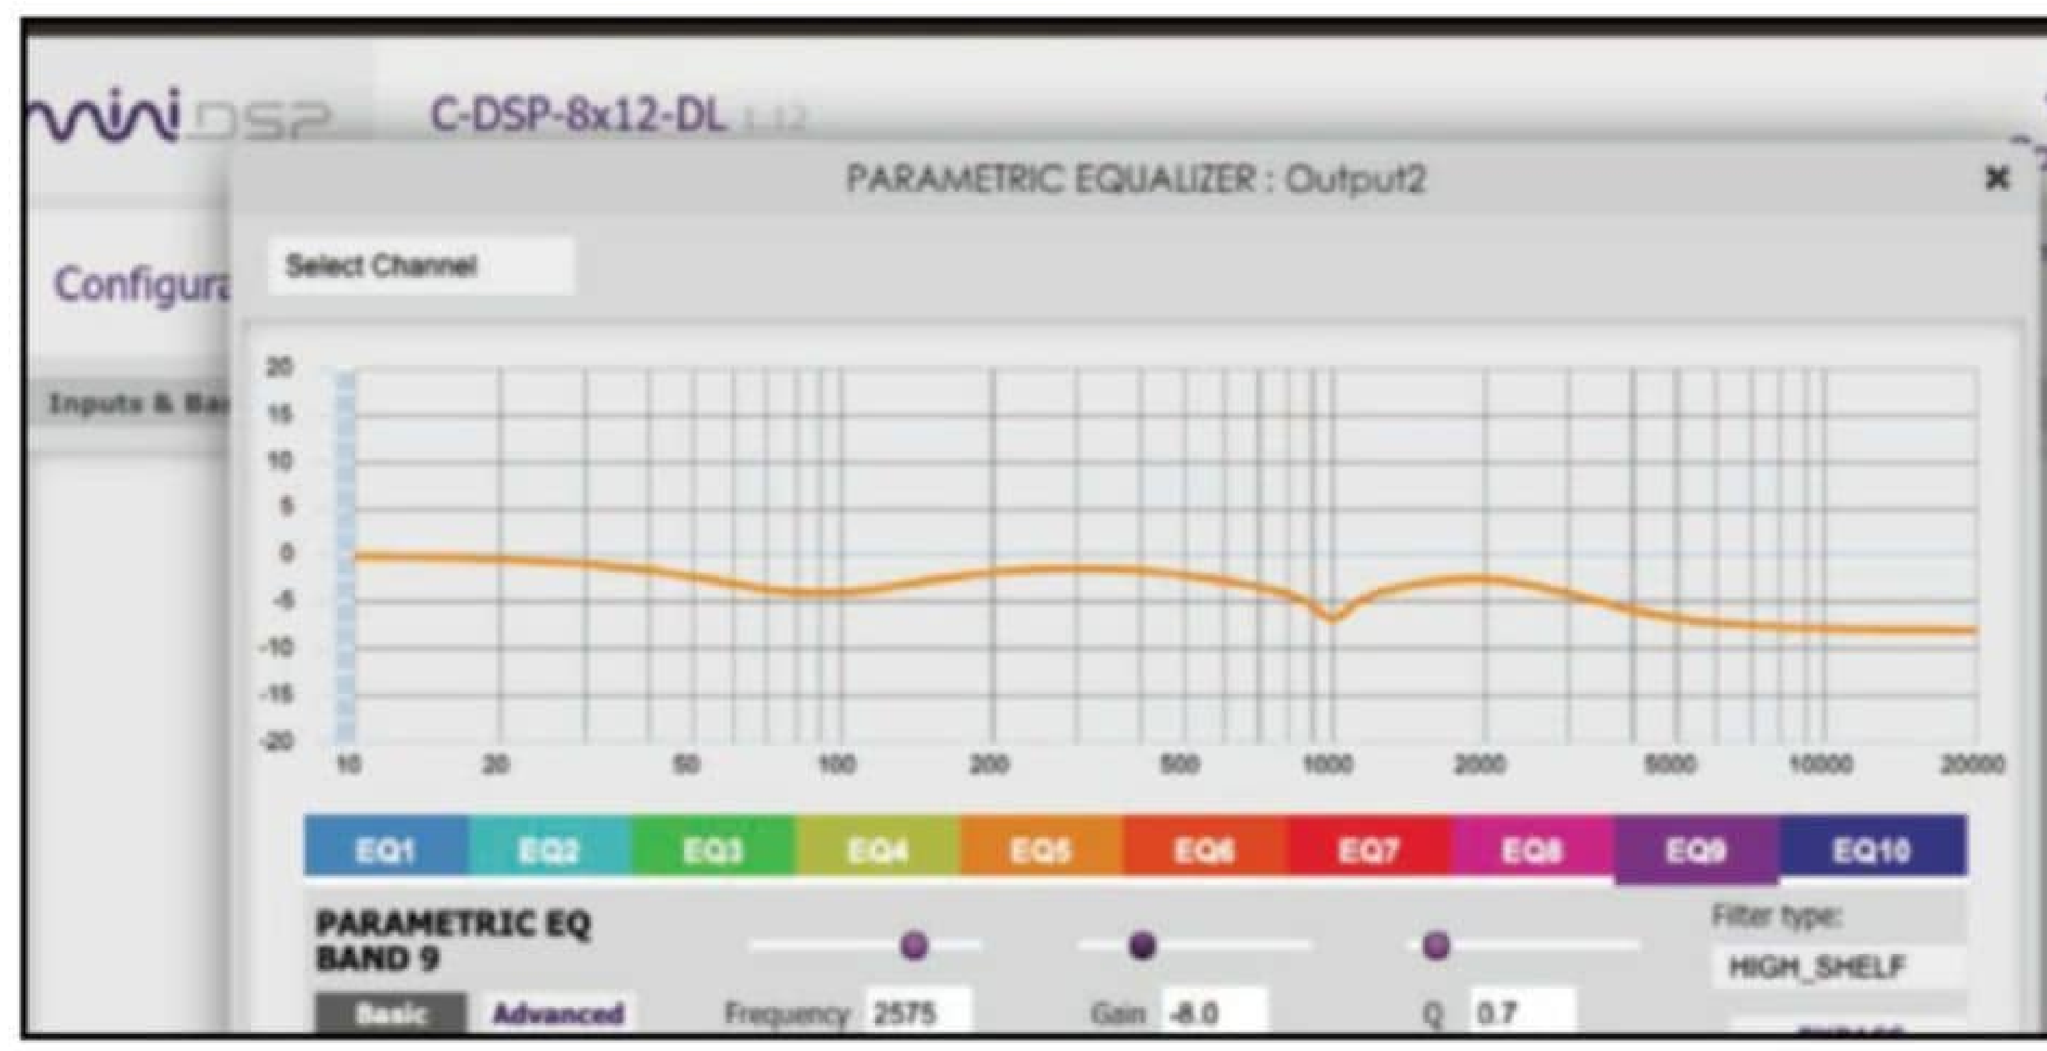

Automotive Audio

Columns

Departments

Websites

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

@audioxp_editor audioxpresscommunity
linkedin.com/company/audioxpress
audioxpress

The Changing Audio Ecosystem in Electric Vehicles

By
Roger Shively
and J. Martins

DTS AutoStage is an AI-powered connected media delivery platform that enables a powerfully personalized, in-vehicle infotainment solution.

The rapid shift to plug-in hybrid and full-battery electric vehicles is having a significant and long-lasting impact on how audio systems are integrated into cars. While some noises are much more prominent due to the absence of the familiar engine noise, the transition offers exciting possibilities for sound systems, passenger entertainment, and enhancing the overall automotive experience.

According to a recent report commissioned by DTS, a wholly owned subsidiary of Xperi, 75% of Gen Z drivers (ages 18 to 25), and 49% of drivers overall, view their vehicles as a third space—a place outside of the home and work, a place to relax and escape the daily stresses. Based on results from a national consumer CARAVAN survey conducted by research agency Big Village, the report finds that safety, comfort, and radio/audio rank highest as vehicle design priorities. According to the report, 91% of Millennials and Gen Zers are interested in mood-sensing technology to adjust in-vehicle audio content. And 69% of drivers are even interested in connected applications, including health and diagnostic updates, confirming the automaker's vision that there are numerous possibilities to monetize subscriptions, apps, and services in connected vehicles.

For Xperi, as we can also see from many other large consumer electronics and technology companies, this means that there are immense opportunities that can be tapped in the future of the connected-car industry in new entertainment, personalization and wellness applications. Car audio plays a major role in the expectations consumers have in terms of available features in a car, expectations that are highest when it comes to features they already have available in personal or home audio.

That is why companies such as Xperi are so invested in exploring new opportunities for automotive infotainment in future vehicles. This includes exploring completely new domains, such as what Occupant Monitoring Systems did with their DTS AutoSense solution (introduced in vehicles in 2021), and using machine-learning technology to perceive occupants' actions and intentions, independently of both the cameras and the platform used in the model. Expanding from this approach, Xperi is now promoting DTS AutoStage as a global automotive-independent, AI-powered media delivery platform, combining traditional technologies that have AM/FM radio and HD-Radio with connected streaming of personalized content. The solution is already present in more than 52 connected car models worldwide, deployed by five car brands including Mercedes-Benz, Hyundai, and Kia, and available in more than 60 countries.

Connected platforms such as DTS AutoStage are just one of the examples of the sophisticated infotainment systems that will require extensive integration not only with all audio systems in the car, but also inevitably with active noise control and other systems in electric vehicles (EVs), such as embedded cognition technology and driver assistance, among others.

The EV Transition

In January, following the CES 2023, *The Audio Voice* e-newsletter reported that car buyers bought fewer cars last year than in any year since 2011, but more fully electric vehicles than ever before. In the US, total new vehicle sales fell 8% in 2022, but EV sales grew by a surprising 65% according to a December 2022 report from Kelley Blue Book's parent company Cox Automotive. In the first quarter of 2023, that number has grown to 70% compared to the

Market Update — Automotive Audio

same time last year. Currently, 7.2% of the new cars Americans bought were electric, up from 6.2% in December 2022 and 3.2% in 2021. Furthermore, it is estimated that there will be a 15-fold increase in EV sales by 2030, with sales of possibly 30 million EVs in the next 5 years alone.

According to a report on EVs from the International Energy Agency (IEA) (Paris, France), global sales of new EVs surpassed 10 million units in 2022, a record 14% of all vehicles sold. The 2022 percentage was a sharp uptick over 2021 and 2020 percentages, when only 9% and 5%, respectively, of new cars and trucks sold around the world were powered by electricity. The IEA also reported that there were roughly 26 million EVs on the road around the world at year-end 2022, with China accounting for about 60% of global new EV units sold in 2022. The IEA expects

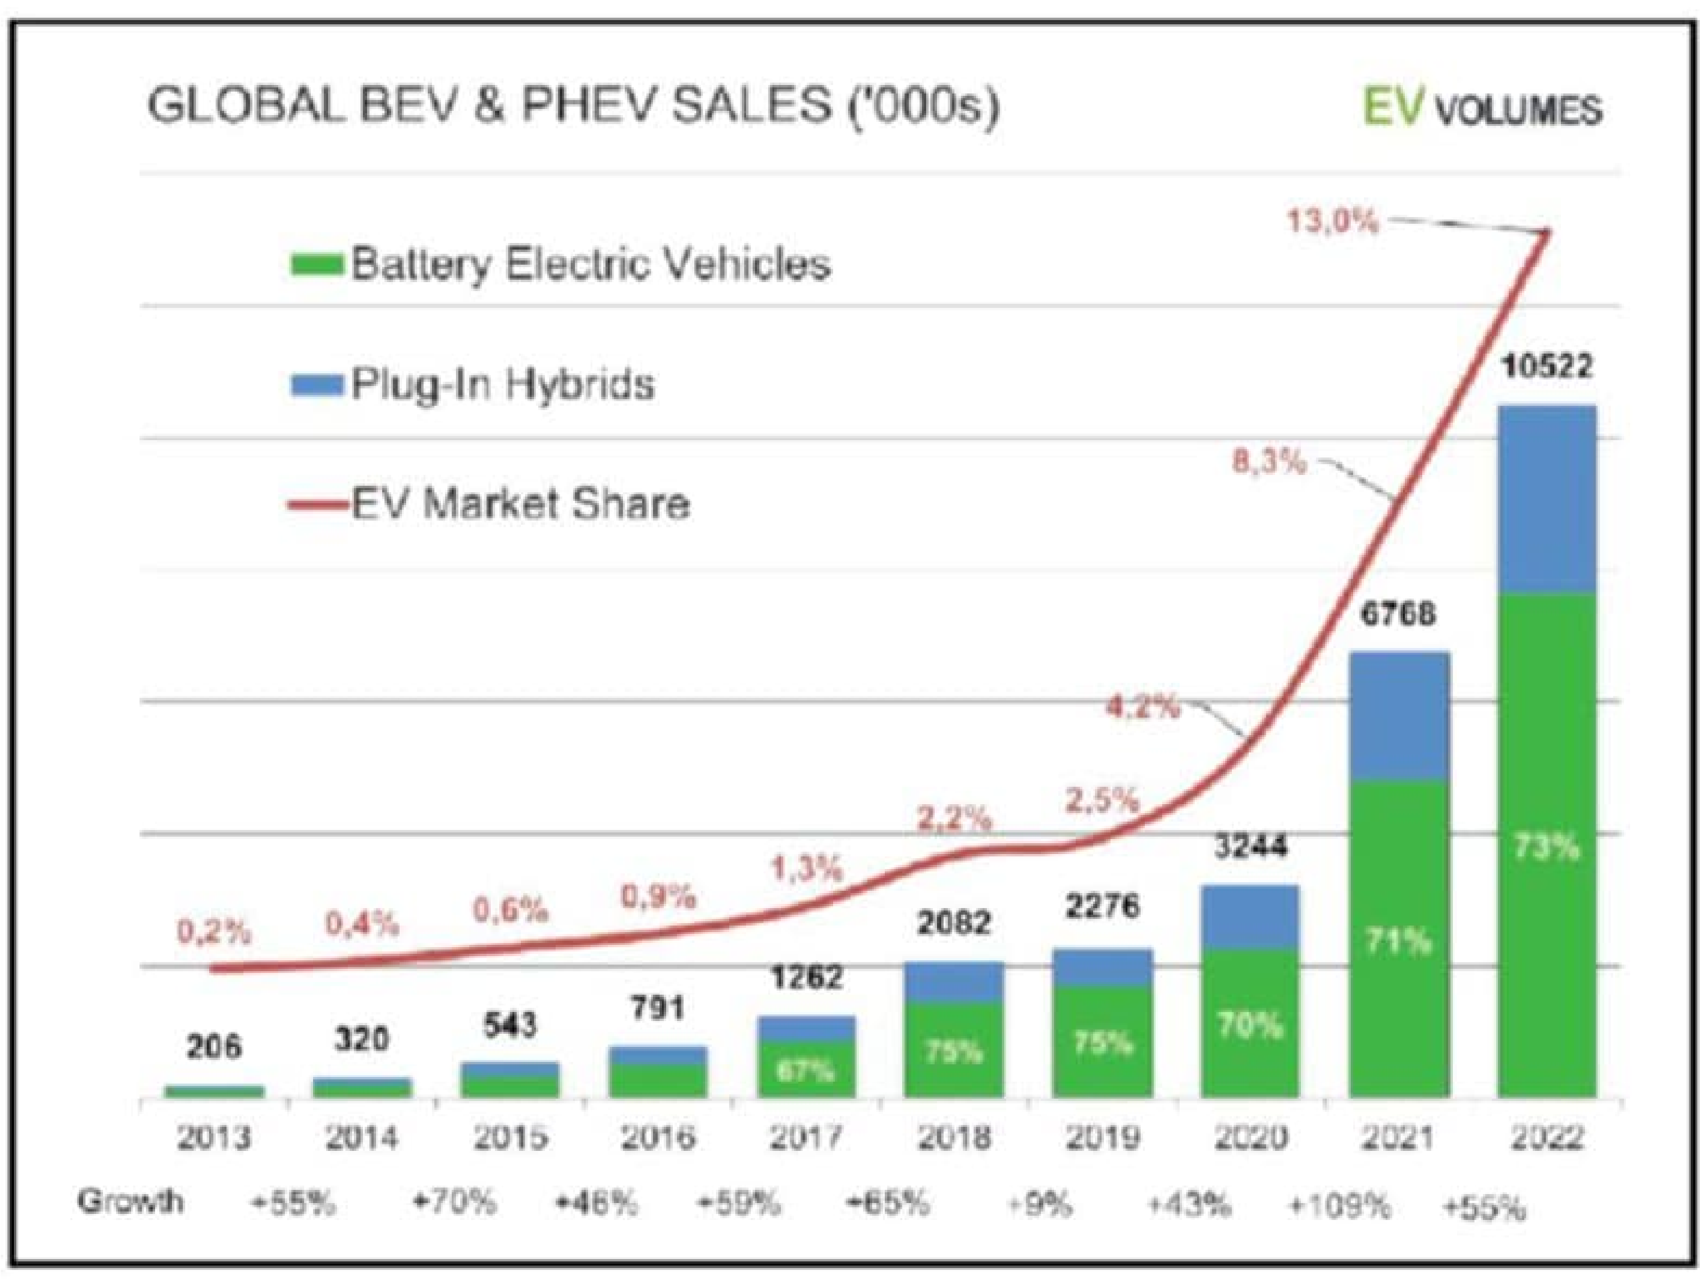

Robust increases of EV sales enabled nearly all OEMs to grow their sales in 2022. Global EV deliveries saw growth of +55% year over year in total; OEMs with higher growth have increased their share in the EV sector. (Image Source: EV-Volumes.com)

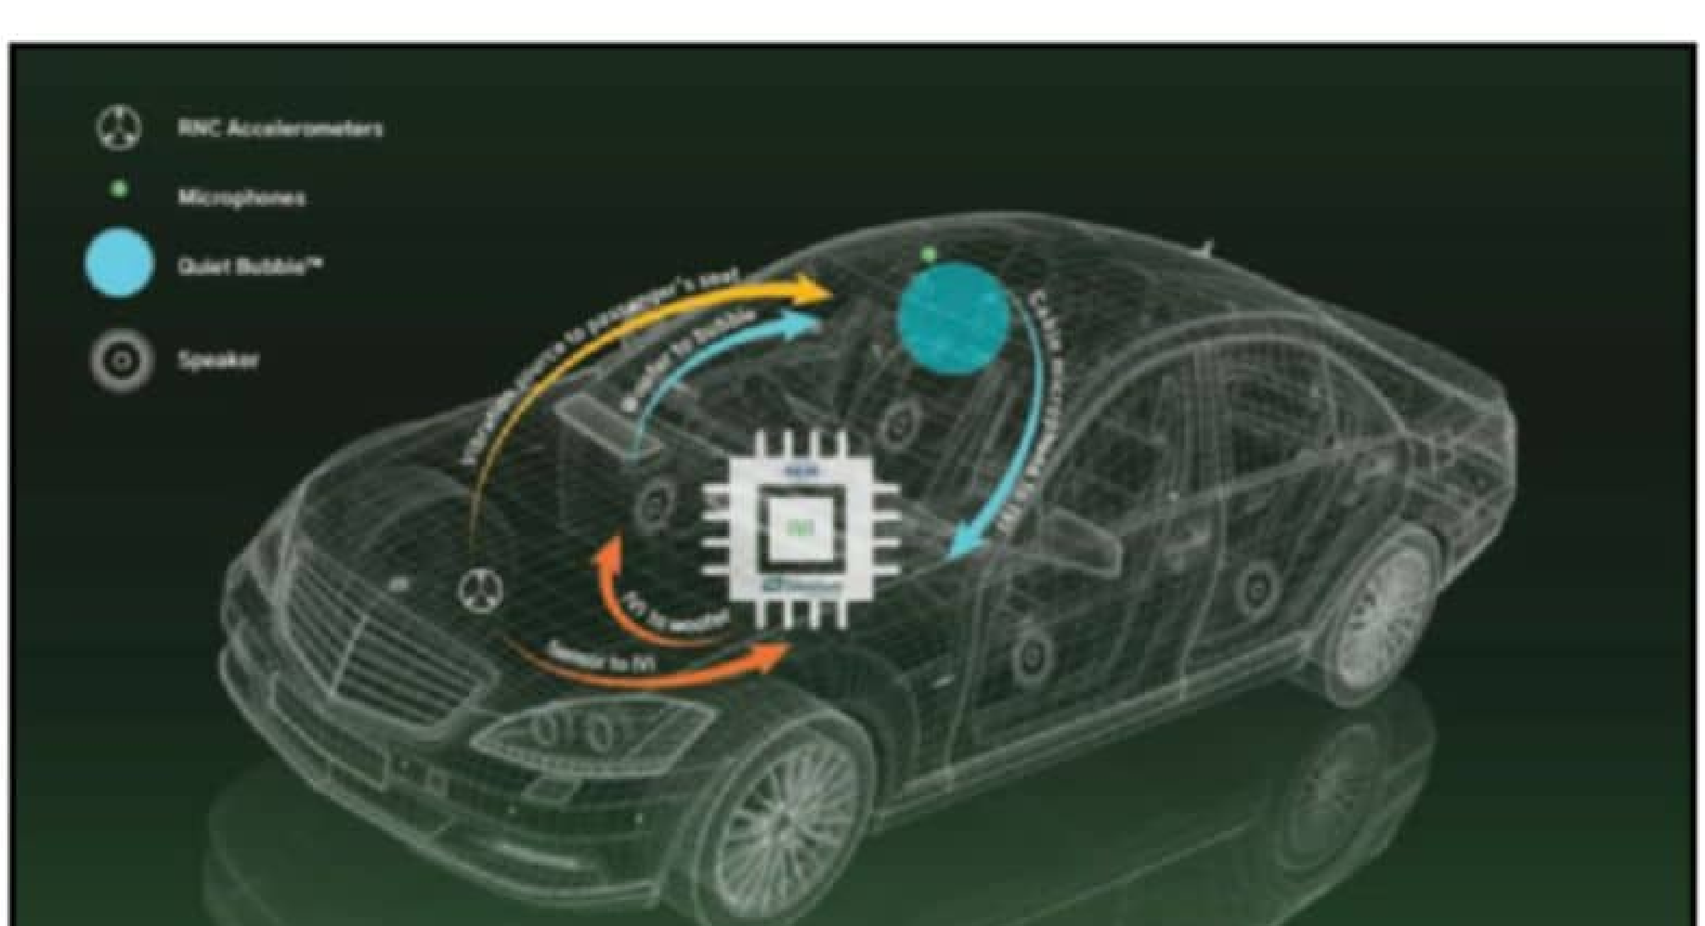

Keeping the inside of EVs quiet is a major challenge for car manufacturers. Active Road Noise Cancellation (ARNC) technology reduces the road noise in milliseconds, ensuring a comfortable user experience for the passengers. In 2021, Asahi Kasei Microdevices and Silentium agreed to work together on the development of a ready-to-use system specifically adapted for engine noise, one of the most sought-after solutions in the automotive industry.

worldwide EV sales to accelerate in 2023, with sales projected to reach 14 million units.

According to a different report from EV-Volumes, plug-in hybrid electric vehicles (PHEV) stood for 27% of global sales in 2022, while fully battery-powered electric vehicles (EVs, or BEVs) stood for 9.5% of global light vehicle sales in 2022.

These numbers are important to bear in mind when considering that almost all of the new automotive audio projects currently under development and scheduled to reach the market within the next 3 to 5 years are designed to explore and compensate for the unique cabin environment of EVs, with all its challenges and variables, where the combustion engine noise no longer masks everything else.

As an example, the growth of active noise control in the car is assured. The EV is the main driver for this technology, because with the disappearance of engine noise, the structural and road-induced noise becomes noticeable.

Noise control was also easier with the existing multiple-piece construction for the car body. Now that the industry has moved to a uni-body for the entire structure, the noise level goes up because the structure is one large vibrating piece. Removing weight to make allowances for the battery further increases the noise.

Automotive infotainment systems now provide an immersive cabin experience driven by consumer desires and expectations developed during the last half a decade or more. In-cabin features such as passenger entertainment, notifications of vehicle and environmental information, and navigation systems all use the audio system in automobiles. If the noise in the car cabin cannot be reduced mechanically or passively, it will need to be controlled with active acoustics. It is about more than just music. It is about comfort and an enhanced experience in the car.

Asahi Kasei Microdevices Partners with Silentium

Answering the demand for Active Road Noise Cancellation (ARNC) applications, and following a period of more than a year working together in development, Asahi Kasei Microdevices (AKM) has entered into a license agreement with Silentium for its proprietary broadband active acoustics technology. The Quiet Bubble software technology from the Israel-based company is now incorporated into one of AKM's audio voice processors and sampling to developers.

The cooperation brings ARNC technology to a new level by combining AKM's expertise in low-latency hardware components and Silentium's know-how in noise-cancelling software. The solution provided by both companies optimizes the control signal path from the reference sensors to the loudspeakers in order to achieve high-speed transmission of vehicle signals and low-latency algorithm processing. This is performed by optimizing specific parts of AKM's Analog-to-Digital Converter (ADC) so that the measured vehicle signals are quickly sent to the company's Digital Signal Processor (DSP) unit, where the Silentium low-latency algorithm's Quiet Bubble software is hosted.

Additional latency reductions can come from vehicle acoustics with the use of near-field loudspeakers in the headrest or at other locations in close proximity to the passenger's head. Since the hardware, software, and acoustic domains are all low-latency paths, this solution delivers in-vehicle broadband noise cancellation performance comparable to the premium noise-cancelling headphone experience.

Pushing the Envelope

How can we enhance the experience in EVs? A company with a vision on that front is Swedish digital audio pioneer Dirac, which has been on the vanguard of some of the most sophisticated projects so far. Dirac's automotive audio solutions are currently used by leading automotive brands such as Rolls Royce, Volvo, Polestar, BMW, and BYD, among many others.

Dirac has also been working closely with NIO, the innovative electric car company that is about to become the first Chinese manufacturer in the segment to become a global player. More specifically, Dirac helped create the sound system for its latest model to be launched in the US and Europe. To ensure the ultimate sound performance, the latest NIO ET7 car has adopted Dirac Opteo Professional, the company's most advanced automotive sound optimization solution.

Following the announcement that Dirac and Dolby had joined forces to design a high-quality immersive automotive audio experience for all-electric vehicles, the NIO ET7 is one of the first vehicles to feature Dirac audio optimization algorithms and a full Dolby Atmos audio system. The NIO ET7 features a standard configuration of 23 speakers and 20 amplifier output channels totaling 1000W—a significant challenge for a battery-dependent car. Its four main channels are composed of high, middle, and bass speakers in three units, with a subwoofer and four overhead channel speakers.

Combined with Dirac Opteo Professional, the automotive audio systems always renders an immersive sound experience, independently of the signal source. In this case, NIO first designed the sound system with other partners, and then turned to Dirac to optimize the performance.

Achieving peak sound performance in cars takes ingenuity and expertise. The car cabin is regarded as one of the most challenging acoustic environments in the world—as in-car speakers cannot perform as intended regardless of their price-point and design. Within the cabin, speakers can't be positioned as optimally as they would, for example, in a home theater. As a result, speakers interfere with each other, causing distortion and reducing audio clarity. The interference increases when more speakers are used. In addition, the car cabin itself adds unwanted sound coloration, resulting in muddy, booming sound that makes it difficult to discern the direction from which sound is coming.

Featuring patented multiple-input/multiple-output (MIMO) mixed phase impulse response correction technology, the Dirac Opteo Professional solution addresses these challenges by enabling all the speakers in a car to work intelligently together and co-correct each other's impulse response.

The solution "erases" the car cabin—removing the unwanted effects of the cabin in ways previously regarded as possible only in theory. It also creates ideal loudspeaker responses for maximum fidelity and bass performance.

Part of Dirac's Intelligent Audio Platform (IAP) suite, Dirac Opteo and Dirac Virtuo include patented digital sound optimization and spatialization technologies, including measurement-based magnitude response correction, virtual bass, and virtual center—to tackle common acoustic challenges such as negative car cabin impact, unbalanced soundstage, and lack of immersion.

As Lars Carlsson, Dirac General Manager of Automotive Audio comments, "The more speakers that automakers use to achieve a signature sound, the more important it is for Dirac to digitally optimize their performance and ensure the perfect listening experience, especially when enjoying immersive content."

"The NIO ET7 includes a sound system worthy of a concert hall," adds Carlsson. And when not enough immersive content is necessarily available to fully explore the car audio system, there's also the upmixing option, which Dirac is increasingly offering.

The latest version of its Dirac Virtuo automotive solution now features the company's Upmixer technology, Intelligent Audio Platform, and multichannel content support to envelope drivers and passengers in high-quality, immersive sound. Dirac's new Upmixer technology uses patent-pending algorithms to transform any stereo content into flexible, customizable immersive listening experiences, free from artifacts, and without requiring any hardware upgrades. This allows automotive brands to consistently equip their vehicles with immersive sound even when sources are based on stereo content.

Dirac and Dolby are working together to promote the possibilities of immersive audio environments in new electro vehicles. A recent example of a high-quality immersive automotive audio experience is the new all-electric NIO ET7, one of the first vehicles to feature Dirac audio optimization algorithms and a Dolby Atmos audio system.

Market Update — Automotive Audio

In 2022, Bose announced its first-ever sound system for Volvo Cars, with the all-new, all-electric Volvo EX90, featuring AudioPilot 3 noise compensation technology.

"Today's vehicles serve a much greater purpose than simply transporting occupants from place-to-place; they're remote offices, home theaters, and entertainment centers, all of which rely on high-quality, immersive sound," states Carlsson.

Car makers optimizing sound with Dirac's Intelligent Audio Platform benefit from evolving software, upon which all Dirac solutions are built and integrated with shared infrastructure and components using a streamlined workflow. It includes tools and software for both audio optimization and audio development, and Dirac's patented audio optimization algorithms are integrated into easy-to-use filter design tools.

Announced earlier in 2022, the Intelligent Audio Platform is a modular solution to easily develop, optimize, and integrate audio experiences with high efficiency. It is fully data-driven, 70% of the audio development and optimization work can be done virtually, and all results can be shared globally throughout an organization. Moreover, the Intelligent Audio Platform is supported by the industry's most common chipsets and frameworks.

In addition to the Intelligent Audio Platform, Upmixer technology and support for multichannel content, the Dirac Virtuo automotive solution also includes cabin sound optimization with advanced sound field control technology, which achieves better speaker performance and minimizes cabin colorations; target sound management, which achieves signature sound with greater consistency and ease; and center image optimization, which enables a stable phantom center for each listener with improved staging and localization.

Cadence Tensilica HiFi DSP now support Dolby Atmos for cars, making it the first DSP IP with this capability, offering a more energy-efficient implementation, according to the company.

Dirac has also leveraged a close collaboration with Dolby to address the challenges in automotive audio from different directions. Dirac uses its digital signal processing to upgrade the sound system performance and overcome acoustic challenges inherent in cars—such as suboptimal positioning of loudspeakers and reflective surfaces—so that the system faithfully reproduces content in Dolby Atmos.

Andreas Ehret, Director Automotive at Dolby, says that the Dirac technology is a great solution for car makers to optimize this experience. "We have found that in our demo cars, the results were excellent—with excellent detail, accurate staging, and a well-balanced sound field in all seating positions."

To accelerate the availability of its automotive solutions, Dirac has also expanded its collaboration with the leading semiconductor vendors. As an example, Dirac recently expanded its collaboration with NXP to combine Dirac's digital audio solutions for entry-level vehicle sound systems with NXP's SAF4000 multi-standard software defined radio processors. This collaboration allows automakers to digitally upgrade the performance of entry-level sound systems, which are typically those supplied in cars utilizing up to nine speakers.

Automotive Sound Platforms

Another key technology provider powering the immersive audio evolution in automotive is Cadence Design Systems, which has optimized the Dolby car experience technology to run on the Tensilica HiFi DSPs. Offloading Dolby Atmos to the low-energy, high-performance Tensilica HiFi DSP enables system-on-chip (SoC) providers to offer automotive OEMs and Tier 1 suppliers a more energy-efficient implementation without affecting the audio experience.

As the most widely used audio DSP on the market, HiFi DSPs are licensed by 19 of the top 20 semiconductor companies, with more than 1.5 billion HiFi DSP cores shipping per year. In collaboration with Dolby, Cadence has already brought Dolby technology to billions of products, including Dolby Atmos TVs. Now, Tensilica HiFi DSPs are delivering Dolby Atmos for cars with energy efficiency and value, while supporting other key processing solutions in voice, speech recognition, and active noise cancellation.

In 2022, Bose announced its first-ever sound system for Volvo Cars, with the all-new, all-electric Volvo EX90. Available by early 2024, the Volvo EX90 will extend to additional models in the coming years. To offer advanced audio options, Volvo will provide a 14-speaker Bose sound system featuring the latest generation of advanced digital signal processing, Centerpoint 360, with SurroundStage and BassSync technologies. In addition, AudioPilot 3 noise compensation technology monitors all sources of sustained background noise and automatically adjusts the music signal for a more consistent and enjoyable listening experience.

The new sound system for Volvo was developed using Perceptual Sound Rendering (PSR), a new software-enabled proprietary tuning approach from Bose, which provides more freedom and flexibility for system engineers to precisely tune audio performance for every seating position in the cabin.

RESONADO LABS

Introducing Res-Core Gen 2.

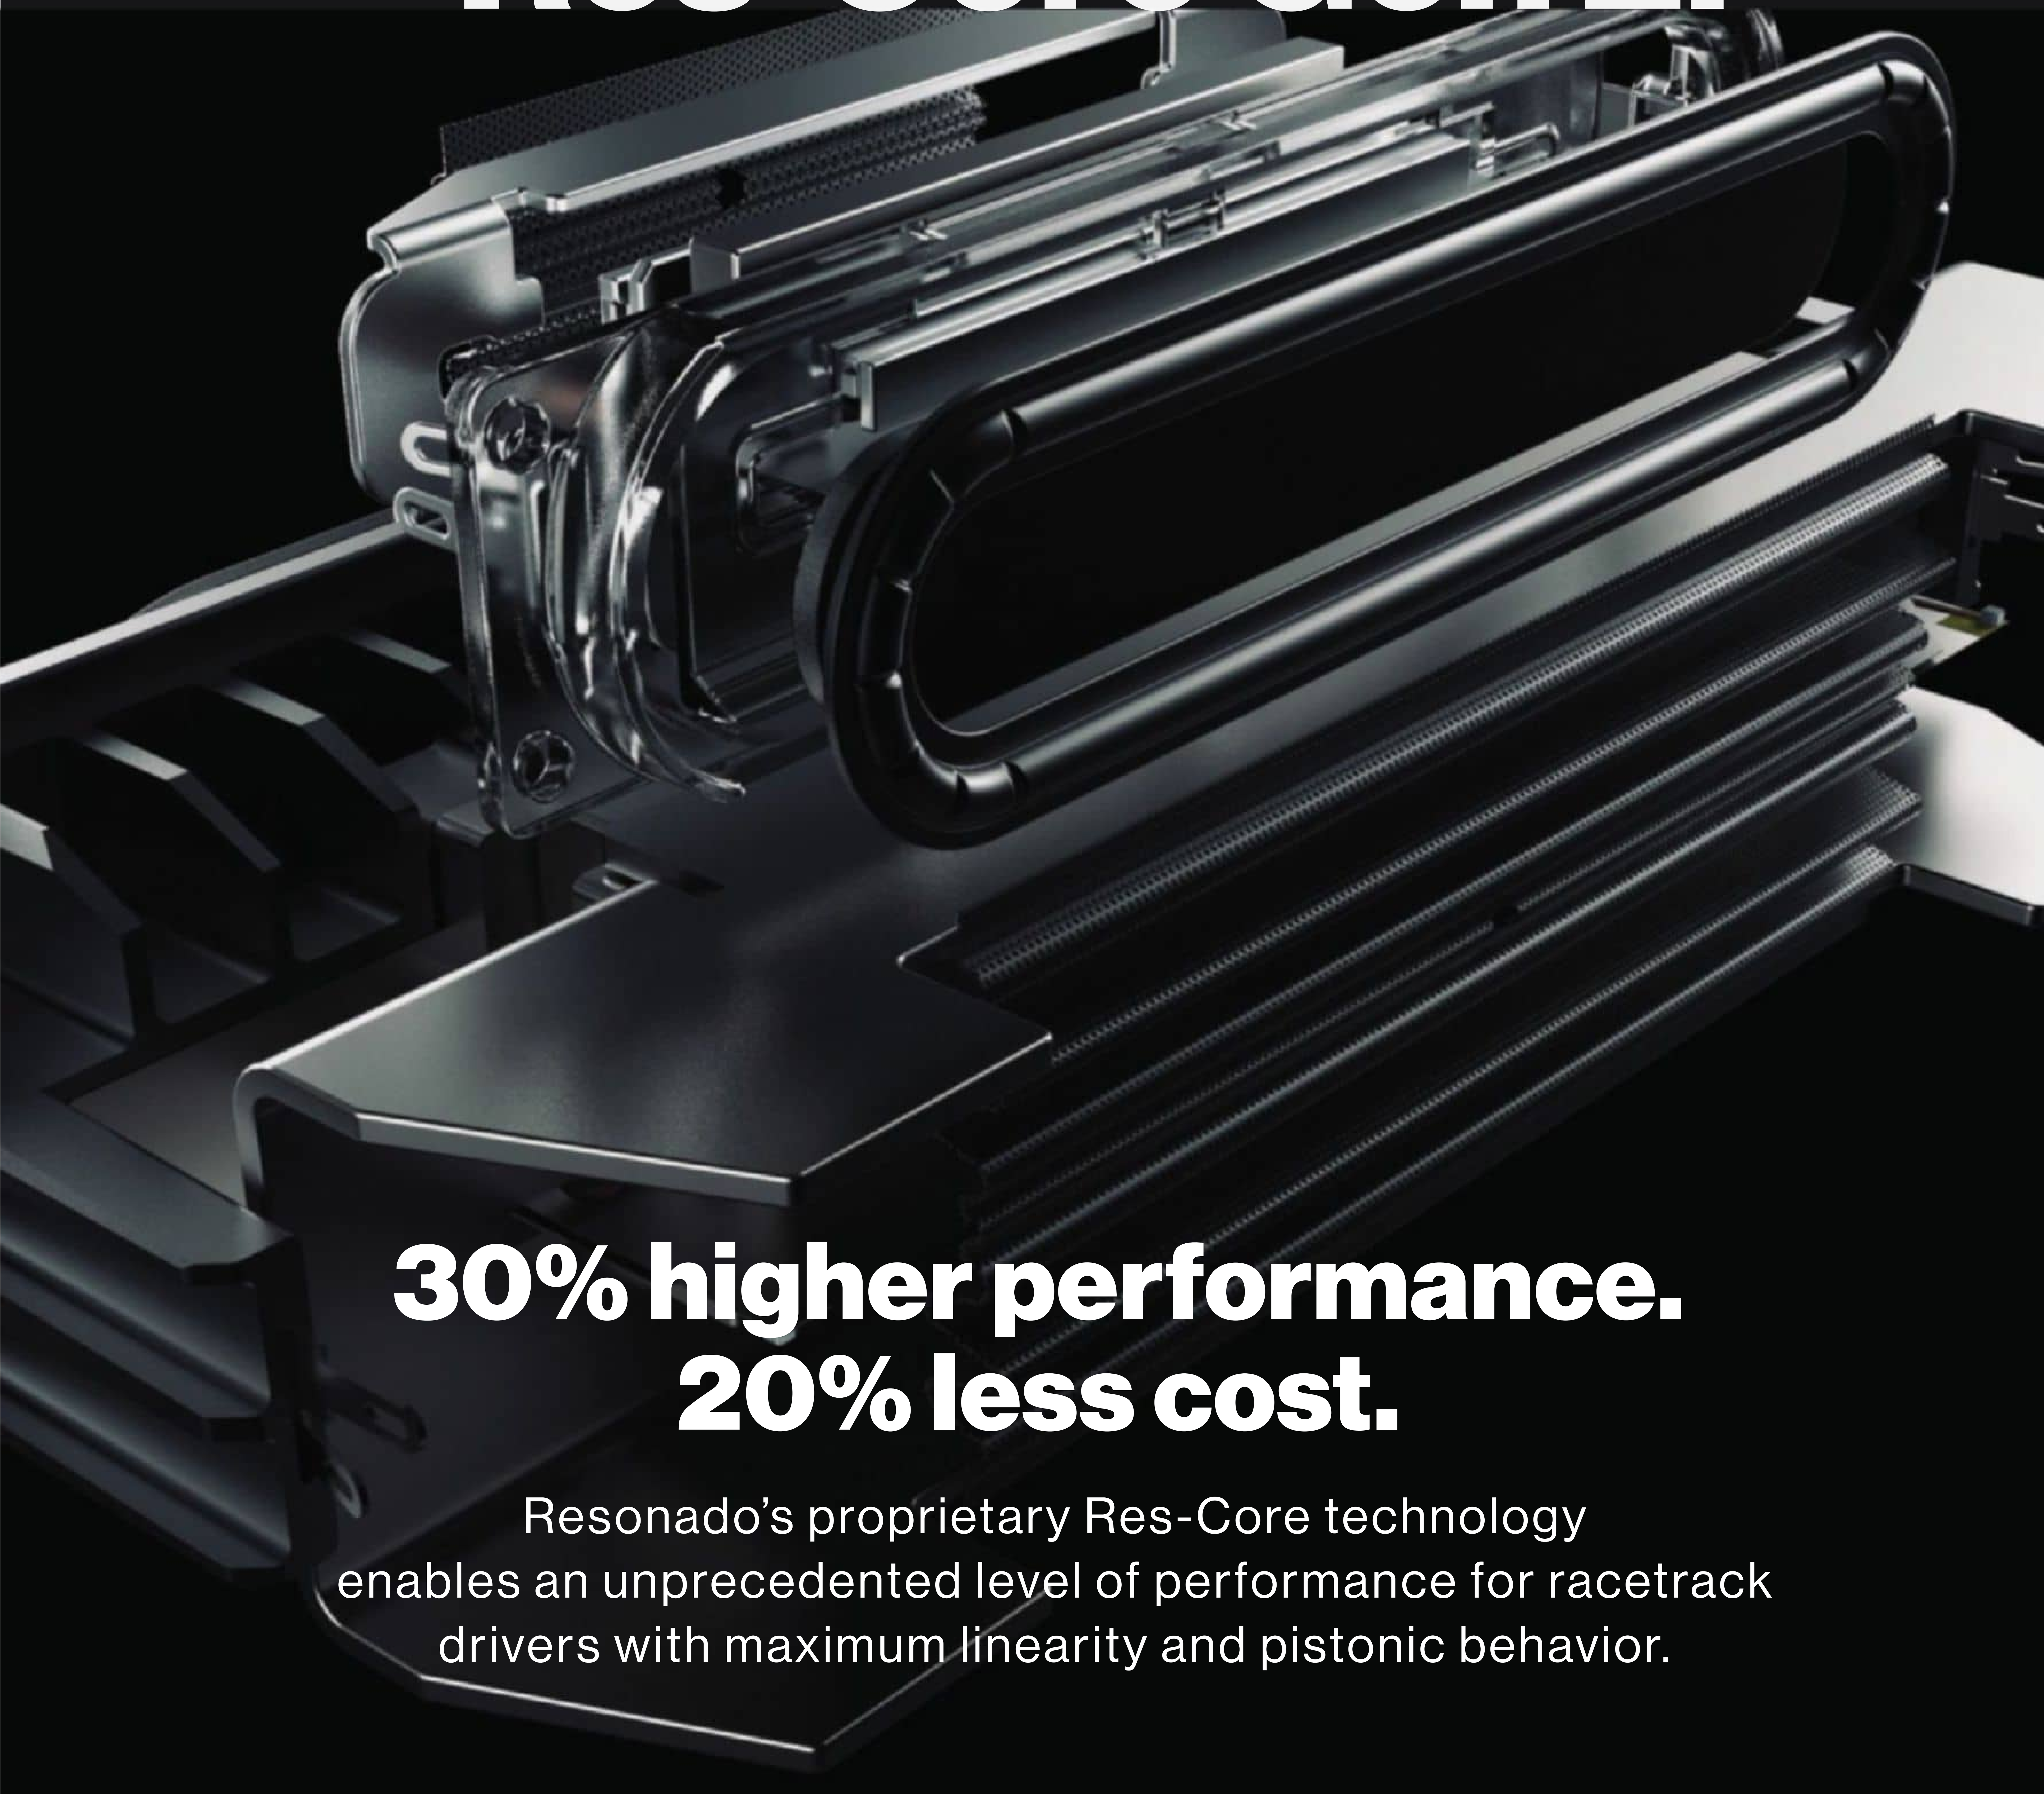

30% higher performance. 20% less cost.

Resonado's proprietary Res-Core technology enables an unprecedented level of performance for racetrack drivers with maximum linearity and pistonic behavior.

Res-Core Gen 2 driver platforms available this summer through licensed manufacturers. Reach out to sales@resonado.com to learn more.

The Maserati Luxury Experience is now fully electric and that poses completely new challenges in terms of sound. Expanding the partnership between the two iconic Italian brands, Sonus faber unveiled the integration of its sound system in Maserati's new GranTurismo vehicle, a groundbreaking new product that is 100% electric, and required a new immersive sound system.

Another key platform solution in the development and design of automotive audio systems is DSP Concepts' Audio Weaver, which covers almost every application need and is hardware and software independent, allowing developers to target any existing silicon, embedded DSP, and software-defined vehicles by offering a domain controller architecture on one high-computing platform.

Recently, DSP Concepts extended its collaboration with Analog Devices on a solution for automotive OEMs aimed at accelerating design cycles and reducing costs for in-vehicle audio system implementation using Analog Devices SHARC+ DSP technology.

The solution comprises a complete signal flow out of the box, with features that simplify audio system customization, tuning, and deployment. OEMs can incorporate a fully featured and pre-configured audio layout with essential user controls and features optimized for automotive cabin applications, including stereo audio output for systems with up to 12 channels, as well as advanced EQ and control capabilities.

The ADI LISTN EZ-AUDIO System by Analog Devices includes AWE Tune for audio tuning, working alongside the Audio Weaver development platform. Featuring an intuitive GUI, AWE Tune makes audio system customization and iteration easier with the ability to manage the entire Audio Weaver signal flow, tune multiple channels at once, create and duplicate presets, lock already-defined presets, and more.

All the tools needed for testing and tuning the system are included. Analog Devices offers the ADSP-21562-AUTO board, designed to provide premium audio for automotive head unit applications and featuring the ADSP-21562 audio processor with a powerful SHARC+ DSP core. The embedded Audio Weaver libraries have been optimized for the SHARC Plus architecture, including leveraging the on-chip FIR/IIR accelerators. CPU load monitoring and tuning report generation helps developers ensure that the system is capable of managing the entire Audio Weaver signal flow with adequate headroom.

Evolving Speakers

Designers of audio systems for electric vehicles have already learned the painful lesson that "power" is not necessarily a parameter at which they should be aiming.

In fact, the whole idea of placing dozens of power hungry speakers in a car for Dolby Atmos, while cutting battery range in half seems preposterous—that is unless the speakers' excursion could actually help move the car. Which indicates the pressing need to come up with new transducers that are exponentially more efficient, starting with... removing all the unnecessary weight.

Warwick Acoustics' electrostatic loudspeaker technology and audio systems are now engineered for automotive applications, where they are able to deliver significant benefits, and perfectly timed with OEMs' rapid transition to electrification.

Within Warwick Acoustics' Product Development and Acoustic Testing Laboratory, the company's engineers have already begun installing electroacoustic panel-based systems into demonstration vehicles, validating how in addition to the weight, efficiency, and design freedom benefits, the acoustics of this unique approach that has been so successful in the headphone arena, translates perfectly into an automotive cabin environment.

Some companies are taking serious steps in that direction and *audioXpress* has detailed some of the efforts in previous features. An example that deserves another mention is British company Warwick Acoustics and its electrostatic panels, already proven in high-end headphones, and now ready for car audio applications.

Following successful R&D and product development phases, Warwick Acoustics has entered 2023 with an accelerating program to commercialize its electroacoustic panel technology to the automotive industry. Warwick Acoustics is hiring and investing in facilities to support the substantial commercial pipeline under development, and believes it is ready to create immersive listening experiences for the premium in-car audio sector.

These new electroacoustic panels have been engineered from the ground up to be applied specifically to EVs as the multi-year development of this new technology is perfectly timed with OEMs' rapid transition to electrification and life-cycle sustainability.

Warwick Acoustics is a spin off of the School of Engineering at Warwick University, founded in 2002. The company is now headquartered at a new state-of-the-art facility at MIRA Technology Park, Warwickshire, in the very heart of the UK's automotive industry. There, Warwick has been expanding its world-class team with vast experience in technology and in-car audio sectors at brands such as Sony, Bose, Harman Automotive, Meridian Audio, McLaren Automotive, and JLR.

And the company is well under way with the commercialization of its electrostatic loudspeaker technology. Like-for-like, electroacoustic panels are up to 90% lighter, while consuming up to 90% less power and offer significant interior design and packaging freedom. Warwick Acoustics' thin driver panels use 100% mass upcycled and recyclable materials, providing OEMs with another key advantage to help them achieve their sustainability targets. Additionally, Warwick Acoustics' electroacoustic panels do not use any unsustainable rare earth metals—something that the design of conventional speakers continues to rely heavily upon.

Of course, other major players are investing in this field, even if they have to start from the beginning. Recently, LG Display announced that it has developed a new Thin Actuator solution with large potential as a new speaker technology for automotive applications. According to the company, the innovative solution is now ready to be commercialized, meeting the requirements from EV car makers looking for new-generation sound solutions. The Korean company, a leading innovator of display technologies, has long researched thin actuators for under-monitor applications. Until recently, when the company was still designing its own range of mobile phones, LG offered an innovative range of screen speakers with improved audio features in its LG ThinQ models. LG typically has explored this approach for flat TVs and smartphones.

While conventional speakers are large and heavy due to components such as the voice coil, cone, and magnet, LG's Thin Actuator Sound Solution is extremely thin and lightweight thanks to its unique film-type exciter technology. The company's actuator technology allows the device to vibrate off display panels and various materials inside the car body to enable a sound generation that can be converted into an immersive experience when combined with multiple units positioned in the cabin—controlled by a central processing unit.

The LG Thin Actuator Sound Solution comes in a passport-like size (150mm × 90mm) with a thickness of 2.5mm, equivalent to that of two coins stacked together, and a weight of just 40g, making it just 30% of the weight and 10% of the thickness of a conventional car speaker.

According to LG, the actuators' compact size and innovative form factor allow it to be installed in various parts inside the car such as the dashboard, headliner, pillar, and headrests, while eliminating the deviation in audio quality for passengers inside. The device's built-in nature not only allows the removal of speaker grills, but also enhances space efficiency by freeing up the space normally occupied by in-car speakers without compromising sound quality.

Harman's Expanding IP Portfolio

As the "automotive technology company and subsidiary of Samsung Electronics," Harman has been busy in its effort to lead the transition of automotive systems for EVs. Among many automotive systems on offer, Harman recently announced a new Sound and Vibration Sensor and External Microphone products to enhance the audio experience inside and outside the vehicle. It also launched Harman Ready on Demand, a software platform and marketplace that allows users to download and activate new audio features.

The new embedded audio products enable a variety of applications to enhance safety and user experience, from detecting emergency

Harman Ready on Demand is a software platform for delivering branded audio value, feature enhancements, upgrades, and monetization opportunities for car makers.

vehicle sirens to listening for exterior speech commands from drivers or traffic controllers to detecting glass breakage or vehicle impact.

From increased awareness of outside sounds inside the vehicle to enhanced communication possibilities for drivers, passengers and emergency vehicles on the road, Harman's Sound and Vibration Sensor and External Microphone offer adaptable and future-looking features. The completely sealed piezo-based Sound and Vibration Sensor can be invisibly integrated into a vehicle's exterior while the External Microphone is designed to withstand environmental elements and can be configured as a single element or multi-element array.

The Sound and Vibration Sensor is a fully developed design that allows for invisible vehicle integration, enabling exterior vehicle speech recognition and acoustic sensing. As future vehicles advance, they will likely be required to include enhanced external sensors to increase awareness of exterior sounds for vehicle occupants and the vehicle systems, such as feeding the acoustic alerts delivered by a Harman Ready Vision system. The devices exhibit a wide band frequency response for accurate speech recognition and acoustic sensing performance, enabling potential AI features such as audio object detection.

Harman has more than 30 years of experience designing, optimizing, and integrating automotive microphones and sensors into vehicles. Both the Sound and Vibration Sensor and External Microphone are unique, patented designs and have been specifically designed to meet the needs of OEMs and consumers, delivering improved safety features and connectivity for all vehicle segments, from entry level to luxury.

Harman is also promoting its Ready on Demand software platform for delivering branded audio value, feature enhancements, upgrades, and monetization opportunities in an easy to-use app. An industry-first product, Ready on Demand is the company's foundation for providing expanded experiences and future upgrades that can be unlocked by the consumer at any time via in-app purchases throughout the life of the vehicle. The expectation is that consumers will be able to improve in-vehicle experiences and personalize their vehicle with premium audio features at any time.

Ready on Demand is pre-integrated with Harman's AudioworX audio processing libraries, and the Harman Ignite cloud back-end for e-commerce and content download. This will significantly reduce required development and increase speed of implementation.

Examples of currently defined packages include personalization and sound enhancements that activate proprietary acoustic environments, adding a layer of ambiance, serenity, or excitement to the content playing inside the vehicle.

Harman's Sound and Vibration Sensor and External Microphone provide an integrated sensor product that enables acoustic monitoring of speech commands and sounds on the outside of the vehicle.

In comparison to conventional audio systems, Continental says that its Ac2ated Sound thin actuator technology not only produces high audio quality but also enables a reduction of the weight and space taken up by up to 90%.

Resources

EV Volumes.com, www.ev-volumes.com

"Global Electric Vehicle Sales Crossed 10 Million in 2022; Q4 Sales up 53% YoY," Counterpoint Technology Market Research, March 2023, www.counterpointresearch.com/global-ev-sales-q4-2022

"Global EV Outlook 2023: Catching up with Climate Ambitions," International Energy Agency (IEA), www.iea.org/reports/global-ev-outlook-2023/executive-summary

Sources

Analog Devices | www.analog.com/EV-21562-AUTO

Asahi Kasei Corp. | www.asahi-kasei.com

Bose Corp. | automotive.bose.com

Cadence Design Systems, Inc. | www.cadence.com

Continental AG | www.continental-automotive.com

Dirac | www.dirac.com

Dolby | www.dolby.com

DSP Concepts, Inc. | www.dspconcepts.com

DTS, Inc. | www.dts.com

Harman International | www.harman.com

LG Electronics | www.lg.com

NXP |www.nxp.com

Sennheiser Electronic GmbH & Co. | www.sennheiser.com

Silentium | www.silentium.com

Warwick Acoustics | www.warwickacoustics.com

XPERI, Inc. | www.xperi.com

With Harman Ready Upgrade, the company intends to offer fully upgradeable hardware and software products that allow OEMs to perform upgrades throughout the life of the vehicle and accelerate time to market, while enabling consumers to add and upgrade vehicle features as easily as they would with their smartphones.

Sennheiser Continental

Another example of an innovative approach to automotive audio starting from the transducers is the technology developed by Continental with its speakerless Ac2ated Sound System, which is based on actuators that excite selected surfaces in the vehicle interior to produce sound. In comparison to conventional audio systems, Ac2ated Sound enables a reduction of weight and space of up to 90%. This technology is further enhanced with Sennheiser's AMBEO software, tuned to the vehicle cabin environment.

Sennheiser is entering the automotive industry with its software technology and sound tuning know-how combined in the AMBEO Mobility suite. Sennheiser's software suite of immersive audio solutions combines solutions for immersive sound, individual sound zones, noise management, and communications for driver and passengers.

The immersive approach is provided by its AMBEO Concert audio software module, with traditional stereo content unmixed to a multi-speaker and multichannel environment. AMBEO Concerto distills the fundamental components of any piece of music, such as different instruments and room information, and redistributes them throughout the car. Concerto can also be used to playback native immersive audio formats, such as Dolby Atmos.

Another solution is an immersive sound zone algorithm for headrest audio. AMBEO Sonata can be added to a traditional multi-speaker setup in the car to enhance the spaciousness. Alternatively, entry-range models can use headrests with Sonata processing to replace a more complex multi-speaker sound system.

In combination with Continental's actuator system, this software suite is able to create an extremely natural and enveloping sound experience for the occupants, who feel as if they are surrounded by sound.

The next generation of Morgan sports cars will benefit from Sennheiser's software technology to deliver a quality sound experience with the support of Continental Engineering Service's Ac2ated Sound System.

Time Keeps Slipping into the Future

By Roger Shively

It takes 4ms for road noise to reach the passenger's ears in a car. Active noise control (ANC) technology is mostly effective when the noise is constant, predictable, and low in frequency, like the narrow-band rumble of an engine or the resonance of a tire when rolling at speed. Random wideband noise events such as short rough patches or potholes can break through if the ANC's signal and calculation speed for road noise can't react quickly enough to generate a counter wave.

The signal latency and processing in active road noise control (ARNC) needs to be fast and quicker than the speed of noise traveling from the road to the cabin. The latency in that is of concern, whole system latency, from the noise sensor and the error microphones, the signal buss, the chip processing, and any audio framework in which the chip is used, and the ARNC's algorithm library. For ARNC to be effective and perceived by the human passenger, that latency usually needs to be 2ms or less. ARNC analyzes the road noise and sends an inverted sound wave via the car's audio system speakers, flattens out the noise before anyone in the cabin hears it. In some early applications, RANC reduced cabin noise by half (i.e., 3dB). It continues to

Tesla introduced Active Road Noise Reduction (ARNC) in the company's two flagship Model S and X vehicles in 2021, activated with a software update.

improve. Depending on the system's combination of speakers, sensors, microphones, and microphone location, 6dB to 9dB is obtainable in low frequencies (<500Hz). There are also high-frequency solutions for frequencies up to 1kHz. High frequencies require closer coupling to the ear for the playback speakers and the error microphone. And latency becomes even more critical. The highest quality algorithms use a number of microphones in relation to the head location and are adaptive to the acoustic behavior of the cabin in relation to speed and road condition. As the noise path changes, the system must adapt. In order to do this, the relevant noise signatures need to be mapped in a vehicle, and the algorithm simulated for reducing noise, which can be made depending on how many occupants there are.

In a demo car, where everything is known and tightly controlled, it is easy create the wow effect, of even more significant noise reduction in decibels, but that demo effect would not hold up in production. In production, the audio system and the car have variability, the noise is not entirely predictable, the algorithm must adapt, and random inadvertent noises, need to be taken into account.

Implementation Approach

To optimize the implementation of an algorithm directly on a DSP chip or on an SoC using their audio framework is another, more direct, set of work. This can be time consuming and labor intensive for an algorithm provider because they will be integrating their solution on any of the DSPs or SoCs that are being used for a given system. The advantage is on the ability to fully protect their intellectual property and manage the latency themselves. This means most of a noise control team is porting code on any given program. If the ANC process is new to a company, it requires an OEM to create a new internal process and the timeframe for the project becomes relative. If it's a new process in Japan, it will take 5 years to implement because of the level of detail. If it's a new process in Europe, it will be 2-3 years. If it's new in China, 1 year. The ARNC team could be as many as five people per car line, dedicated to the years needed to develop the car. But, most OEMs are getting more used to ANC, have developed a process, and time is getting shorter.

To reduce the cycle time and personnel required for software integration across DSPs an SoCs, it is also possible to work on a software platform, such as QNX's Hypervisor or DSP Concepts Audio Weaver. The question then becomes where does the low-latency feature reside and how is the rest of the system interacting with it? The selling point of a software platform should be that you can be flexible and not dependent on the DSP processor. So, you make one generic DSP block of your algorithm and then you can scale it according to the audio level in the car.

The software platform provides the code-optimized DSP blocks for the specific DSPs and SoCs. The algorithm providers also need to be flexible in terms of their DSP RNC libraries so they can scale their solution accordingly. The first step then is to use the software platform to define the channels of the components, loudspeakers, microphones, and accelerometers. The second step is to run a simulation to predict the algorithm settings based on vehicle noise data to test the RNC performances with the specific ARNC system that it was defined in the software platform. If the second step meets the acoustic targets, then the software platform compiles the ARNC block of code for the specific SoC/DSP audio framework.

As an example, Qualcomm is offering four Hexagon DSPs on their Snapdragon SoC. They used to work with four general-purpose DSPs in an implementation, but have now moved to their own SoC. For an audio framework, Qualcomm offers Audio Reach and they have Audio Reach Lite for the low-latency applications. For a system based on Samsung's Exynos processor, everything runs on Tensilica HiFi5 DSPs, without separate audio frameworks. Both systems have all the multi-core threading worked out and can be integrated into a software platform.

It's All About Clock Management

In the automotive market, audio processing for active noise control will use anywhere from 8 to 16 samples at a 48kHz frame size (0.16667ms to 0.33333ms). The desired latency from microphone input to speaker output is under 1ms to impact noise content up to ~700Hz or a little more. The actual benefit also depends on the correction mic and speaker locations relative to the listener ear.

Some algorithm providers will work at much lower clock speeds to implement FIR filters more efficiently. This would be taking the 48kHz clock speed and down-sampling it for the processing work on a sample by sample basis and then upsample the system clock. This requires working directly with the DSP/SoC or creating multiple levels of software modules in a software platform so the system level instructions (not requiring low, low latency) can be managed at the higher clock speed and the ARNC instructions (requiring low latency) can be managed separately. There would need to be not only multi-rate support but multithreads for instructions so low latency can be maintained. So, for example, an algorithm developer could down sample to 1.5kHz (2ms latency) or 3kHz (1.6ms latency), depending on the number of samples they needed to work with for the different quality levels of ARNC that would be provided.

About the Author

Roger Shively is the owner of Shively Acoustics International in Seattle, WA. He has more than 35 years of experience in engineering research and development, with significant experience in product realization and in launching new products at OEM manufacturers around the world. Roger co-founded JJR Acoustics in 2011 and has been a principal since its founding. Before JJR Acoustics, Roger worked as Chief Engineer of Acoustic Systems as well as functional manager for North American and Asian engineering product development teams in the Automotive Division of Harman International Industries Inc., a journey that began in 1986.

At 3kHz, this would allow six samples to work with in 2ms, while 1.5kHz would allow two samples to work with in 1.6ms, and they could work in blocks of eight samples to achieve their goals of applying multiple FIR filters with multiple inputs and outputs to receive and control state variables of the ARNC system.

Once the DSP/SoC level latency is controlled, the rest is system or software platform concerns. If a software platform is used, the multi-threading can be critical. If everything is maintained there, an additional latency concern for the algorithm integration, can be that for CAN bus, or other bus system messaging, such as that provided by A²B.

Integration is also required in the case of controlling an audio algorithm that might not be done in the software platform without actually writing the code for it and reading the messages from an automotive MCU, which decodes the bus messages. ADI's A²B claims the industry's lowest latency at 2ms. This on a high-speed digital interconnect technology, which distributes audio and control data together with clock and power over a single, unshielded twisted-pair wire.

Current deployments of automotive road noise cancellation need the availability of a cost-effective, low latency networking technology that efficiently connects required input sensors to a central processing unit.

Demonstration of the ARNC integrated solution from Asahi-Kasei and Silentium at CES 2023.

Perspectives on Automotive Sound Design

In the Era of Electric Vehicles

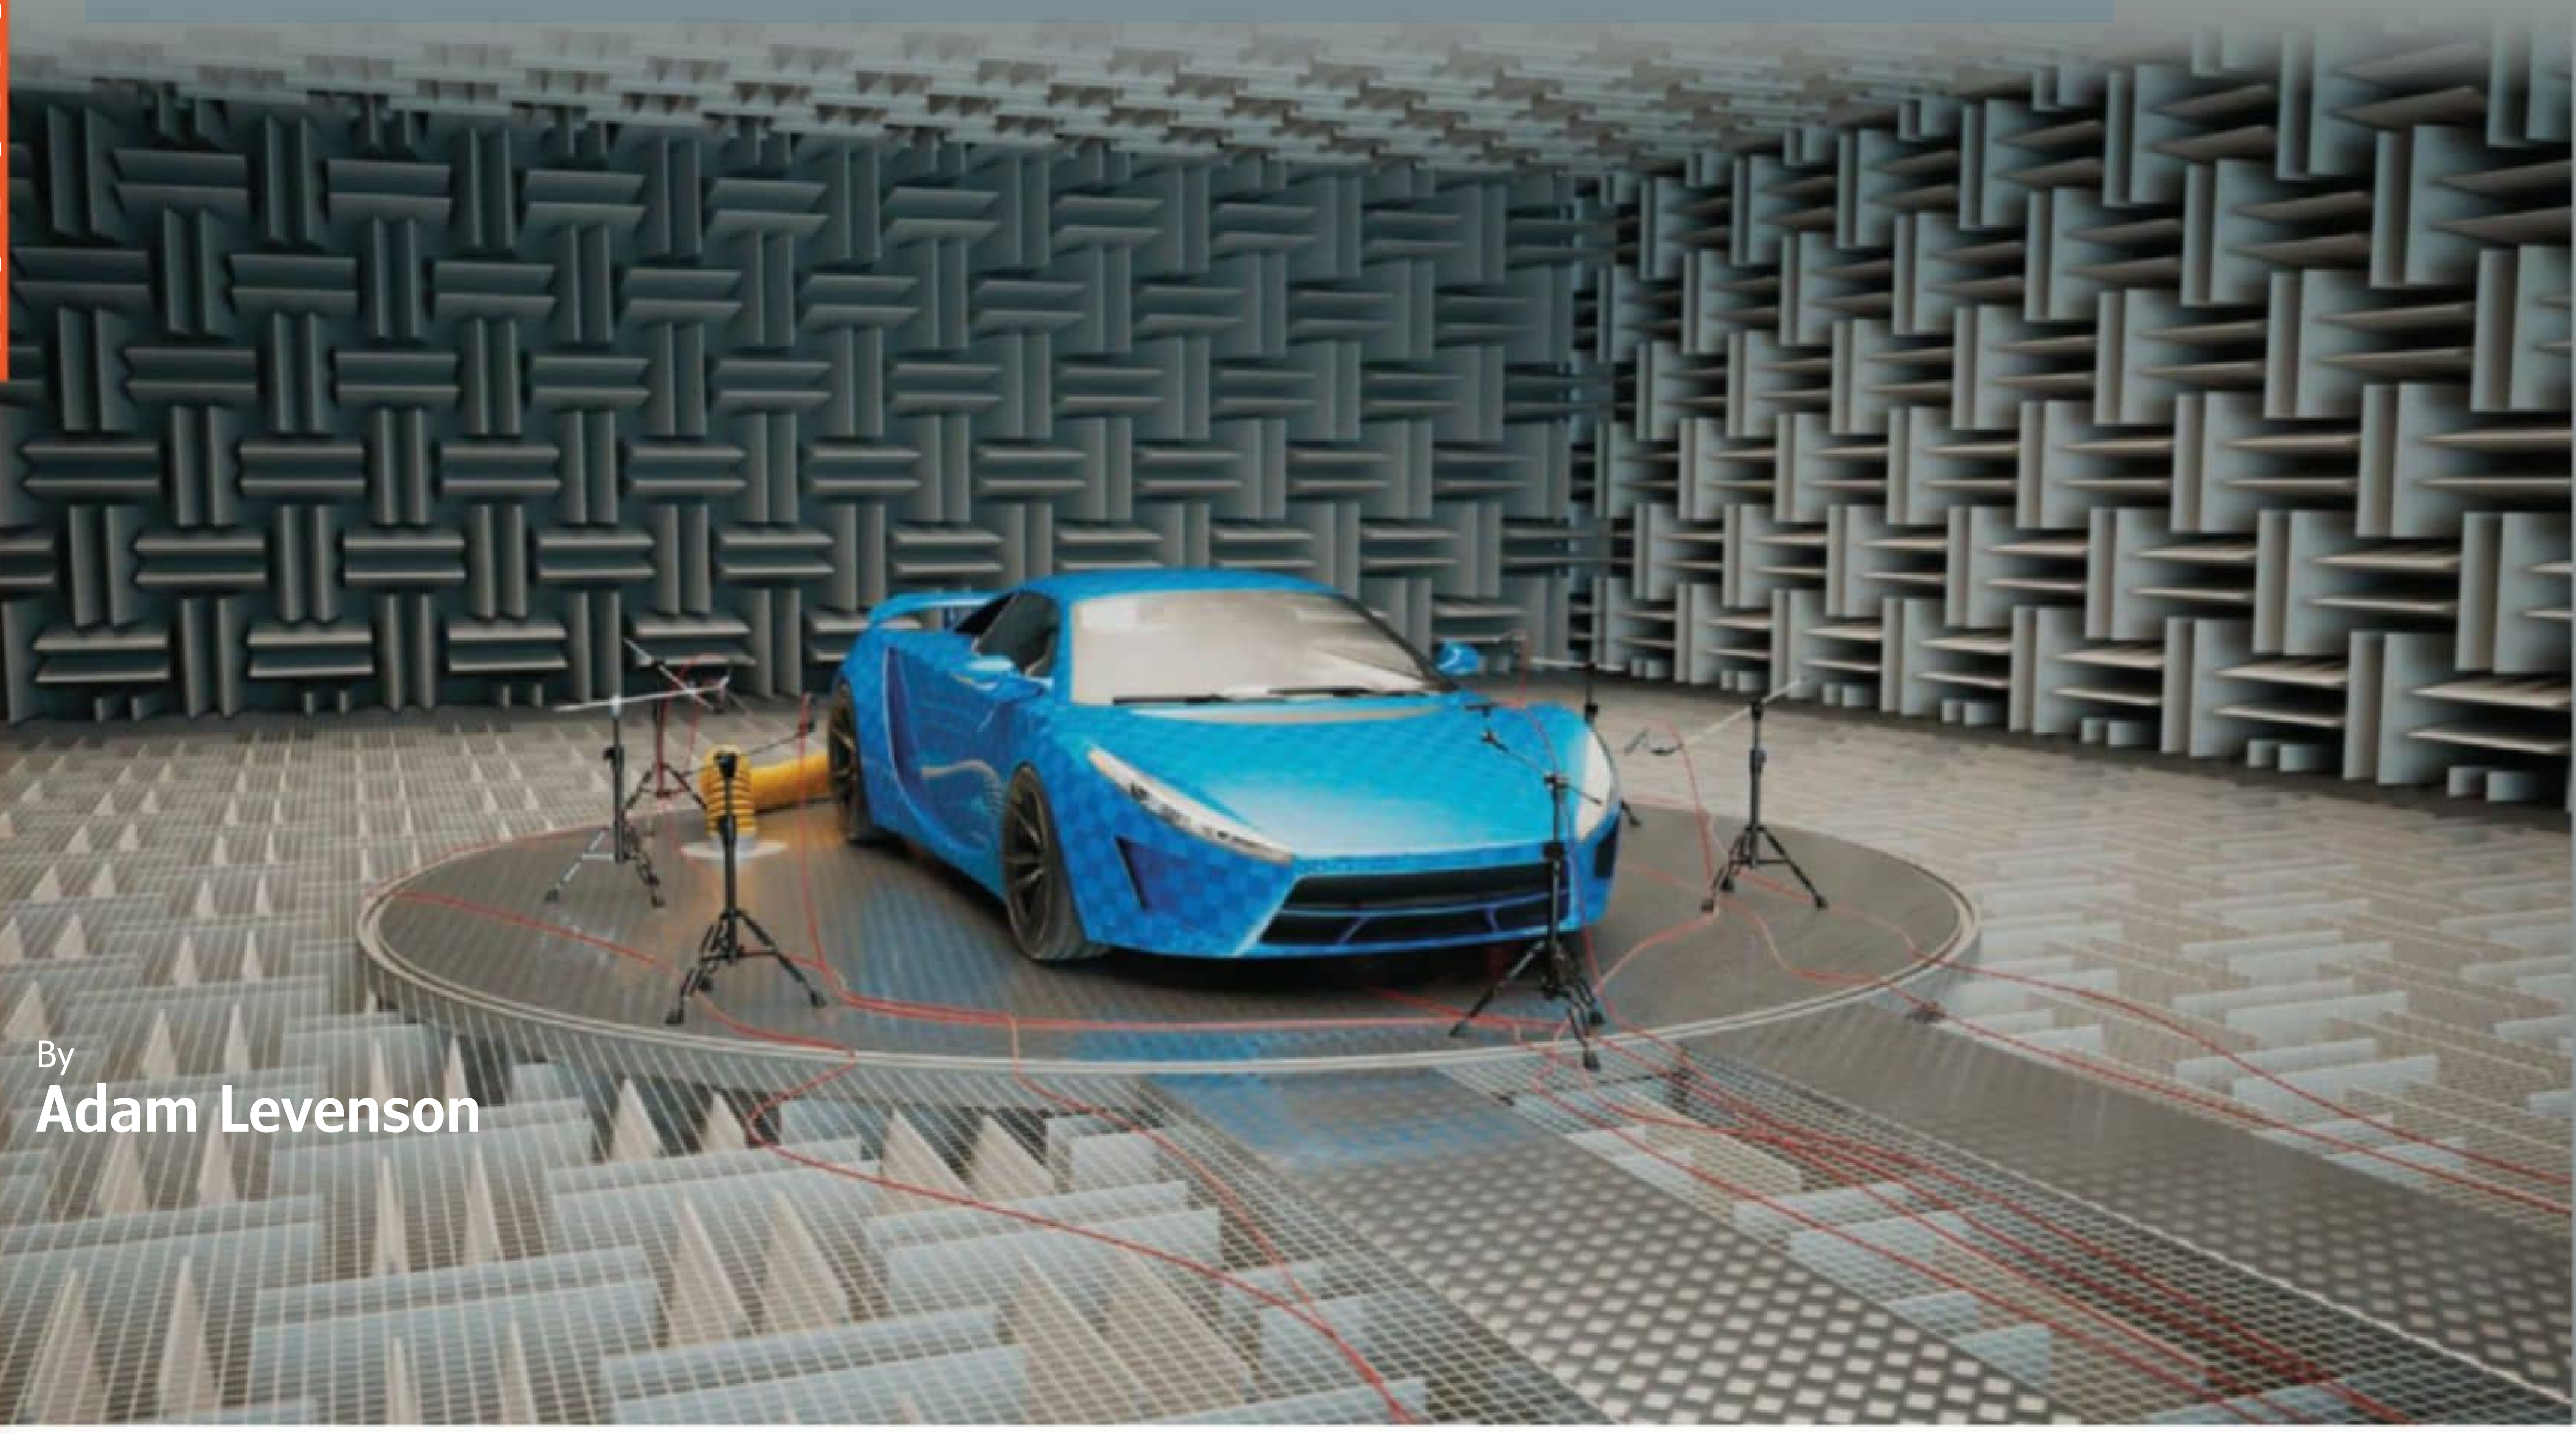

By
Adam Levenson

From combustion engine vehicles to fully electric vehicles, and from sound systems in the car to sound design and synthesis for the car itself, everything is changing in the automotive industry. This conversation with two members of the DSP Concepts automotive team with vast experience in the field, John Whitecar, VP of Product Management, and Steve Ernst, Head of Business Development, offers valuable insights on how things are evolving.

As the automotive industry moves toward a fully electrified future, sound design is becoming increasingly integral to the driving experience. Phasing out the internal combustion engine is better for the environment, but the unintended consequence is the loss of the sonic masking of the engine mechanics. That means exposing the tire, electric engine, and other unfamiliar and discordant sounds. It also means losing the built-in safety aspect of the gas-powered engine: Its unmistakable sound draws people's attention, especially in close proximity when moving away is a life saver.

In addition to safety, sound design has opened up another way for OEMs to build up brand character. By developing the internal soundscape and a distinctive external sound profile, manufacturers can differentiate from competitors while deepening the sense of connection for owners.

A wave of sonic innovation and creativity is likely to follow the rising popularity of EVs and the gradual advancement of autonomous vehicles. Aside from the lofty brand and artistic developments, the legal regulations now in place in major markets make sound design mandatory for EVs.

Back in October 2022, Dodge announced the new Charger EV and demonstrated its synthesized engine sound in front of a live audience. This wasn't an effect built in post-production for a sci-fi movie or a video game. This is a real car with real-world engine characteristics, and muscle brand equity in its iconic growl. Could it possibly rattle windows, set off car alarms, and give users that shot of combustion-inspired adrenalin like its gas burning predecessors? Will present-day drivers embrace synthesized engine sounds and appreciate the safety and interactivity that's designed into modern EVs? Is the driving public ready for cars with fully orchestrated interior and exterior soundscapes that start even before you open the doors?

Steve Ernst, Head of Business Development (left) and John Whitecar, VP of Product Management (right) from the DSP Concepts automotive team

We got together with John Whitecar and Steve Ernst from the DSP Concepts automotive team to discuss automotive sound design and synthesis. John Whitecar's experience in automotive infotainment architecture, system design, vehicle acoustics, digital receiver and digital signal processing spans more than 40 years and has impacted the audio experience in millions of products, including the first EVs. Before joining DSP Concepts, John Whitecar was Senior Manager Wireless Engineering, Principal Engineer at Tesla.

Steve Ernst, Head of Automotive, is a 25-year automotive business development and sales veteran leading global account teams, product strategy and product launches including his tenure as a global sales director at Harman International. His real-world view of the automotive industry is invaluable to understand the current trends.

Terminology

- Acoustic Vehicle Alert System (AVAS)—Pedestrian warning system that meets Functional Safety (FuSa) requirements
- Engine Sound Synthesis (ESS)—External sound designed to augment combustion engine noise
- Interior Sound Enhancement (ISE)—Designed to augment combustion engine sound or create soundscapes for electric vehicles
- Exterior Sound Enhancement (ESE)—Exterior sounds other than AVAS

Adam Levenson: Essentially, there are two sides to building the automotive sound experience: design and technology. Let's begin by describing the design side—the creation and composition of the sound elements.

John Whitecar: Sound design is the creative, artistic side. It's an art form focused on representing the vehicle, and the vehicle brand idea in a soundscape. So whether it's just welcome sounds or engine sounds or even the chime notes, that's all developed through an artistic process. It's the creative work to formulate what the car maker wants the user to experience, and how the brand is recognized on an auditory level. OEMs build brand identity that includes drivability, powertrain, interior trim, and UI. All of that is coupled very tightly to the brand, and the sound design is really just one more aspect of that. So it goes far beyond just normal audio playback: it's all about creating an environment.

Steve Ernst: When design is done holistically, for both internal and external sounds, and it reflects on the brand itself, there's a continuity between all these sounds. You really didn't get that in the past because these systems were managed differently and by different divisions and organizations. For powertrain, there's the engine sound, there's the exhaust sound, they're making mechanical changes to create a sound that they want. They're not thinking at all about what's happening with the chimes or the alerts.

Today, all of this is converging into a common design platform both technologically with tools like Audio Weaver, and artistically with a common design concept. There's now so many different sounds for so many more purposes, and the design platform is like a canvas that allows the OEM to have continuity between all of these sounds. All of them play into the brand instead of being disjointed.

Levenson: Is that a relatively new way of thinking for car makers?

Whitecar: I think it's changing. You used to tune the muffler. You would tune the exhaust system by tweaking resonance in the muffler and the tailpipe to enhance notes from the engine. It was all done acoustically and mechanically. Effectively, they were designing an instrument, like a tuba. There wasn't an idea that these tunings paired to the interior. For some brands the goal was to achieve absolute quiet in the cabin, for others, like Porsche, the idea was to connect with the car. I think the key thing that is different now is the capability to create sounds. In the past, most of the alerts came out of a chime module and it was very limited: just beeps and very simple tones. Now, we have this capability to create a synthesizer and play anything you want at any time based on the vehicle feedback. It's like a blank canvas now for sound designers.

Ernst: As they transition to electric vehicles, OEMs are now having to rethink what the sound of the vehicle is going to be, especially externally. It's a tough challenge because you want to be genuine, but it's hard to attain. How can you be genuine when all you have is a very low-pitched whine that's coming out of your electric engine? They're having to rethink how to project some type of sound quality that makes that user think of power, performance, luxury, or speed. It's not going to sound like it used to so OEMs have to reinvent to be as authentic to the virtues of the vehicle and what the brand represents.

Levenson: The current reference point for authenticity is still the internal combustion engine. Gradually, those sound characteristics will become a distant memory, further and further away from our communal consciousness. When the reference point finally recedes into the past, what's going to happen to the sound of a car?

Whitecar: You can envision the sounds of Tron and the Jetsons. The automotive industry is very good at tying back to legacy. Tying back to retro designs, making cars look like something, or sound like something iconic and revered. That's a very common approach in the industry to reinvigorate brands. But you also have big companies that break the mold and try something new. At the end of the day, you have to go back to the customer consciousness, where the buyer is asking themselves, "how far out there do I want to be?" Some car buyers love being on the edge, but the majority of people just want their car to look like a car and sound like a car, whatever that means in their frame of reference.

Ernst: There's the safety element too. As the manufacturers are coming up with what this exterior sound is going to be, they'll be bound by legislation that regulates tone, frequency, loudness, and more. OEMs will have to meet and incorporate those requirements into whatever profile they come up with.

Levenson: Where do you think we are in terms of the maturity of these external sounds? Do you think we're in the early stages of experimentation or do you think that techniques, technology, and regulation have progressed toward some degree of standardization?

Whitecar: With very few exceptions such as the Taycan from Porsche, it is still very early. Some of what you hear on the street now, because of our collective, predetermined notion of what a car sounds like, seems like something completely different, like something unidentifiable might be rolling down the road.

Ernst: I think the OEMs are still exploring what they want to do. They're not certain of what consumer acceptance is yet so they're treading lightly. The OEMs at the cutting-edge are trying ideas on one car, such as the Porsche Taycan, not the whole fleet. Those OEMs are leading the market when it comes to experimenting with what this total synthesized sound experience could be inside and out of the car. No doubt, they're learning a lot from their customers, and that will put them in a good position as they go to proliferate the sound profile throughout their EV model lineup.

Whitecar: In my previous work at Ford, we were looking at AVAS before it was even a requirement. We did a bunch of studies where we would conduct test drives in urban and suburban settings to understand what people actually notice. These were behavioral observation studies looking at pedestrian awareness. Basically, does this pedestrian know that a car is coming? Does the car sound like a car to this person? Does the sound of the car capture their attention or is the sound lost in the overall city ambiance? That was way before any legislation was passed addressing what was really required for AVAS to be effective from a safety perspective.

Levenson: Do you remember some of the results of those studies?

Whitecar: What we found is if it was mechanical sounding, people would pay attention. If it was musical, like a harmonic or choral effect, people tended to ignore it. If the vehicle sounded clearly mechanical, then people registered that this is something requiring their immediate attention.

Ernst: It's 100 years of cognitive conditioning that the industry will have to contend with and re-familiarize people to a new, recognizable engine sound profile. The other reason why we're still at the nascent stage of this is because the hardware aspect isn't worked out yet. We have many OEM customers asking questions like, "What type of products can I use to generate these sounds?" "Do I have to use a traditional transducer?" "What size enclosure gets me you know what level of performance?" "What frequency range?" As an industry, these questions are still in flight. We haven't yet arrived at the standardized, industrialized hardware answers.

Currently, sound synthesis playback uses the traditional ways of reproducing music, which is not the long-term solution. We all know it. Perhaps, some of these newer ideas, like these new actuators that can leverage exterior body panels to reproduce sound, could be one of the paths forward.

Whitecar: Think about trying to create an exhaust note at 40Hz with sufficient power, right? You need to move a lot of air, like a subwoofer.

Ernst: That's a lot of enclosure. It's a lot of space, and a lot of SPL is needed to make that downward firing speaker heard at the pedestrian level. Plus, with placement underneath the car, that transducer would need to be protected. There are a lot of hardware challenges to work out.

Whitecar: You almost need to drive a tuned pipe right, which is what you're trying to get rid of in the first place! You could put an organ-type pipe under the car, but now you've added a lot of hardware, a lot of cost. And remember, manufacturers like Tesla don't want noise in the cabin. The whole premise for Tesla is the quiet electric car. So if they're trying to create noise on the outside, they don't want it to come inside, and that's a hard design and engineering problem. It's very hard to isolate low-frequency sound.

Levenson: What about the related need to cancel some noise out, such as road noise?

Whitecar: Getting rid of the gas engine made other noises much more apparent. Road noise, wind noise, and mechanical noises. Tire noises in particular are really problematic. Tires have a particular tonal quality to them, especially at highway speeds. Fundamentally, that's not desired in the interior space independent of whether you want to add noise to it or not. So together with advancement in sound design, the evolution of the sound experience means there's going to be much more adoption of noise cancellation tech just because you've lost that masking quality of the IC engine.

Levenson: Do the OEMs consider how sound pairs and connects with visual and touch feedback systems?

Whitecar: There's definitely the consideration of how everything needs to work in concert. Something as simple as lane detection: the mirror blinks, the audio alert triggers, and you might have a haptic effect with a steering wheel vibration. So, OEMs have to combine all those things so that they're intuitive to the user, and combined, the event elicits an appropriate reaction. It's a system design problem in ergonomics, human perception, and behavior.

Levenson: Is sound spatialization becoming a critical element?

Ernst: There's a lot of interest in it, but still forward-looking at this point. It's that next level of audio reproduction that OEM's are looking at because spatialization provides that third dimension of context that will help the driver. So in the lane changing scenario, if it's an alert that there's a vehicle in your blind spot, you get the benefit of having the acoustic alert coming from the direction of the oncoming vehicle. You don't have to take your eyes off the road to know the location of the hazard. OEMs also need to consider what's natural or what's comfortable to the user, what adds context, but doesn't distract or annoy. It's a narrow path for the designers to actually provide benefit.

Levenson: As automotive sound design matures, will brands continue to pursue individualized sound profiles or will the industry standardize around a generic set of sounds?

Ernst: In the next handful of years, we're going to see a lot of different interpretations. Long term, the market will sort out what works and what doesn't, and we'll see a convergence. For the soundscape, the overall orchestrated mix of sounds, that will continue to be specific and identifiable for brands.

Whitecar: It's going to be an iconic brand differentiator, just like the grille was in the past. You could tell a Pontiac on the road vs. a Cadillac just by looking at the grille. It's an immediate brand identifier. Going way back, there was a spec on GM cars called Click Clack. In the radio, the audio team had to create the turn signal sound based on the original mechanical relay sound. We had to create a tick-tock, tick-tock sound and it was iconic to GM. Click Clack was in every GM spec.

Levenson: There's a high level of familiarity with these sounds and something different could create a sense of dissonance, of discomfort. What happens when the sound design is poorly executed? Do owners notice and complain?

Whitecar: Yes, the OEMs certainly hear about it. The good thing about the whole software-defined vehicle approach is that the OEM can just flash an update and fix the issues. That's the advantage of the new paradigm: If somebody dislikes the standard sound set you can offer them other tested and validated choices. Or the OEM can listen and respond to feedback and send out an OTA update and fix the problem.

Levenson: Is the best sound design the experience you don't really notice, or do consumers get excited about a good soundscape?

Ernst: There will be a positive reaction to it and I think the OEMs are playing to that. OEMs are having lots of conversations to design that very first sonic interaction between customer and vehicle. As I approach the car, is it giving me a welcome sound acknowledging that I'm there. Manufacturers are looking at all of those experiences. For example, when you enter, you sit in the driver's seat, and you close your door. Then there's the welcome to the car sound. The OEMs are carefully thinking through and designing that customer experience from first approach, to driving, to exiting the vehicle. They know that customers derive satisfaction from these interactions.

Levenson: The design and synthesis of sounds is obviously complex. What are some of the most difficult technical issues and how are they best solved?

Whitecar: First, the question is how to cost-effectively create a synthesizer. If you get down to basics, that's what we're doing: we are creating a synthesizer. So how do you do that without spending $3,000 USD per car? And how do you make it so that it can scale across low-end cars all the way up to super premium electric cars. So that leads to the question of how to get signal processing down to minimal footprint, and minimal CPU.

The other challenge is how to achieve ease of development. The design team is asking how to get the concepts to run on a safe embedded platform. Sound designers work in powerful Digital Workstation environments. They've got endless plug-ins, they've got a great sound monitoring system, they've got controllers, and multiple displays. Sound designers have a great environment for creating sounds, but then there's the challenges of implementation. They have to apply the sounds in real-time simulations based on vehicle data triggers including speed, transmission, doors opening, and more. Finally, all of that needs to translate and move over to an embedded platform.

So you go from endless memory and endless CPU in the DAW environment, but in order to make it cost effective, you have to fit all of the designs onto a small MCU or DSP or it has to fit inside the main entertainment system. Car makers have to have a scalable tool and design environment that allows them to take the design running from a PC to the target hardware platforms.

That ability to translate from the creative world to the physical world and then to the embedded world is the trick.

Levenson: And what's the solution?

Whitecar: Honestly, Audio Weaver is the best solution today because of the flexibility of going across platforms as well as running on a PC. There are some tools that run on an embedded platform, but don't have a good simulation environment, and there are good tools that work in the simulation environment that don't translate to embedded. Audio Weaver is one of the few tools that crosses those boundaries. Today, the DSP Concepts team is working on mechanisms to actually make that coupling tighter: making it more seamless to work with DAWs and the output from DAWs.

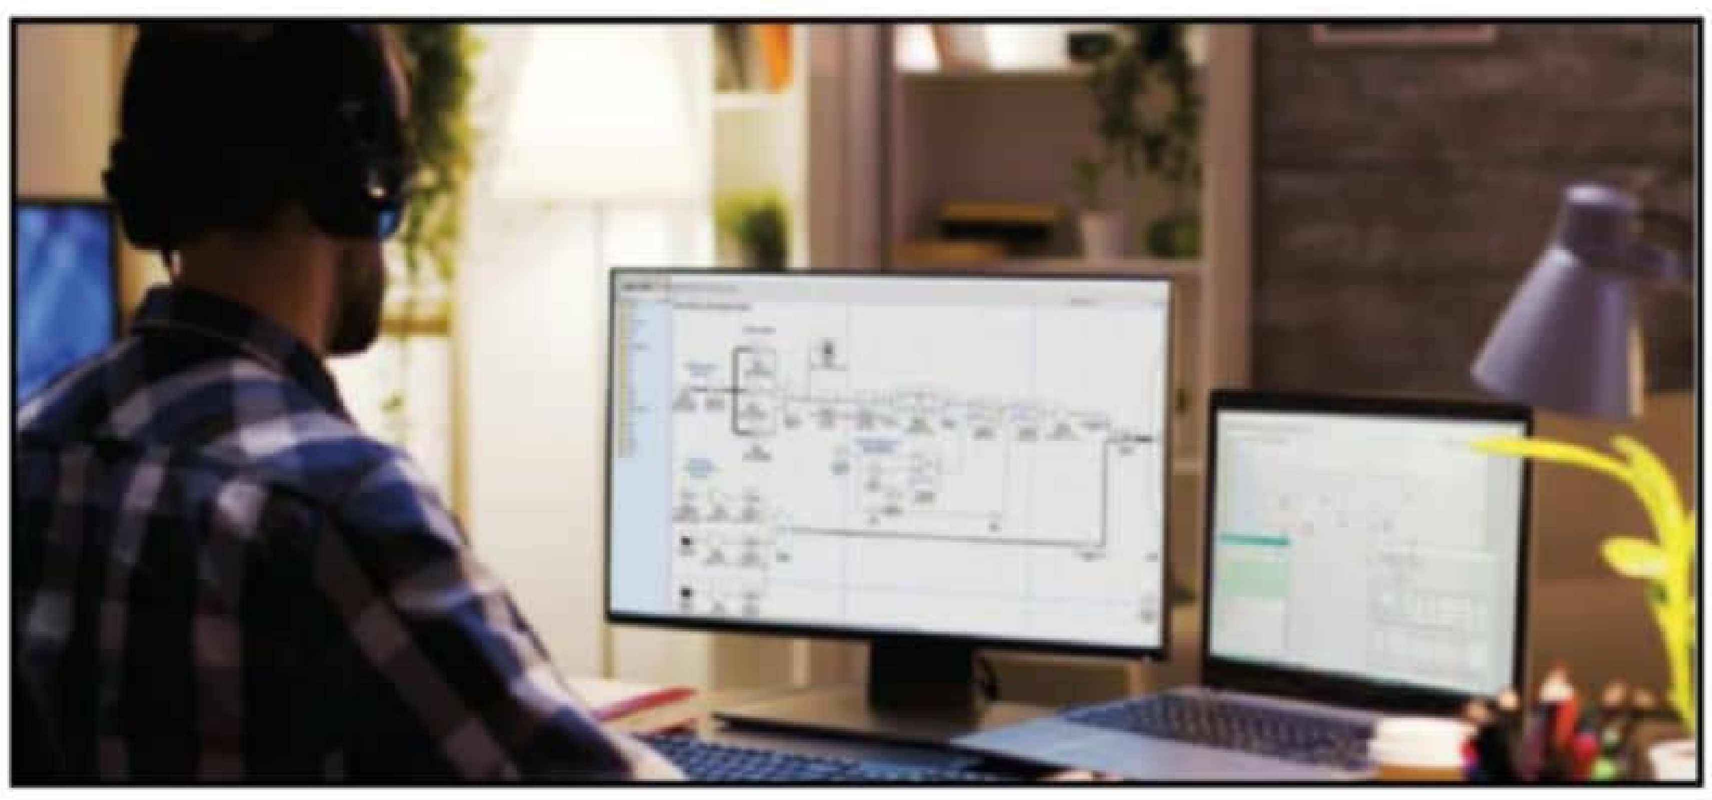

Levenson: What's the future? What do you see coming in the next decade, especially as autonomous becomes more prominent?

Whitecar: My wildest prediction is that the vehicle will become more in tune with the passenger/traditional driver. The experience will become more of a dialogue and less of an alert. For example, "Hey, John, I need you to take control of the car right now if you don't mind." We've been presented with this type of future concept for decades. When infotainment was introduced in the car, the promise was to make life easier. We're going to give you navigation so you don't have to hold the map. We're going to add a voice assistant so that the car knows where you like to get your coffee. The industry keeps taking small steps toward this future vision.

Today, processing power is super strong and getting even stronger—exponentially stronger every year. Plus, with 5G there's going to be off-board processing that is instantaneous. That promise of a relationship that you can have with your vehicle is going to follow. I see a much more interactive relationship between the vehicle and the owner/passenger/driver compared to what we have today. When we hit autonomous at Level 4/Level 5, and that's not happening tomorrow, but someday when we get there, the vehicle is going to be more like an extension of the owner.

Levenson: Are we going to need some powerful AI?

Whitecar: Yes, AI is going to drive that. AI is going to apply contextual awareness, look at the situation in real-time, and give you feedback. For example, "The weather's getting pretty nasty out there: we're going to have to try a different route." The conversational aspect of that experience is the key rather than just firing off beeps and alerts. You don't have to read the screen to receive the guidance, the AI just has a conversation with you.

Levenson: It would be a lot safer, frankly.

Whitecar: Very much safer, because today, you still have to do a lot of interpretation. If you think about how navigation works now, I can say, "Take me to the Home Depot" and the voice assistant will respond with a list of 10 options. It's going to be much more helpful when the AI can tell you which one is closest and just ask if you want to go there. These concepts are not new, it's just the ability to fully realize the concepts, that is what has changed.

Levenson: And to pull it all together into a coherent experience, that's also new.

Whitecar: The ergonomics experts and the psychology experts are working to understand the human experience and psyche to figure out what works with human beings: what experiences are natural and which ones are distracting. Since we're constantly being studied and observed, Machine Learning is closing that gap. They know where you're looking. They can tell if the road is bumpy. They can tell if the weather is good. They can tell if you're talking to somebody. All of this awareness builds the intelligence so that the vehicle is just interacting with you almost as another being.

Another challenge for OEMs is to introduce these new experiences in a way that acclimates the end-user so we can come along for the ride. Otherwise, these changes could be jarring. If you try to introduce too much of this at once, people will disengage. They won't appreciate it, but if the experience evolves gradually and naturally, then people will just go along for the ride, literally.

Ernst: And following up on that idea of intuitive experiences, even as it relates to what's being done right now with soundscapes and 3D sounds that help contextualize things. Those sounds still have to be intuitive. It can't just be a noise, it's got to make sense, has to be instantly understandable. That's something that is only going to come over time—time and experience.

Whitecar: Owners don't want to read a manual that explains what the sounds indicate and what function they correspond to. To reach that level of intuitive responsiveness, manufacturers need to work out the design concepts, and gradually develop a sonic language that is instantly understandable. Achieving that goal requires that special combination of artistic skill to create the sounds and a technology platform to effectively implement and orchestrate the soundscape.

About the Author

Adam Levenson's 30-year audio career has evolved from sound effects design, to business development, to marketing with roles as Director of Audio at Activision, GM at CRI Middleware, VP of Product Marketing at Waves, and now Head of Marketing at DSP Concepts.

DSP Module Application and Tuning Procedure in Automotive Audio

By
Philipp Paul Klose

This article offers a primer on how to approach acoustic measurements and the practical application of DSP in car cabins, one of the hardest acoustic environments for quality sound reproduction. A continuation of the article focused on microphone techniques published in *audioXpress* March 2023, this article details how to make use of the collected measurement data.

Tuning audio systems requires a structured and strategic approach in the first calibration phase. Following the signal flow order of modules from input to output or picking random parameters to tweak is not a constructive line of action. One general approach that will give reliable results is described here.

For visualization, I used the tuning graphical user interface from the automotive DSP/amplifier combo miniDSP Harmony 8x12 [1]. The take-aways from this article can be used for DSPs from other manufacturers too and the learnings are not limited to automotive but are applicable to other audio systems and product development.

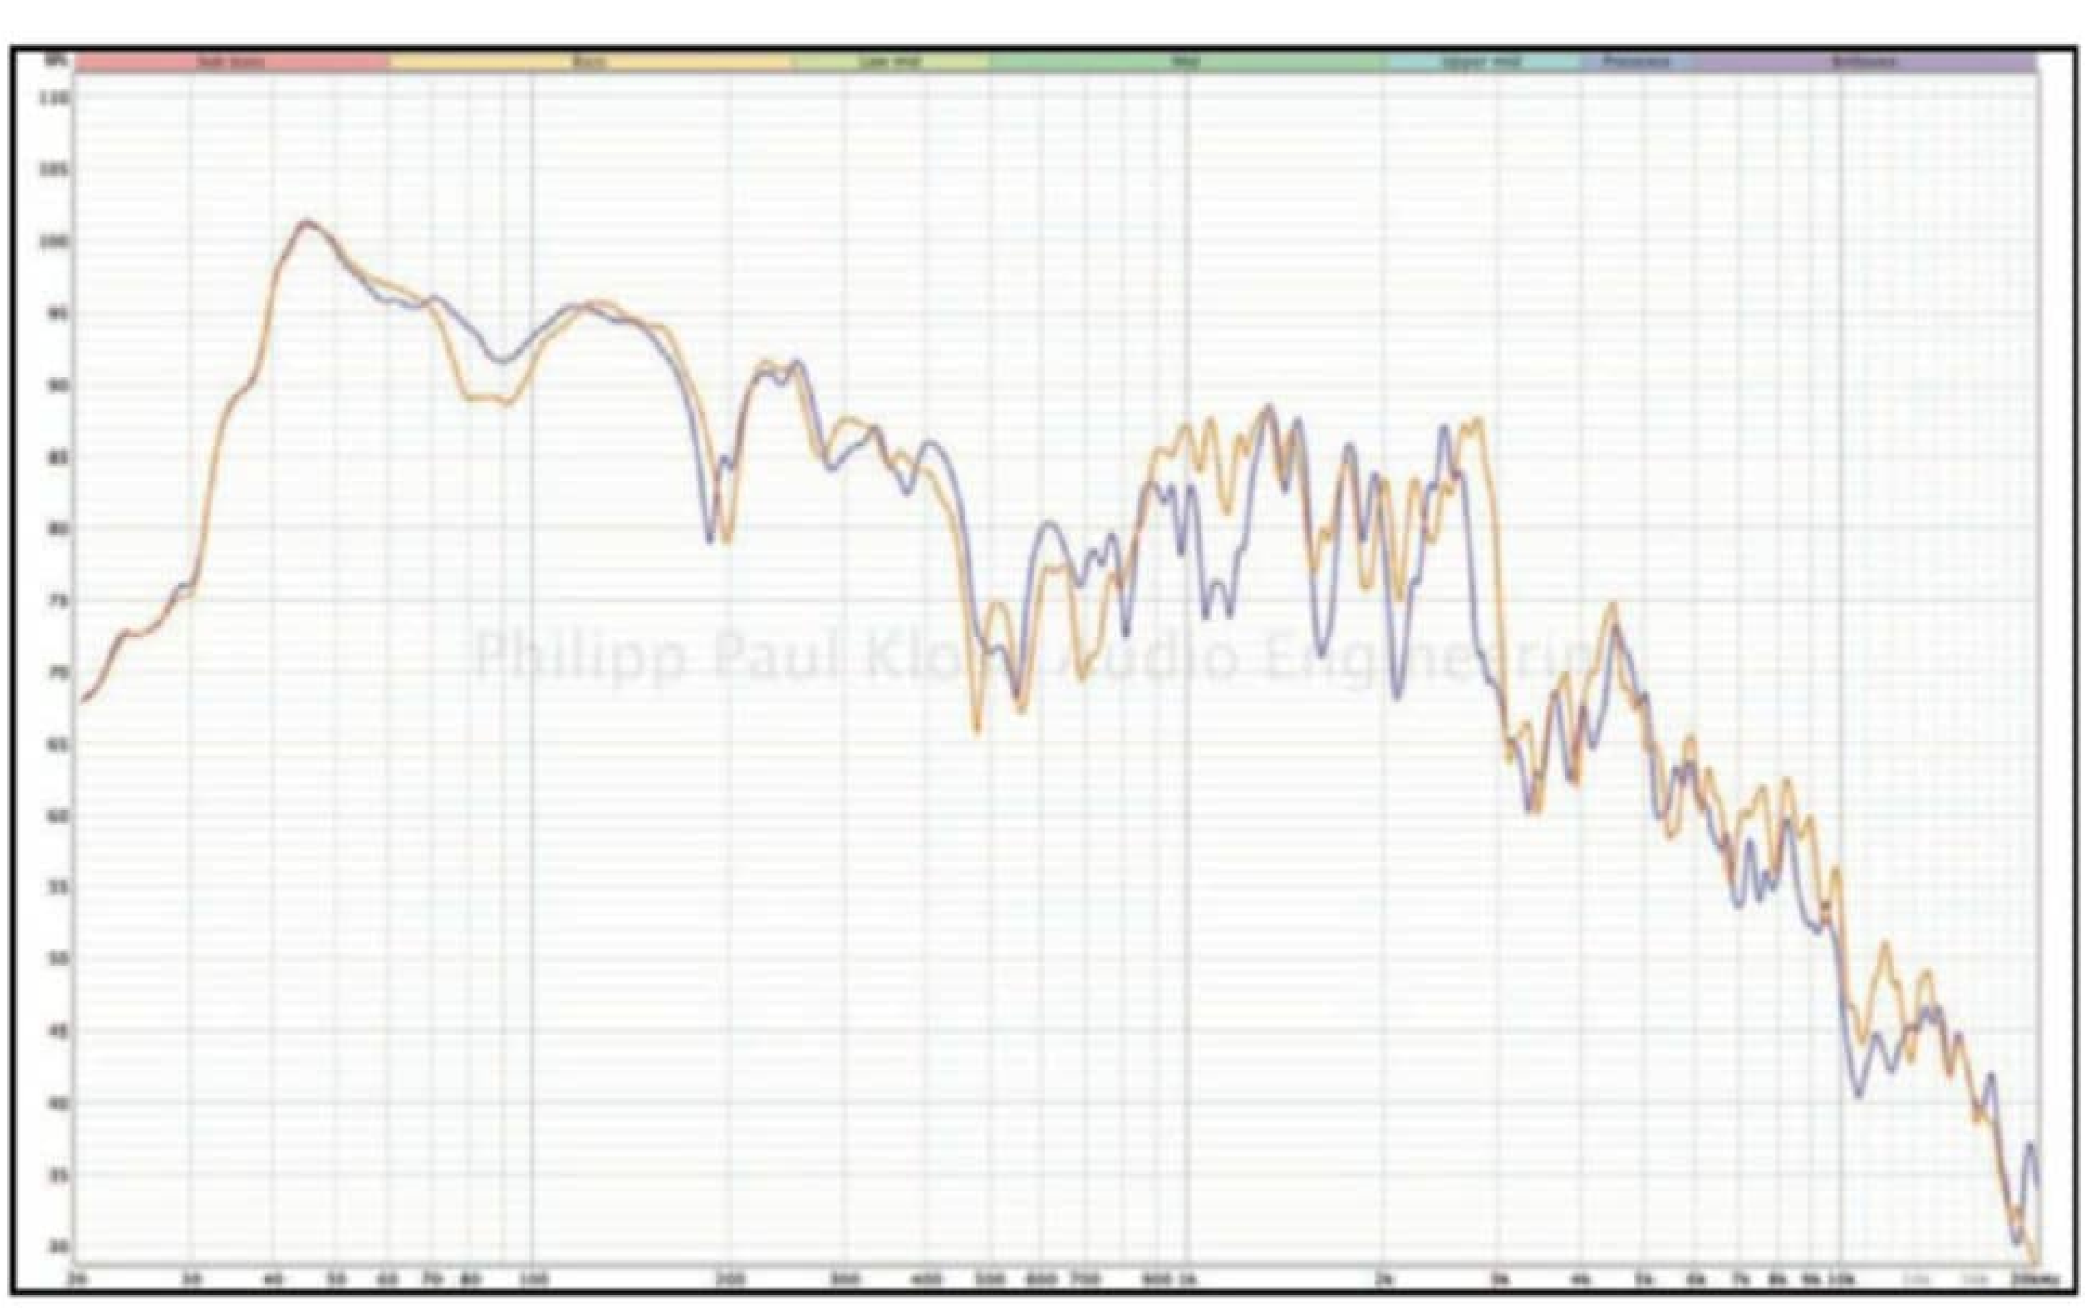

Figure 1: Sensitivity check of left and right door woofers with a single microphone in the middle of the car's front row. Both speakers match nicely in their sensitivity.

Sensitivity and Polarity Check

After the successful installation of speakers, DSP, amplifiers, and cabling in the car plus the setup of the microphones and the analysis system, you need to check the sensitivity and polarity of the speakers.

To check the sensitivity, use a single microphone in the middle of the car's front row. Although the single microphone technique is not the best for setting up EQs and filters, it is useful for the initial setup process.

My article in the March 2023 issue of *audioXpress* discusses more about microphone techniques. Measure every speaker with the same output voltage of your amplifier (somewhere between 1V to 3V) with a log sweep (**Figure 1**). Compare them in pairings that make sense. That can be door woofer left vs. door woofer right, tweeter left vs. tweeter right, etc. The frequency responses and levels should be similar in the comparison. Aberrations of 1dBr to 2dBr are considered normal. This way you can rule out problems in your installation (e.g., acoustical short circuits or faulty wirings). Do not delete the measurement data from this step, it will come in handy later.

In the next step, the sensitivity of loudspeakers used for the playback of a logical channel should be checked (**Figure 2**). If a woofer, midrange, and tweeter is used for the playback of the left channel of a stereo signal, then these speakers must work together appropriately.

The next step will check for the correct polarity of the system (**Figure 3**). The impulse response or step response of loudspeaker measurements will be used. Note that when you are using an OEM head unit to feed your system, the processed signal of the head unit could cause deviations and unclear results in the measurements [2]. In some situations, it can be necessary to make an additional measurement just a few centimeters from the speaker away to get a distinct result.

Speaker Protection Through Output Level and Crossover

Speakers can break when used in the wrong frequency band and/or handed too much power. This happens either through mechanical or thermal stress. To be safe from the get-go and during the calibration phase (where overload mistakes can easily happen), speaker protection should be addressed as early as possible during the tuning.

Here is where the datasheet of your speakers comes in handy. Check for information about the frequency band where your speakers can be used and maximum power ratings. Measure the output voltage of the amplifier and calculate the power delivery to your speaker.

An RMS multimeter can be used to measure the voltage when your system is turned up to maximum and all other signal processing (EQ, crossover, etc.) is turned off. For example, you have connected a 4R midrange speaker that is capable of handling 20W to the miniDSP Harmony 8x12. Since the amp can drive up to 40W into 4R, the output must be set to -6dBFS to halve the output power effectively (**Figure 4**).

If your DSP has a limiter or compressor—a compressor can be turned into a limiter when the ratio turned up to 20:1 or higher—you can add a second line of defense for the speaker through voltage limitation.

In the car it is helpful to have a look at the overlap band. The overlap band can be found through individual measurements of the speaker drivers with the same voltage. Then, compare them in an overlay view. Shift the overlap band to a point where all speakers are in their comfort zone (**Figure 5**).

It is best practice to start with 24dB/octave Linkwitz-Riley (LR) in automotive due to the good summation function. If the crossed speakers have a narrow overlap band, a steeper filter should be used, a wider overlap band can be managed with a shallower filter. When both speakers are crossed,

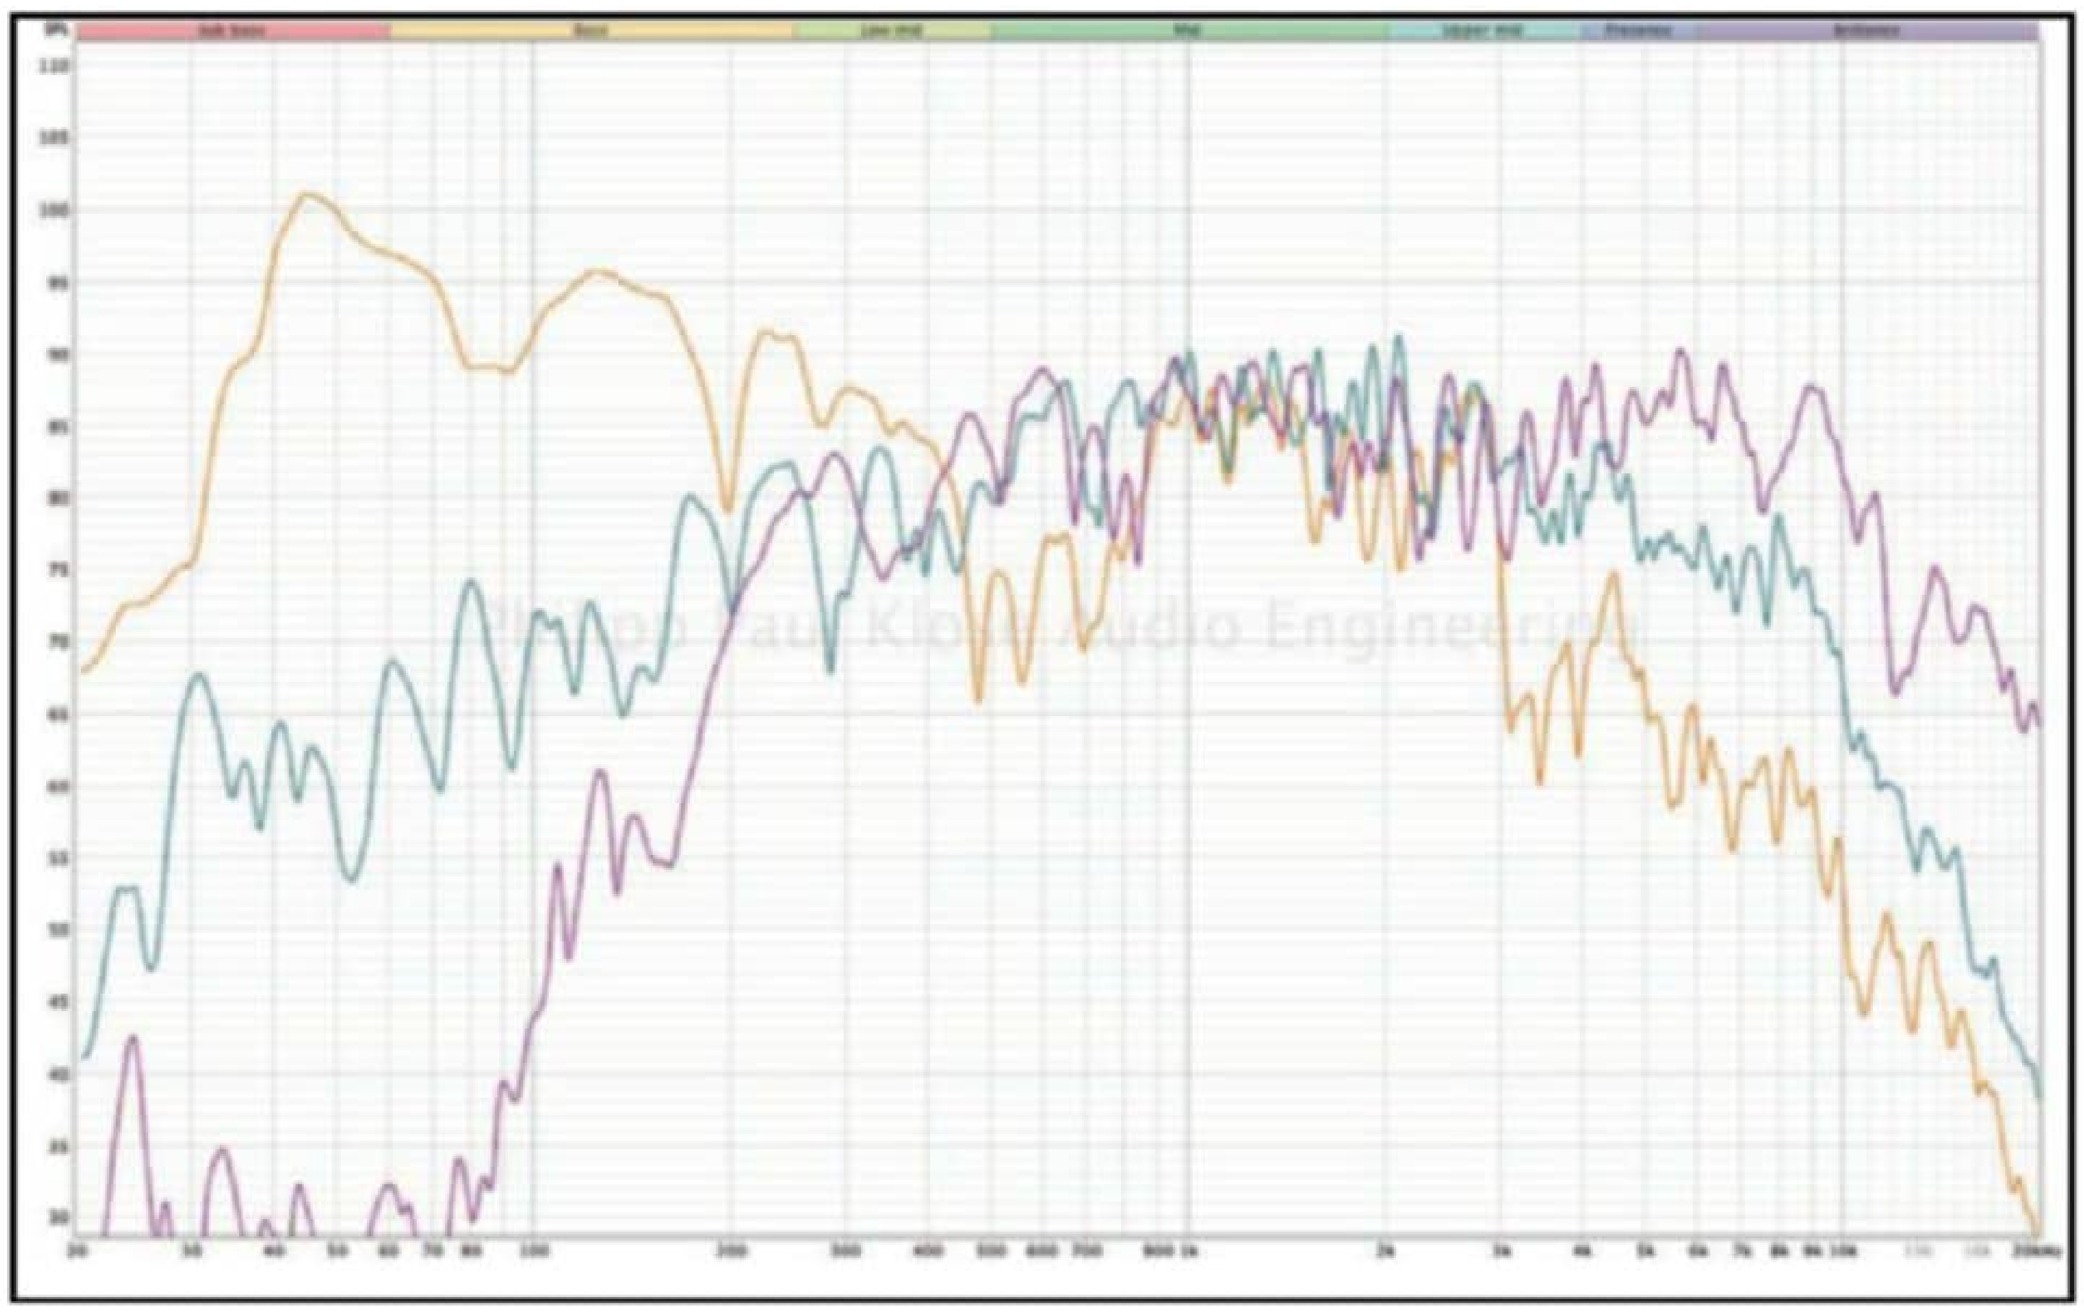

Figure 2: Sensitivity check of the single ways of a three-way system for a logical channel (woofer, midrange, and tweeter). The sensitivity matches at the tested 2.83V and the wide overlap of the speakers gives plenty of options in the tuning process.

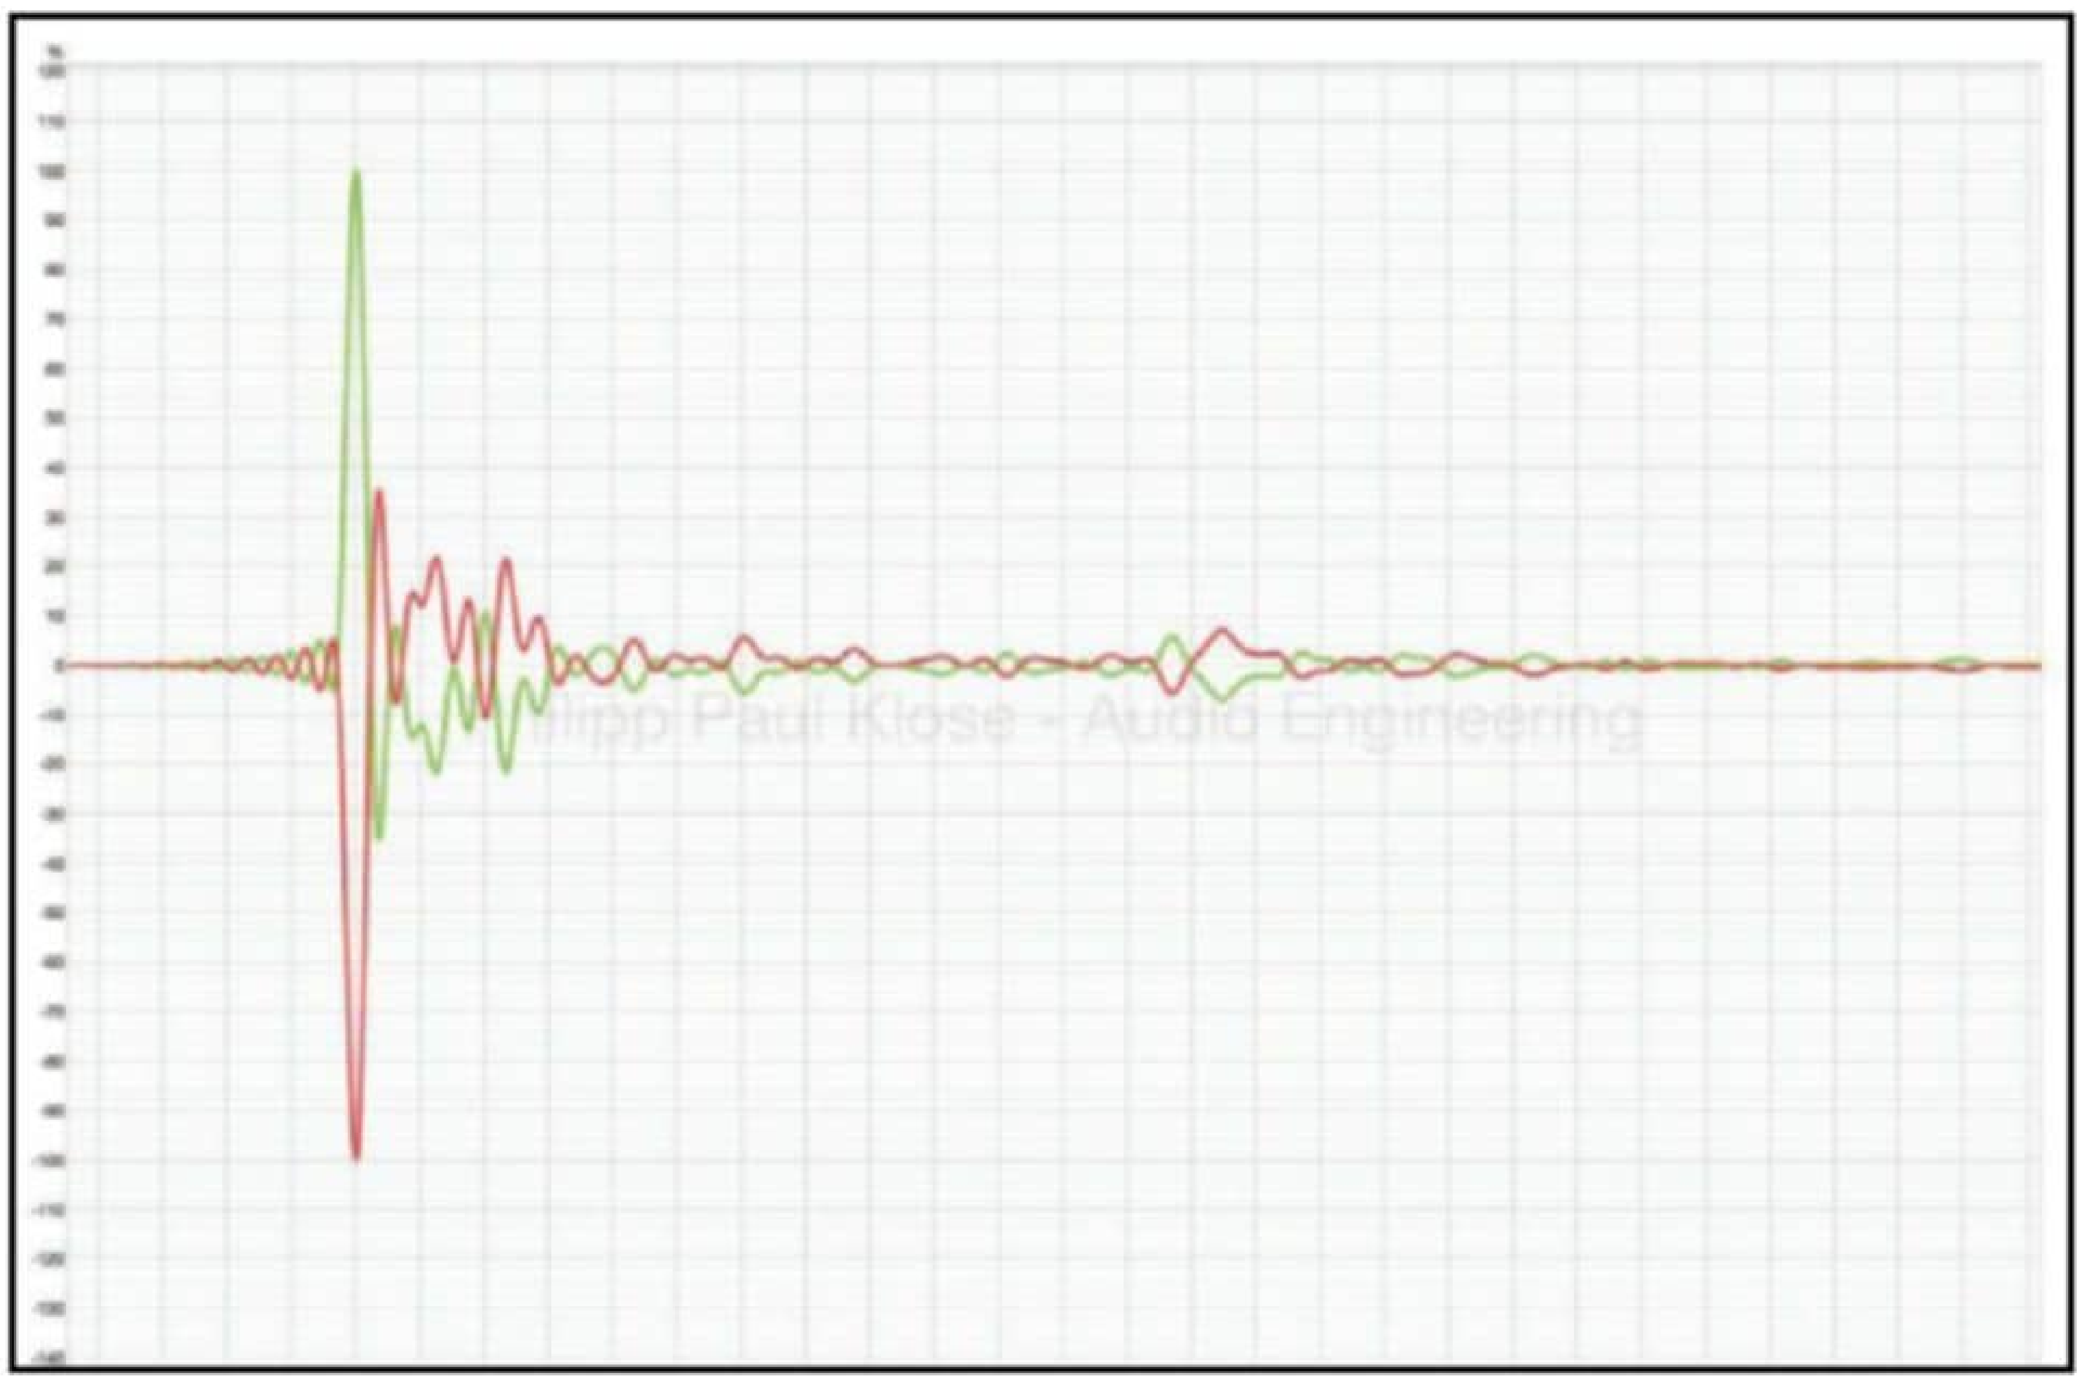

Figure 3: Polarity check through an impulse response from a log sweep measurement: Positive (correct and expected) polarity in green, vs. negative (incorrect and unexpected) polarity in a speaker setup.

Resources

P. P. Klose, “Acoustical Measurement and DSP Amplification in Automotive Audio,” *audioXpress*, March 2023.

References

[1] miniDSP, www.minidsp.com

[2] A. Wehmeyer, “How to Test for Polarity,” audiofrog, August 2021, www.audiofrog.com/how-to-test-for-polarity

[3] A. Wehmeyer, “Crossovers: How They Work and How to Choose Them,” audiofrog, January 2021, www.audiofrog.com/crossovers-how-they-work-and-how-to-choose-them

[4] Index of Multiple Sub Optimizer, www.andyc.diy-audio-engineering.org/mso

Figure 4: Crossover setup for a midrange speaker. The frequency range of the speaker is cut off at 150Hz and 2550Hz with Linkwitz-Riley (LR) filter.

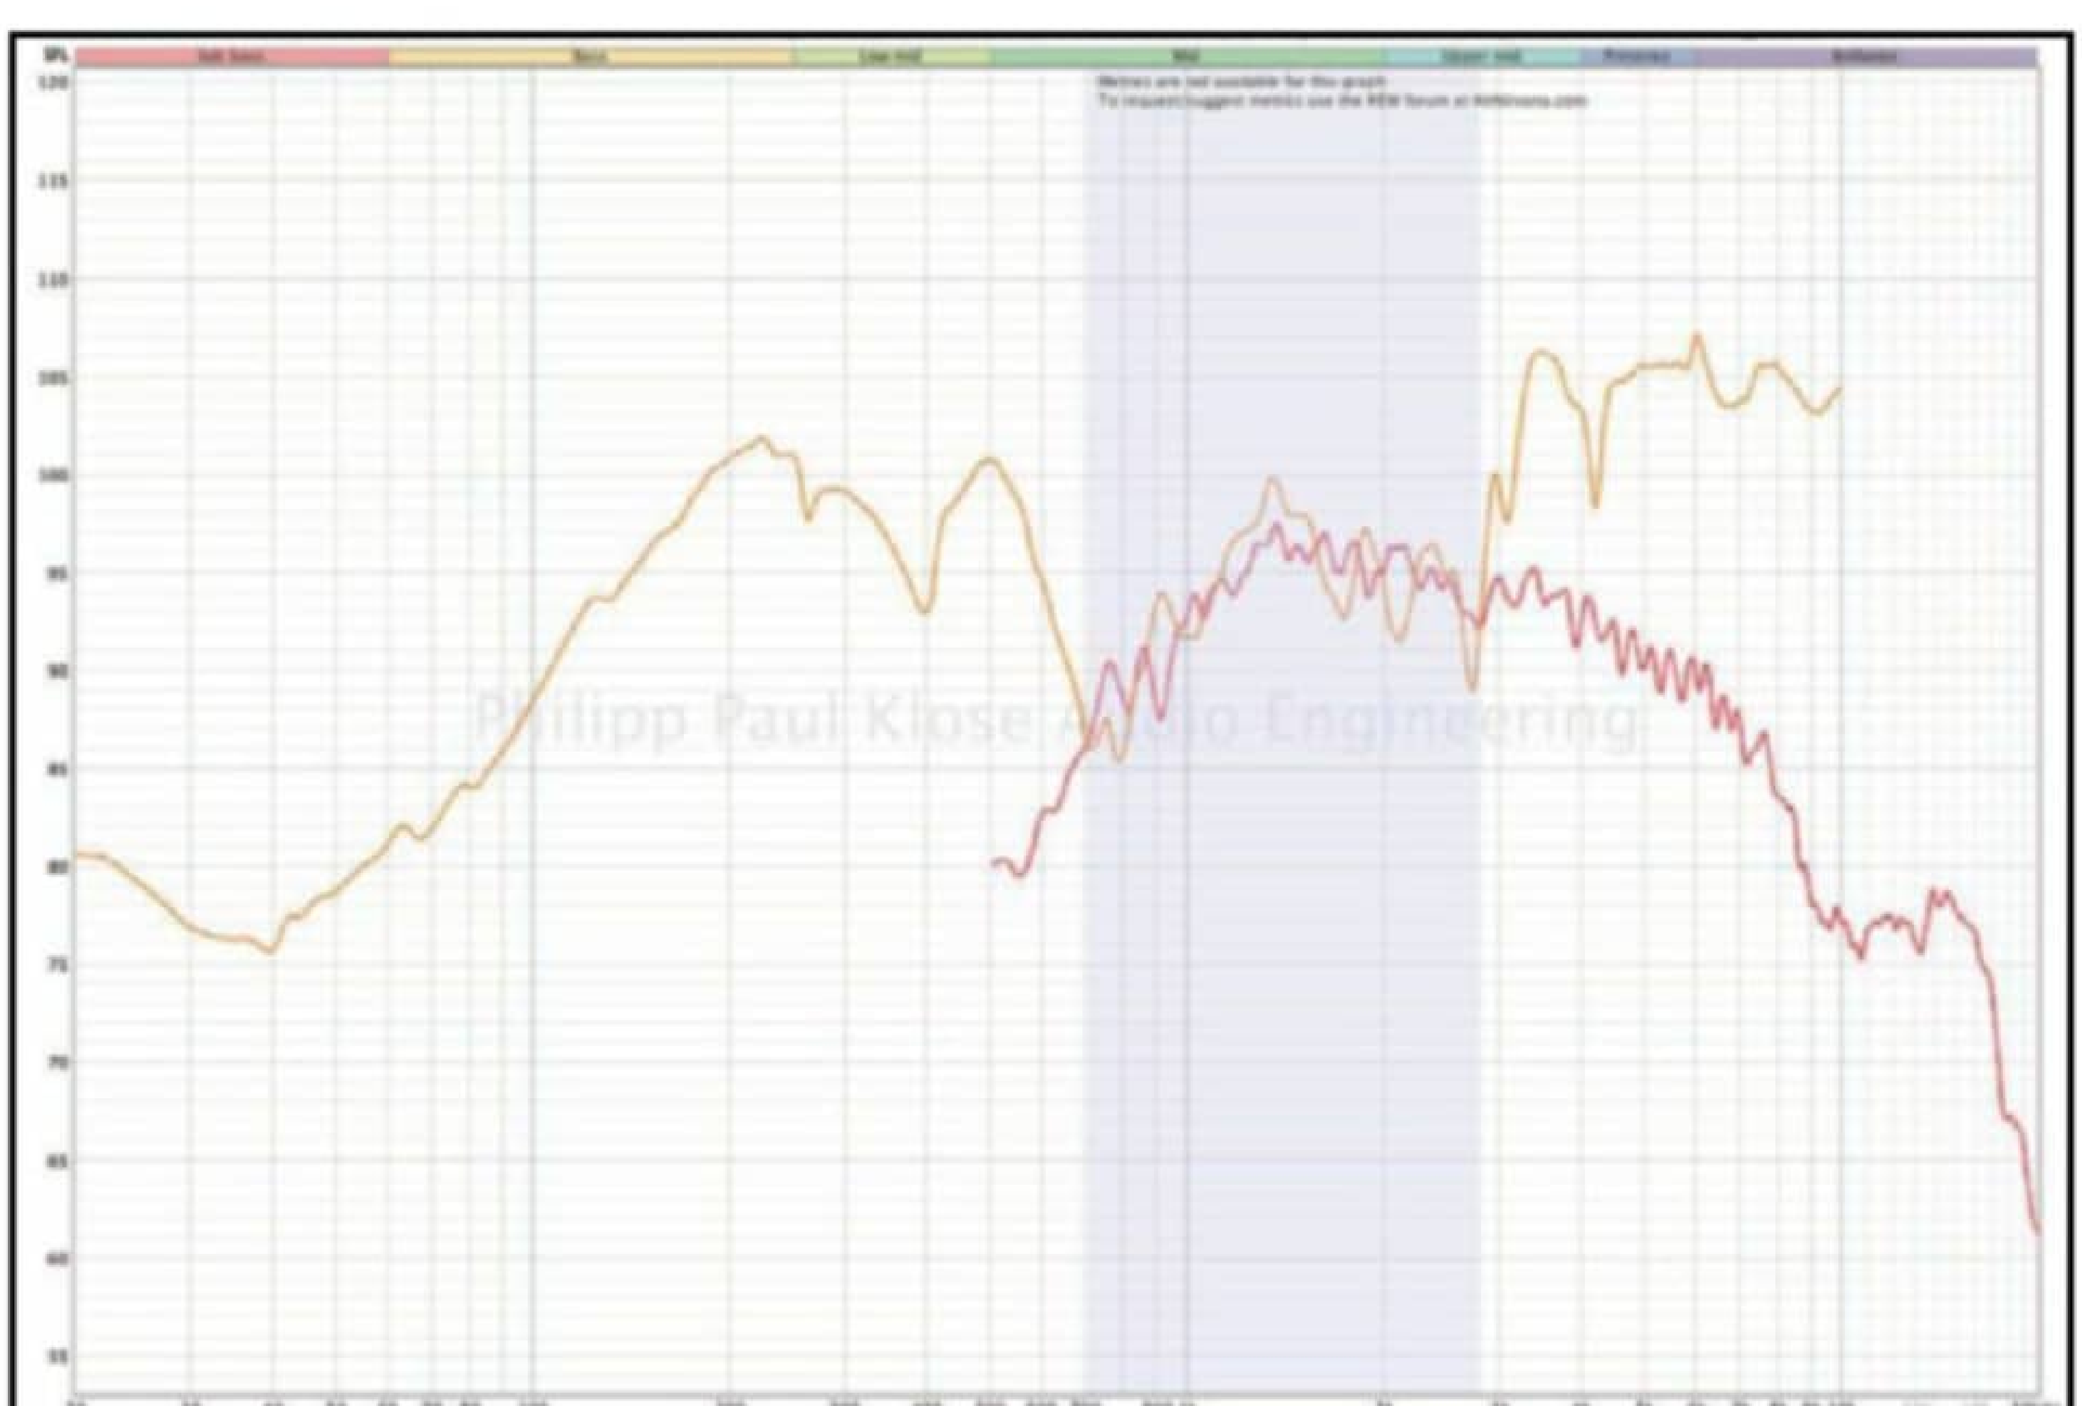

Figure 5: Overlap band measurement of the midrange and the tweeter with both speakers measured at the same voltage. In the wide band between 700Hz and 2.8kHz (colored blue) both speakers match and allow for a good crossover design.

About the Author

Philipp Paul Klose is a senior engineer for automotive audio concepts and functions at Cariad SE and previously served in positions such as audio system development engineer, project lead, and infotainment test engineer. He holds a Bachelor of Arts (Hons.) degree in Audio Production. Within his 10 years of automotive experience, he has worked for customers such as Audi, Volkswagen, Porsche, Lamborghini, and Ferrari.

their integration function is checked with both speakers active in the measurement (**Figure 6**).

Timing and Phase Correction with Delays

In the car cabin every speaker is at a different distance from the listener(s). To improve the overall impulse response, quality of the imaging and improve crossover efficiency, delays are applied.

The idea here is to find the speaker with the biggest physical distance to the listener and delay every other speaker accordingly to the same distance. The goal is to even out the arrival time of all sound events from every speaker to create one acoustic center. This will massively improve the overall impulse response of the system and imaging quality in general. Although it is stated in its module name, it is worth remembering that delay values are always negative in time for real-time DSP.

To determine the speaker delay, measure the distances from the speakers with a tape measure from the acoustic center or you use the delay data from your sensitivity measurements in the first steps (**Figure 7**).

After delays and crossovers are done, the summation function between crossed and delayed speaker pairs should be checked (**Figure 8**). Make sure that both speaker channels from the comparison pair have been calibrated in their crossover and delay settings and are playing properly. When the crossover band is measuring flat, timing and crossover frequencies are fine. If the summation band dips or comb filters are visible, cross frequencies and/or delays must be adjusted [3].

Bass Management and Woofer Mono Bass Optimization

Although car cabins are relatively small on volume (between 2.5m^3 and 4m^3), it has been proven in praxis that only the use of several woofers and subwoofers delivers an overall coherent result between seats. The load is distributed to several speakers and position dependencies can be reduced (less changes when moving the head in a seat) and the consistency between seat positions is better. This concept is called "mono bass." To get an input signal into the bass array, the bass management low-passes every incoming signal and "collects" the bass energy that way and hands a high-pass filtered signal to the other DSP modules downstream.

The optimum result can be found either through EQ'ing delaying or with algorithmic help. For computer-based help I can recommend the Multiple Sub Optimizer (MSO) tool, which is capable of iterating through several thousand DSP settings and checking for linearity of results [4]. REW also offers an auto EQ section to calculate filters and EQ to achieve a certain target. The advantage of the software is that also complex overlapping EQ (EQs "over" each other) can be calculated (**Figure 8b**).

Routing and Distribution

Routing is often realized in trivial way: Left signal to the left-hand side speakers of the car, right signal to the right-hand side speakers of the car. Routing is an important tool for creating proper staging (**Figure 9**).

The drastic placement of the speakers in the car creates often extreme acoustic imaging. Try to distribute the signals from your source over several speakers. Often a matrix mixer is offered for that purpose. This creates a more immersing experience and lets the speakers work more like a "team," than a bunch of single speakers handed over individual signals.

In general, it can be said, that the more input channels your source has (stereo: 2, 5.1: 6, Dolby Atmos: 12), the more discrete the routing to the speakers can be done to achieve the artistically intended acoustic imaging (**Figure 10**).

Equalization

An Equalizer is manipulating a signal in the frequency domain and allows for the balancing of frequency bands (**Figure 11**). Less is sometimes more...but sometimes very necessary: Although we have a lot of DSP power it does not mean we have to use every EQ band or function.

When approaching filtering, it makes sense to start with the biggest peaks or dips and validate

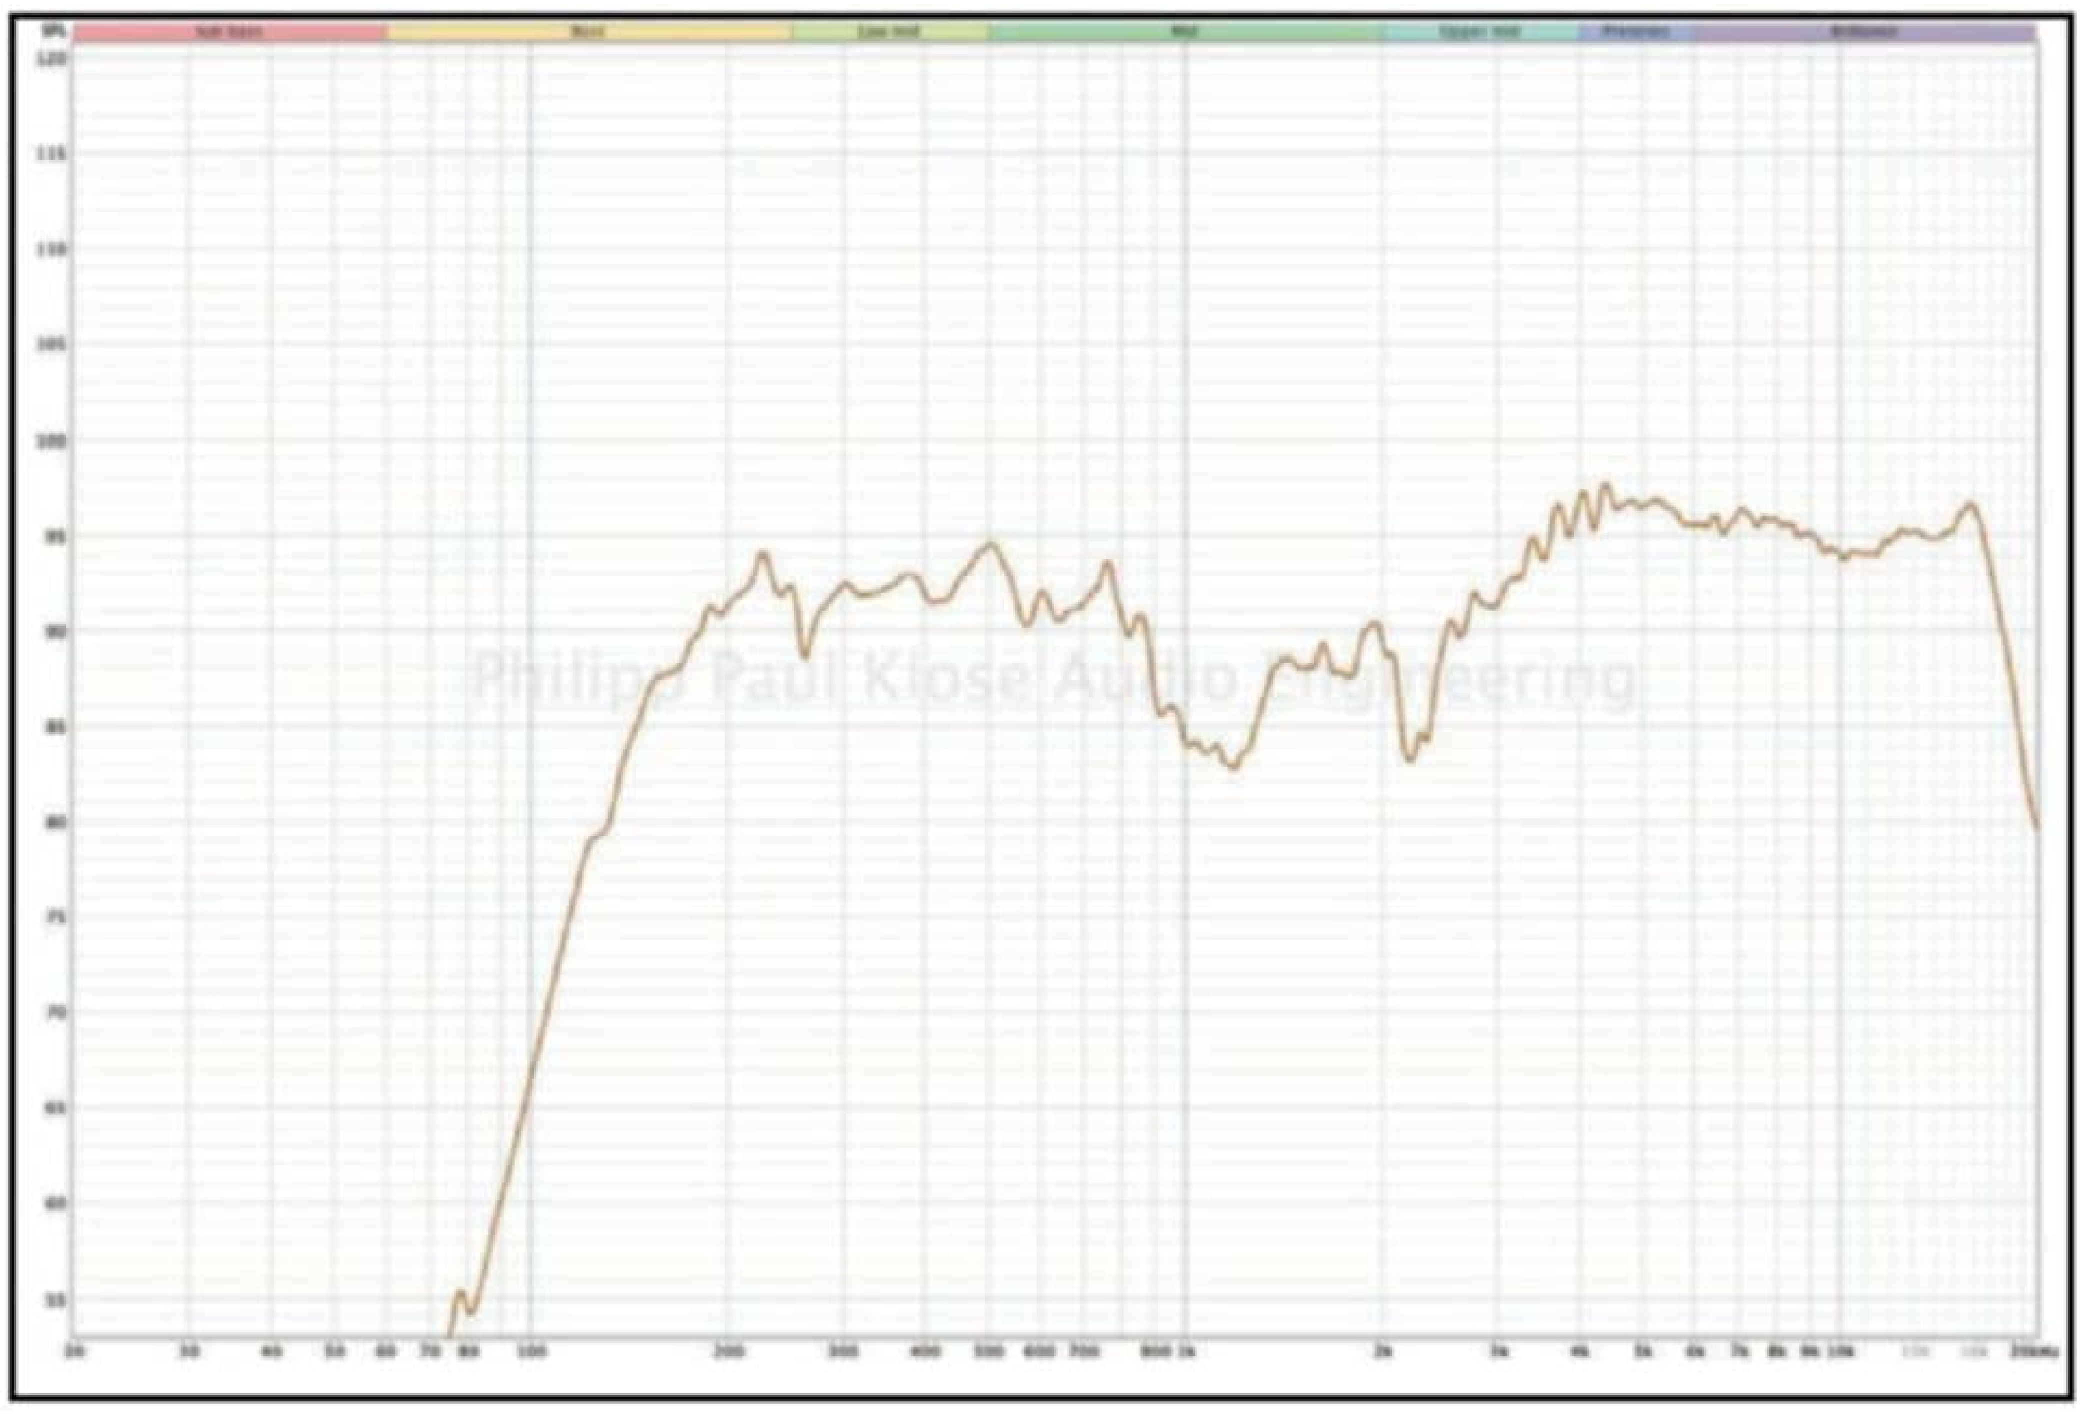

Figure 6: Crossover integration test part 1: Tweeter and midrange measured together. The dip in the crossover region at 1.5kHz is corrected when we introduce the delay in the next step.

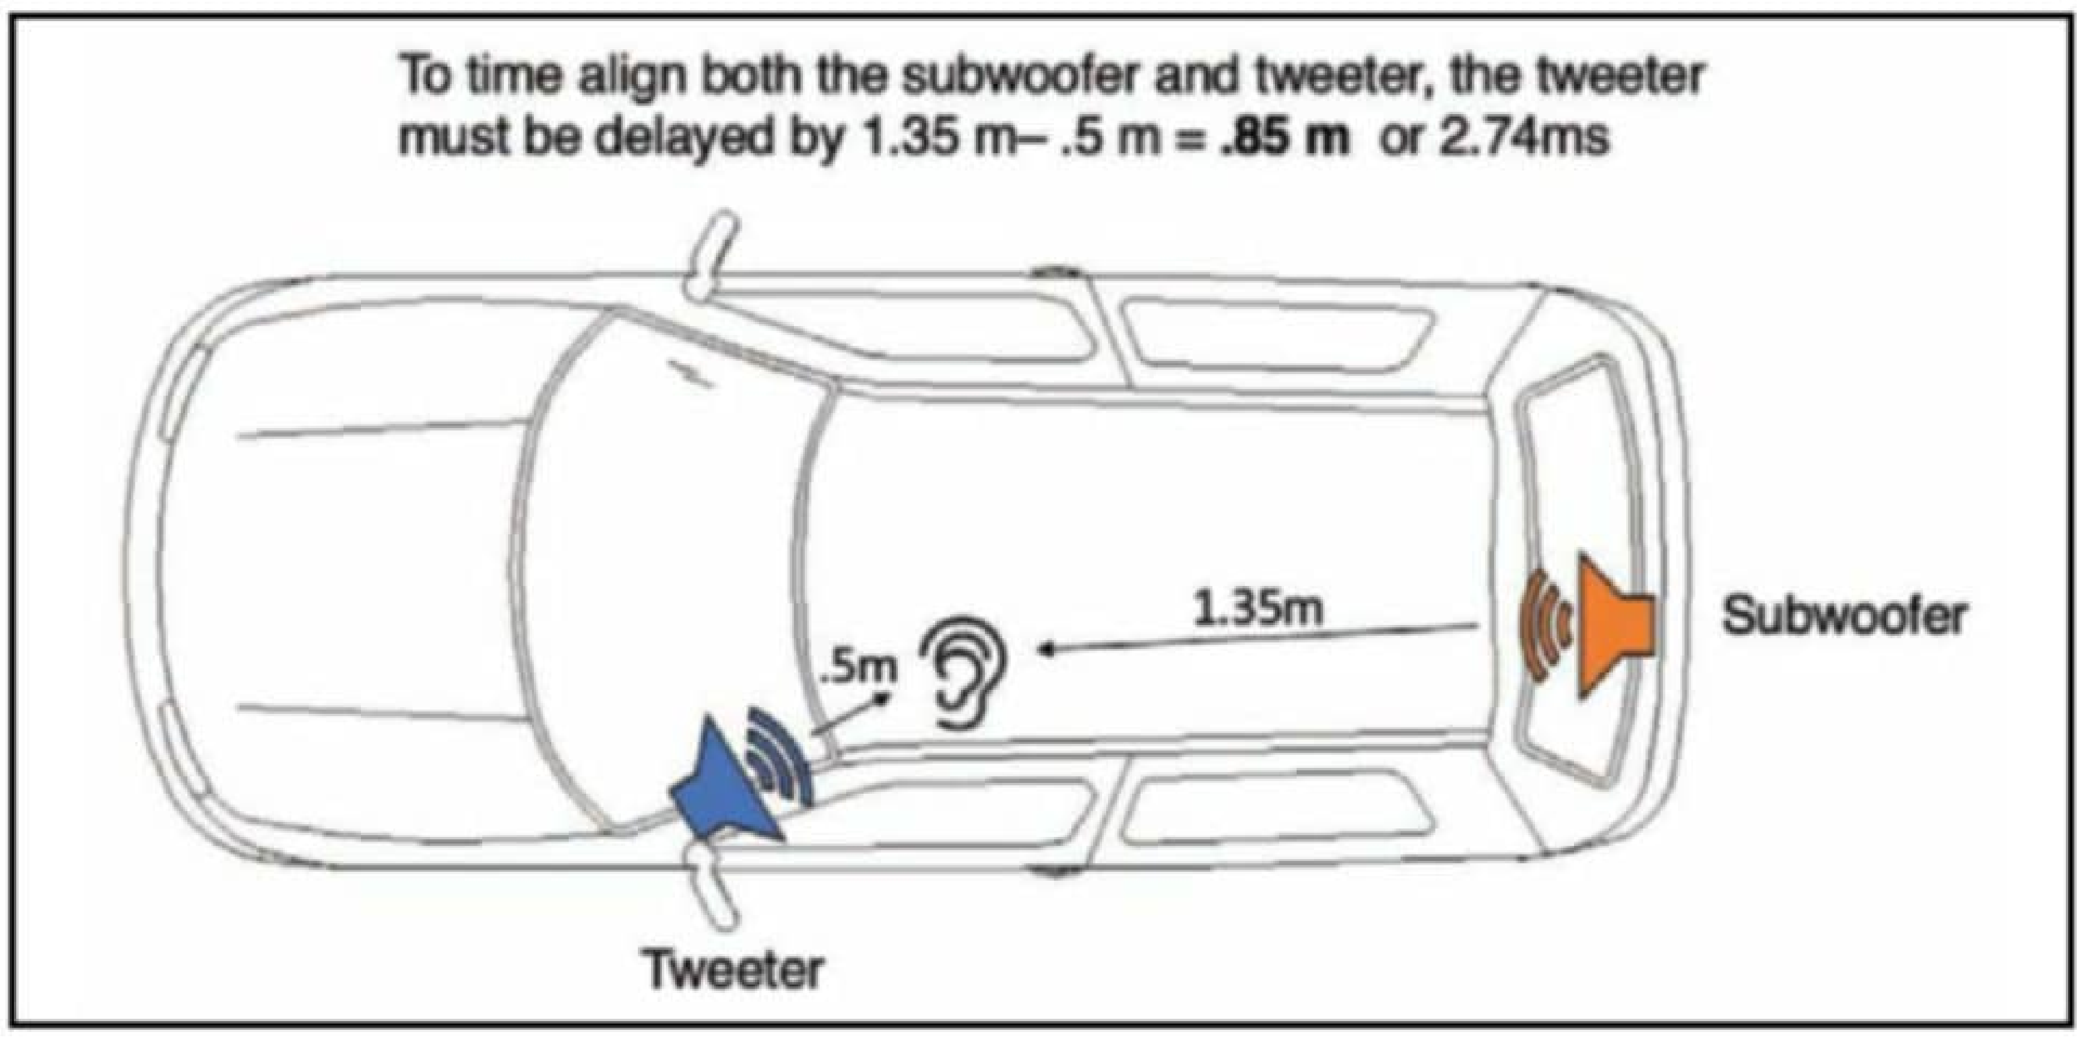

Figure 7: The acoustic center of the car as a symbolized ear. The subwoofer in the back is furthest away, the other speaker(s) must be delayed to become aligned. In this example the tweeter must be delayed to a total distance of 1.35m, achieved by adding a 2.74ms delay.

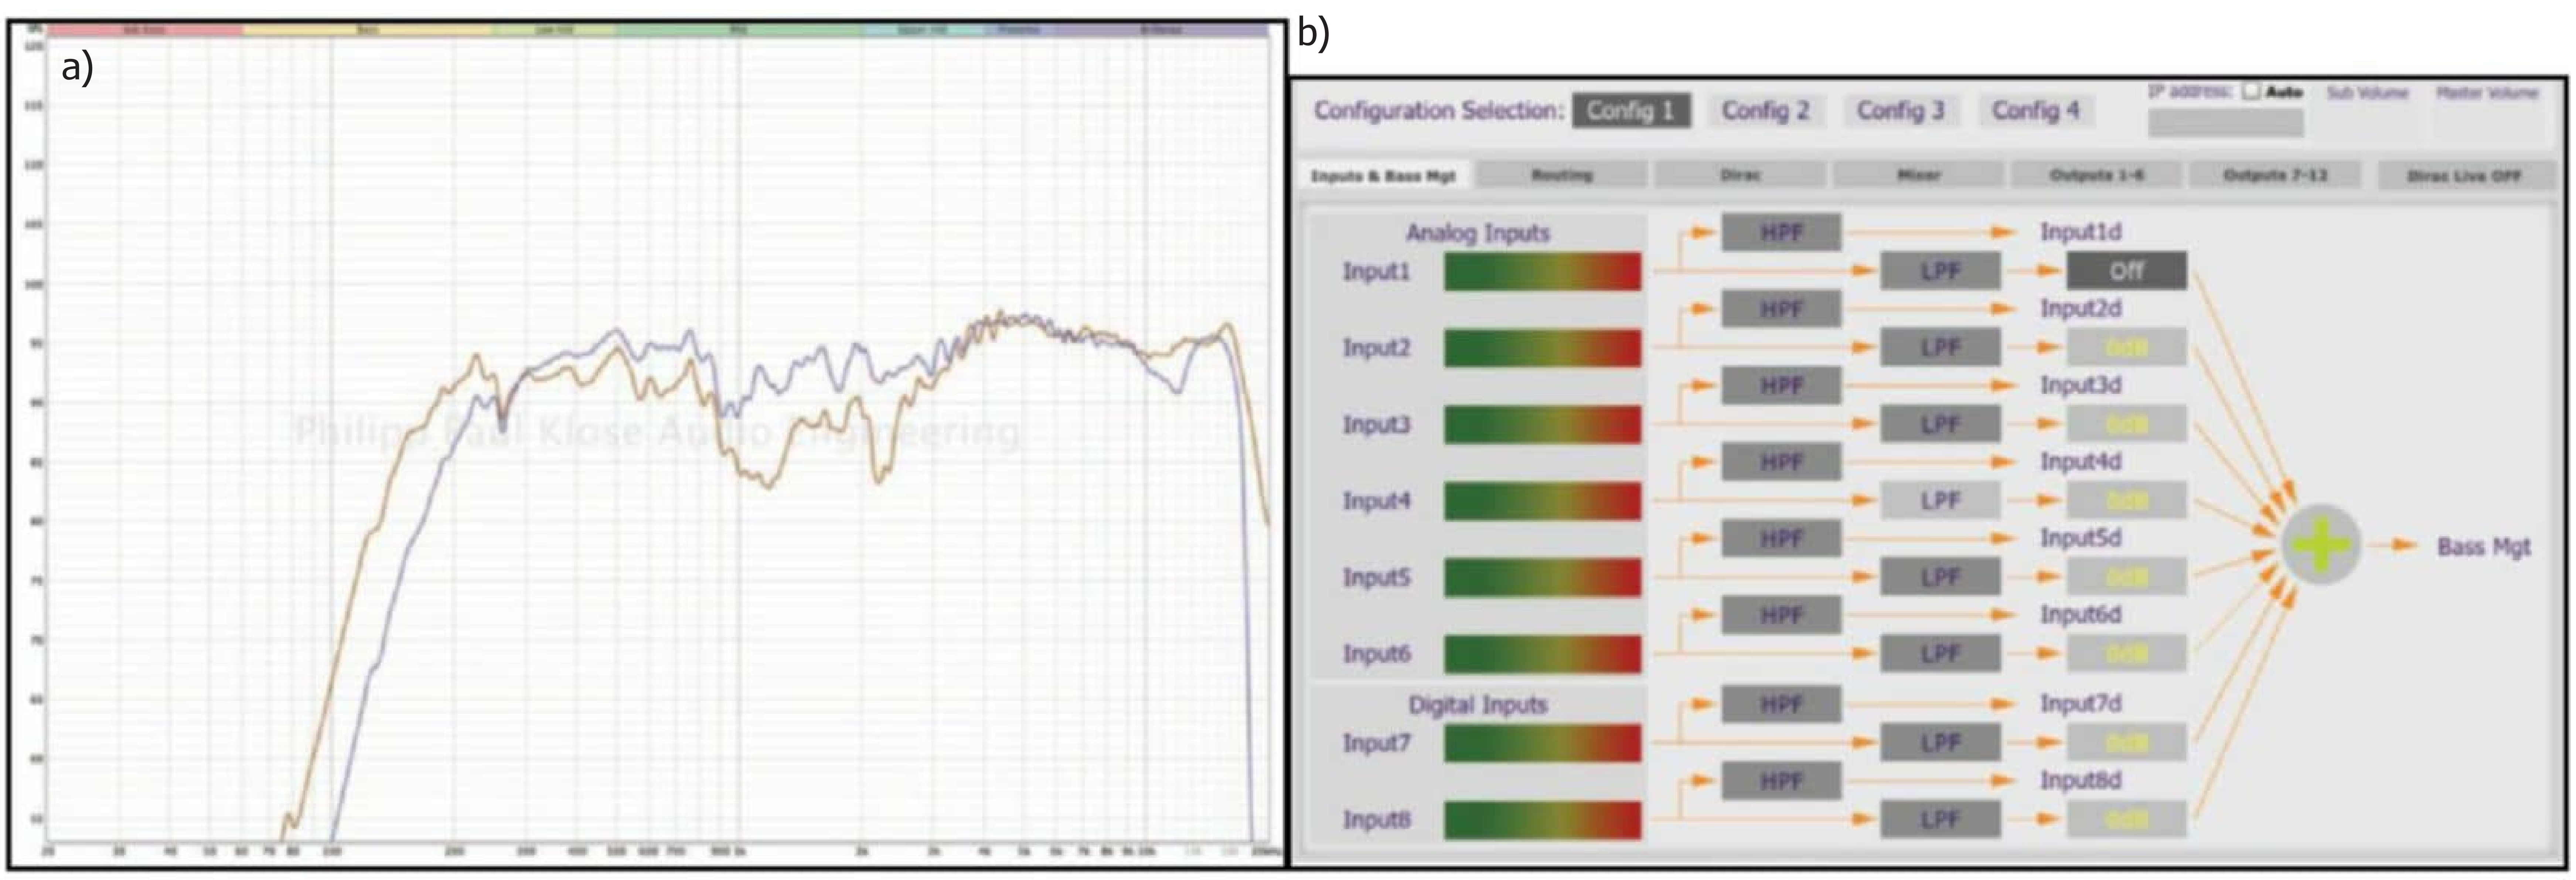

Figure 8: Crossover integration test Part 2: Without delay optimization (brown) vs. optimized crossover (blue). The delay improves the level performance in the band between 800Hz and 3kHz. Figure 8b: Bass management pipeline in the miniDSP Harmony 8x12. The signals of the incoming channels are low-pass filtered, level adjusted, and then summed

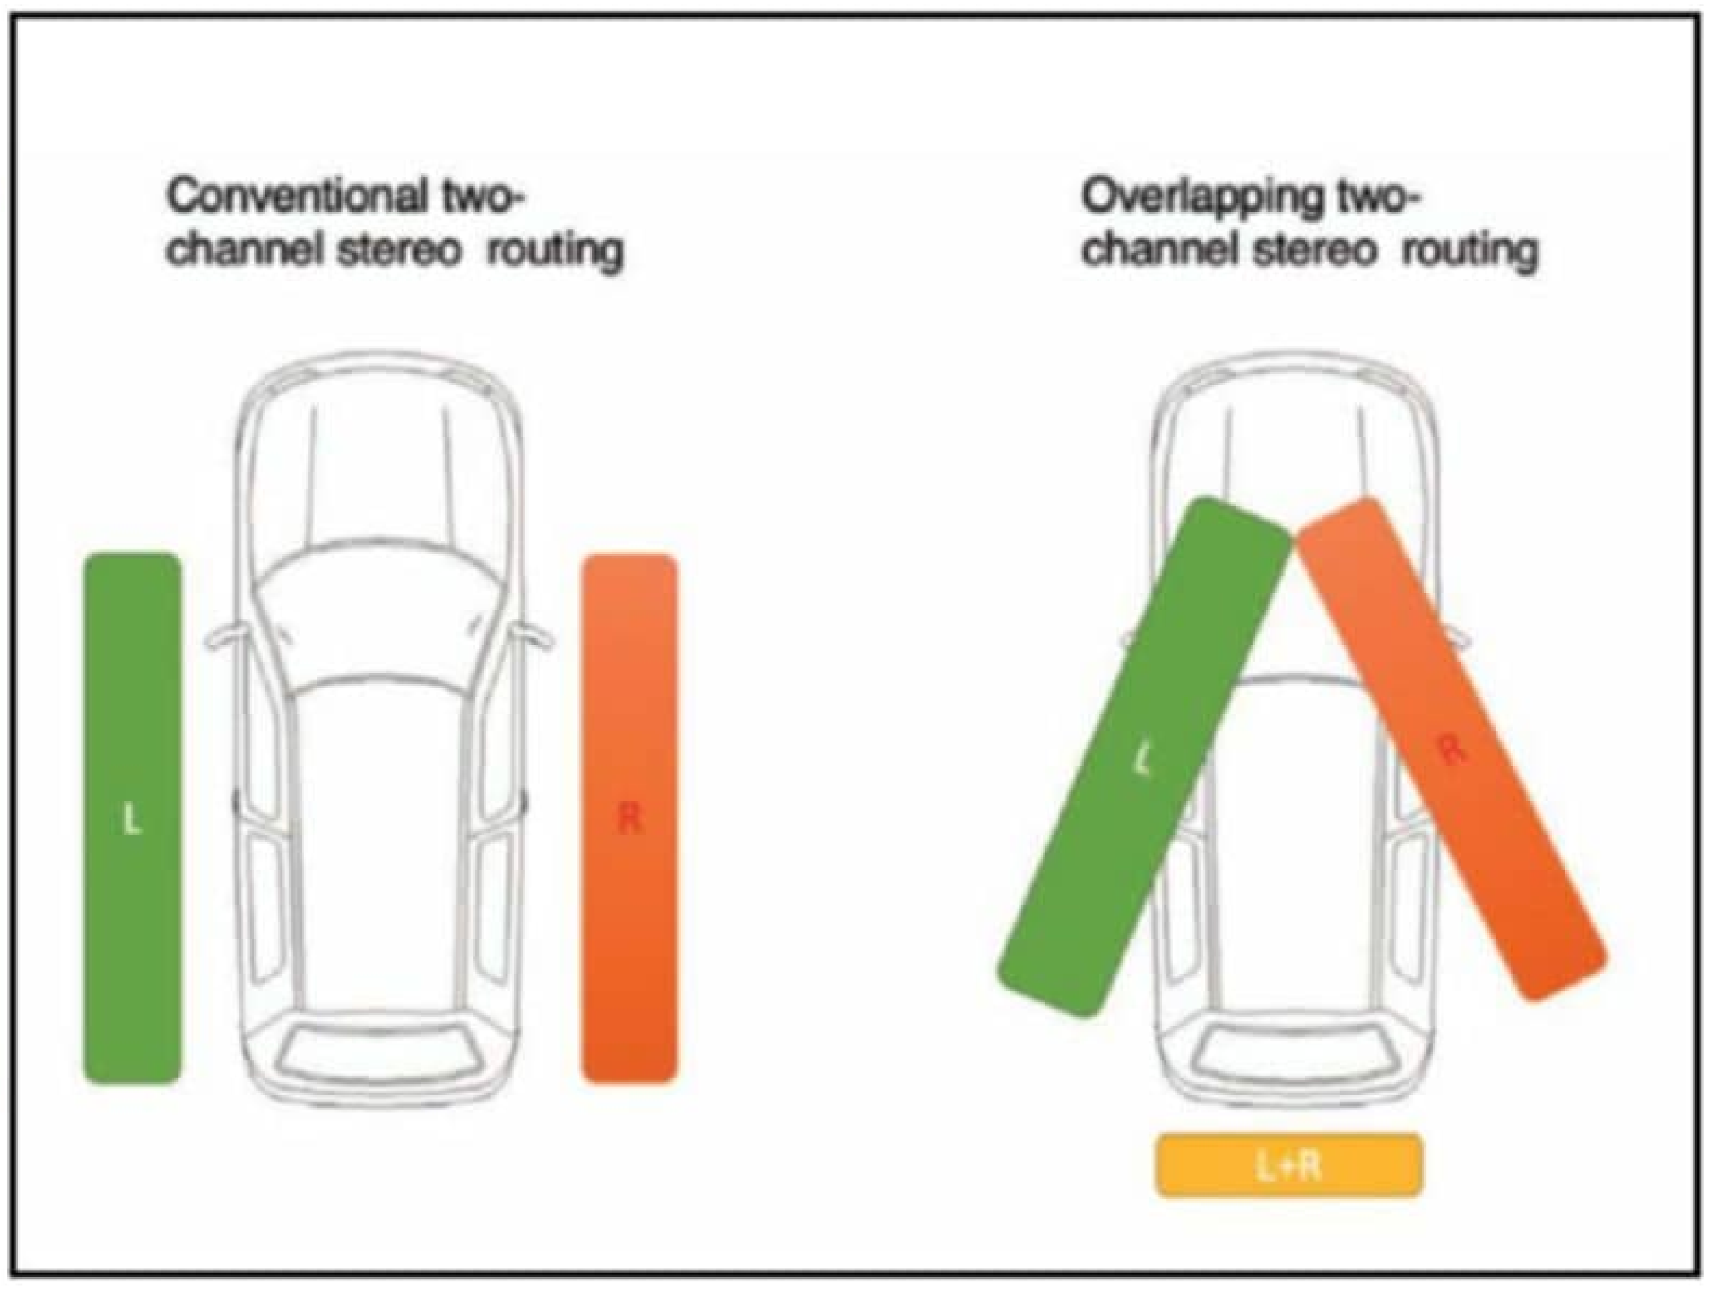

Figure 9: Routing and channel distribution approaches. The example on the left is a "hard routed" example. The example on the right uses overlapping and distribution to create an enticing acoustic atmosphere. The yellow L+R field symbolizes the bass-managed subwoofer.

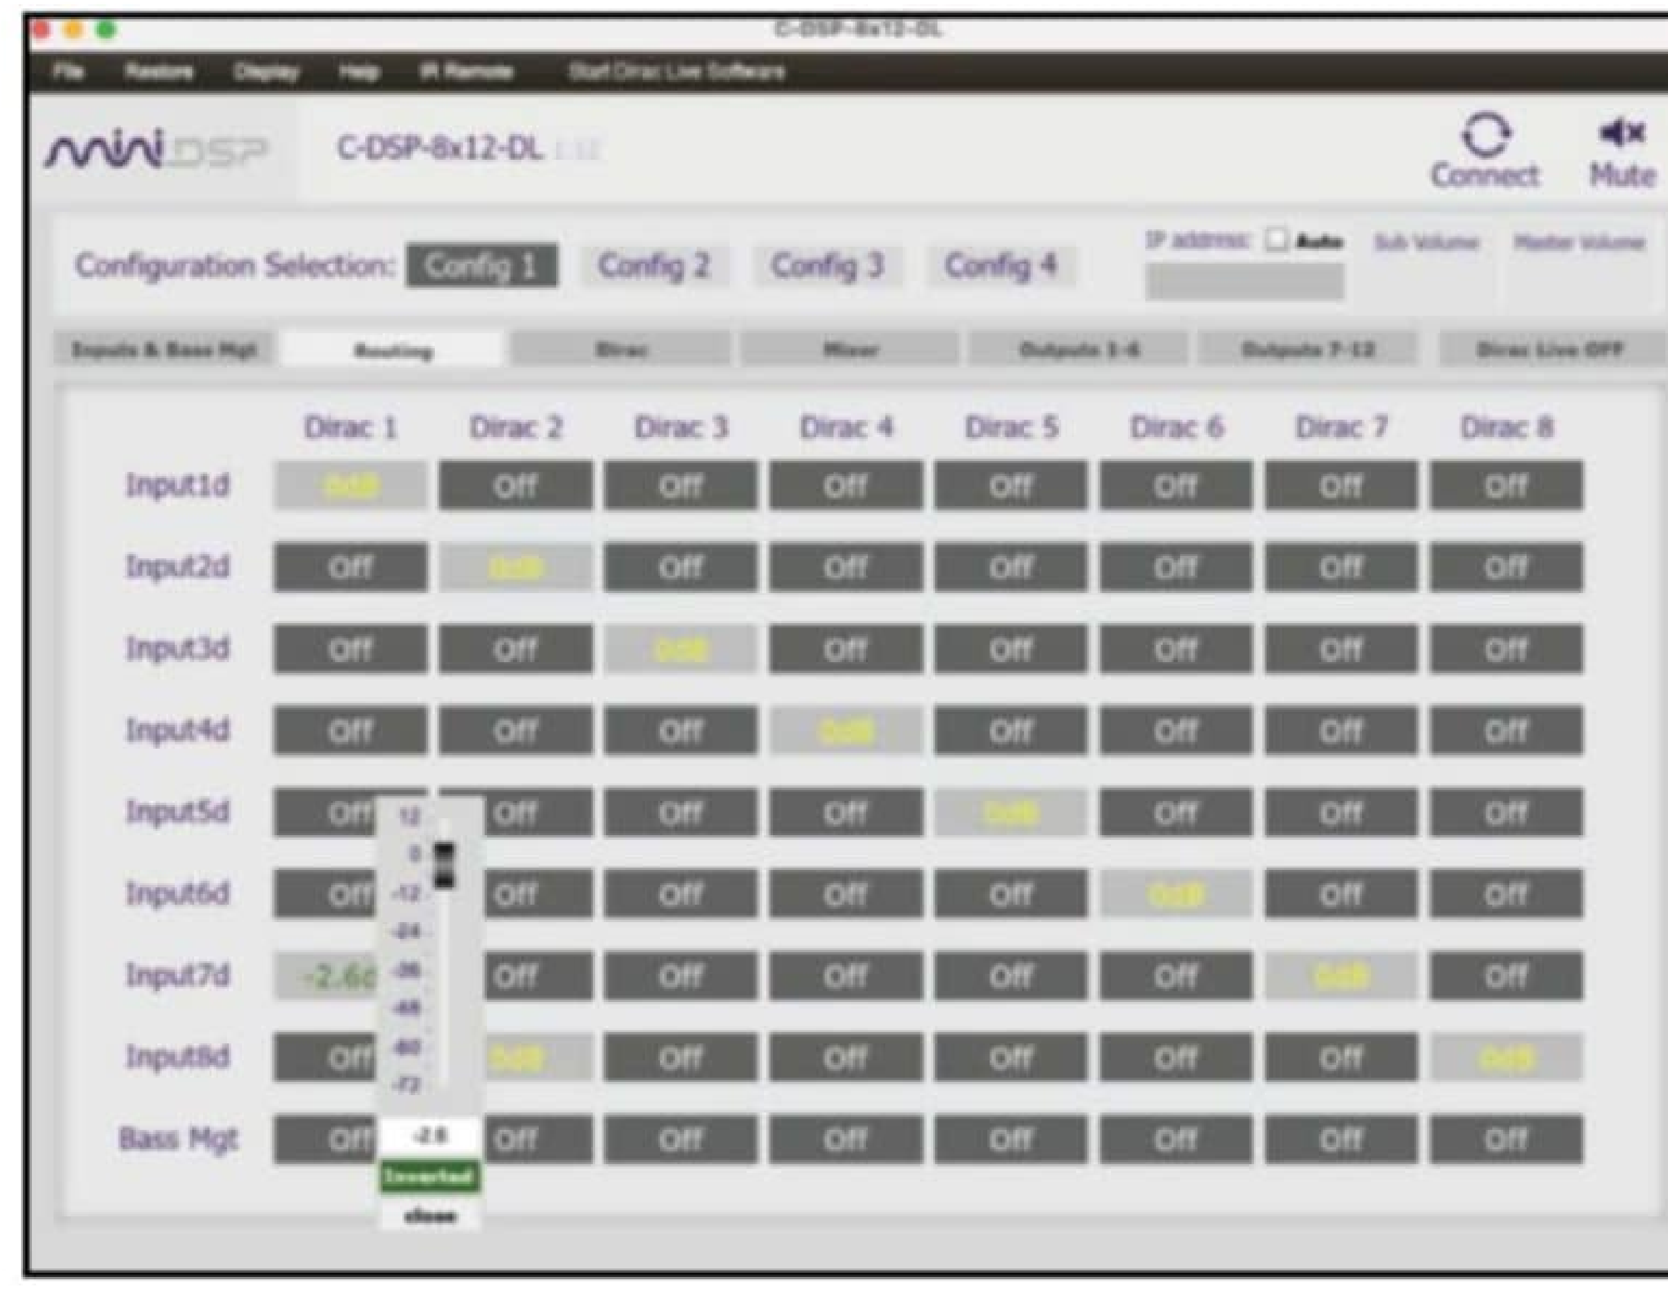

Figure 10: Matrix mixer in the miniDSP Harmony 8x12. An input signal can be distributed to several outputs and manipulated in level and polarity.

through listening first before adding another filter band. Due to the many changes in the spectrum, even for someone with a trained ear it's difficult to even out spectral deficiencies by maxing out all the EQ bands from the get-go and listening.

Narrow Bands

Narrow bands are often not useful: It is tempting to even out the smallest peaks and dips. But often filters with Q > 9 are not helpful. Reason being that the spectral resolution of our ear (which has an approximate resolution of a little more than 1/3 of an octave) is limited and the spatial effect of your filter turns negative or can null when your head is moving even the slightest from the sweet spot position.

Null Points

There are some frequencies that simply cannot be EQ'ed: Null points. They are found in every car and can be maddening. At these frequencies it simply makes no difference if you boost or cut. Nothing, little, or very little is happening in the frequency domain. It can be necessary to restructure the tuning of a channel around a certain null point to keep the desired inner spectral balance of a system intact.

Better Negative than Positive

When reading graphs, it seems to be a natural reflex to add something when there seems to be a lack of level in a certain frequency band. But it is better to think the other way around. Remove what is "too much" instead of adding what is not there.

Digital sources leave only very little headroom, that is why adding gain with EQs is not recommended: The digital path could be overloaded. Not that it can't be done or that it results in instantaneous bad sound quality, but consider it twice before using it. Most likely an attenuation in the input stage of the filter is needed instead.

Allpass, Maybe?

Allpasses are probably the most underrated filters and can be a lifesaver in automotive. Although they do not affect the spectrum, they affect the phase of the processed signal. This can be helpful on midrange and midbass channels to make staging and integration perfect, when delay in just a certain frequency band is needed to shift a certain aspect in the imaging.

Gain Staging

During the calibration and tuning process, you should keep an eye on your gain structure and make sure that you are not clipping the DSP system at the inputs and outputs of different modules or using unnecessary digital gain to boost analog noise. Note that not only boosts, but also cuts in EQ bands, can increase the output level and therefore lead to distortion!

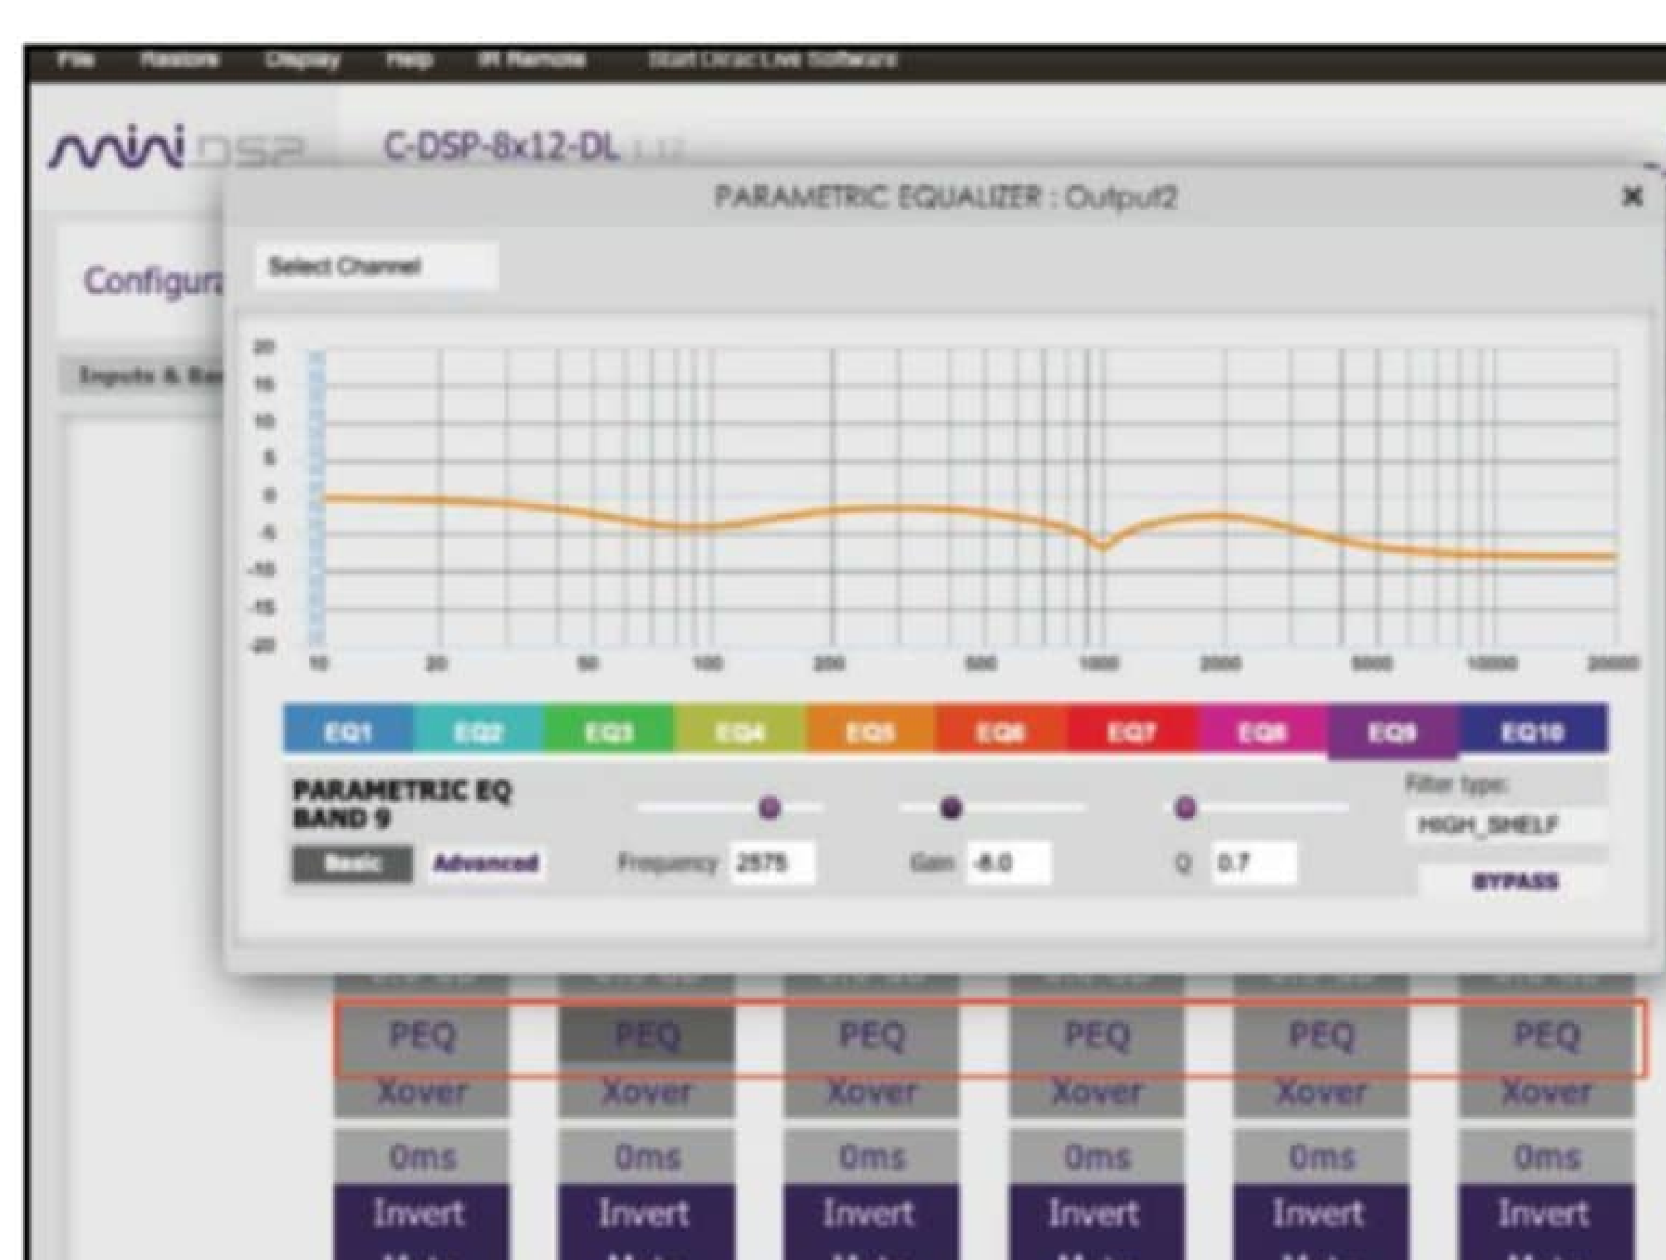

Figure 11: 10-band EQ per output channel.


HEAD acoustics
I can finally select the info I need...
using Bluetooth® LE Audio Auracast™ – Reference-tested by HEAD acoustics
REAL-LIFE-PROOF.
www.head-acoustics.de

Audioscenic: From Vision to Product

Audioscenic is an innovative 3D audio beamforming technology company spun off from the University of Southampton. It is a remarkable example of perseverance, from early development to the first commercial implementation, which *audioXpress* considered Best of Show at CES 2023. This article shares the story so far, as Audioscenic continues its groundbreaking journey exploring consumer applications for its forward-thinking approach to spatial audio.

By
Marcos Simón
(Audioscenic CTO and Co-Founder)

The S3A Project

Audioscenic started in 2015, when I met Professor Filippo Fazi, working together on the S3A audio project. The S3A program was a 5-year UK research project that aimed to bring spatial audio into people's living rooms and develop innovations that would increase the in-house adoption of spatial audio consumption, which at that time was stagnant due to the slow sales of home cinema systems, mainly caused by surround home systems being complex and bulky and hard to set up. This was in 2015, so ambisonics for headphones and virtual reality (VR) had not yet arrived on the market.

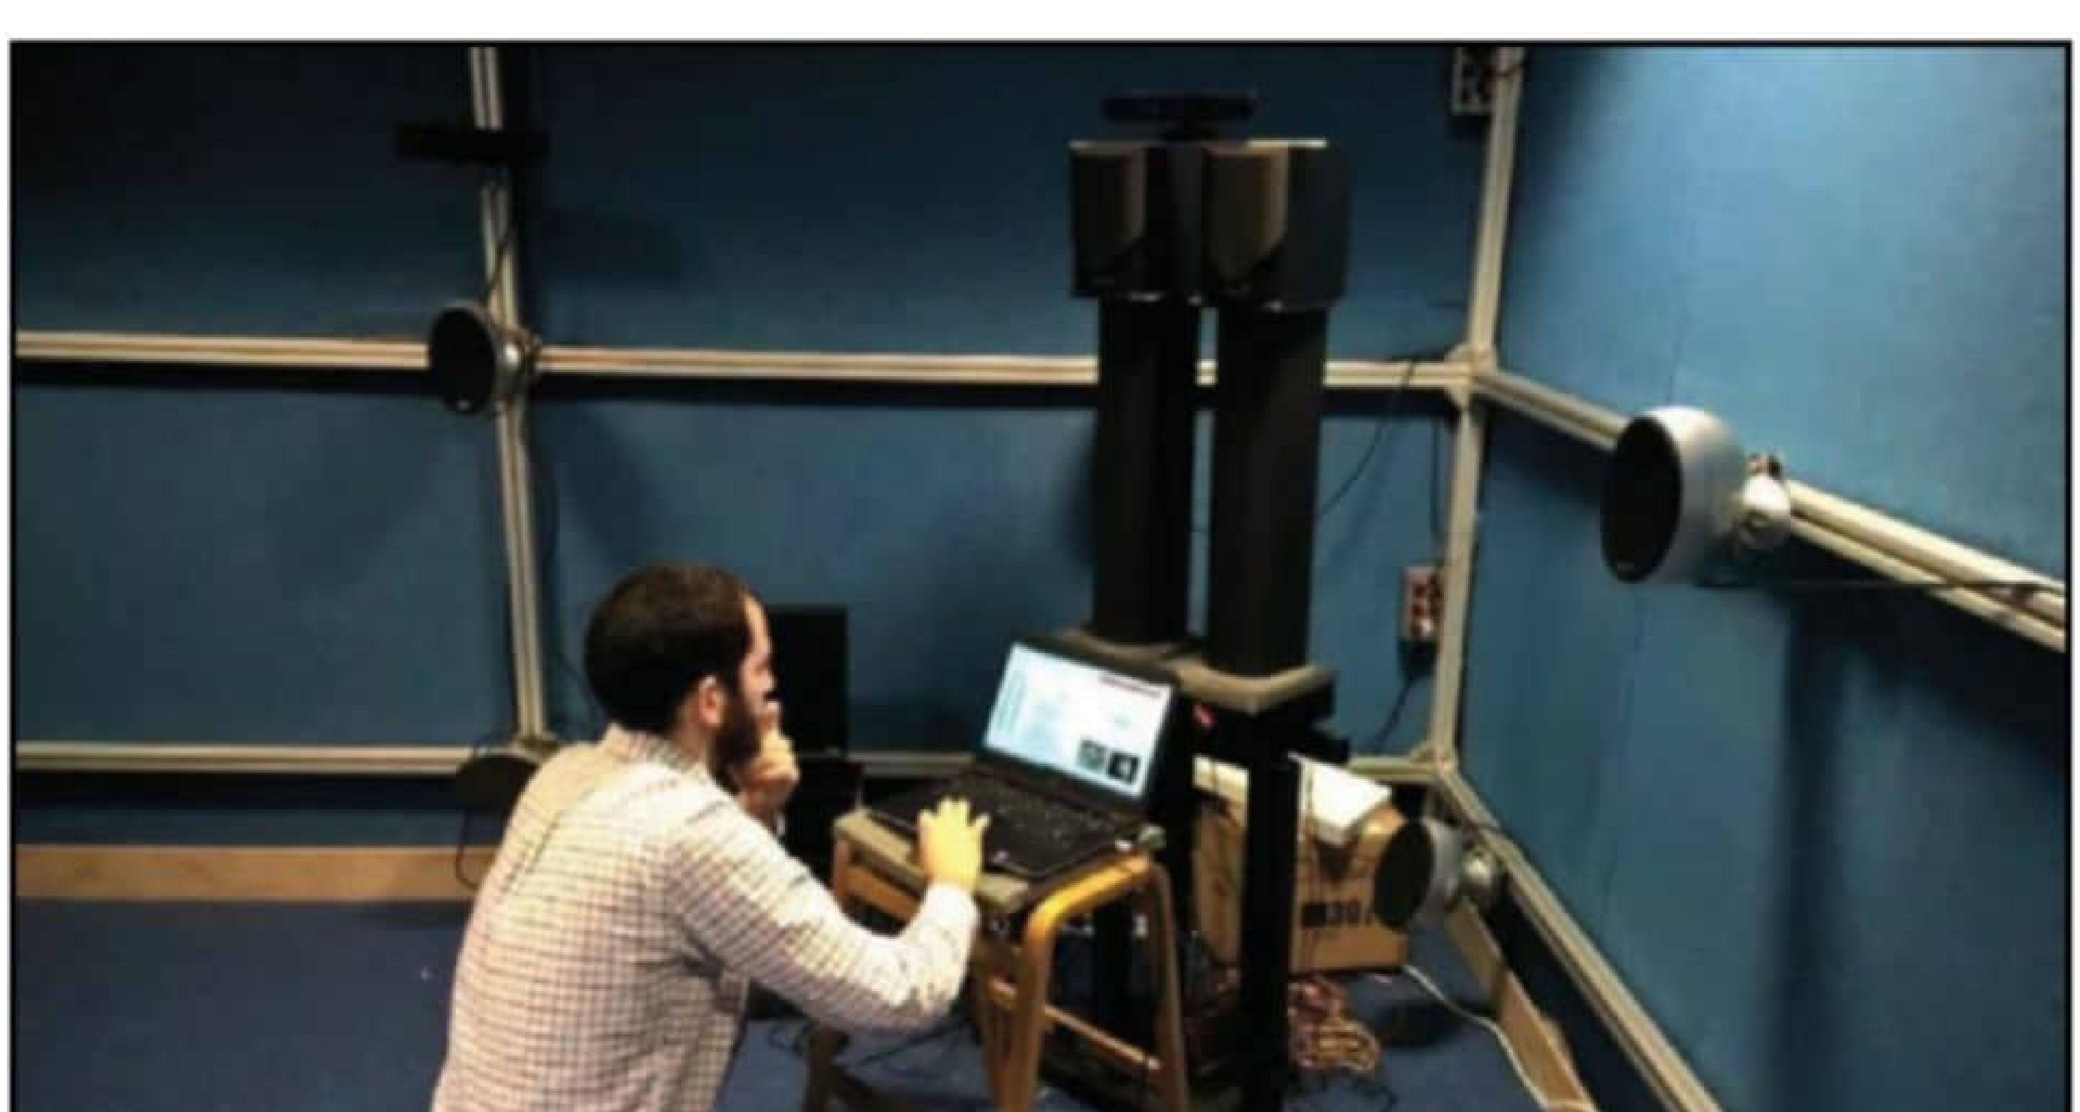

Early days in 2015 at the University of Southampton working on a prototype of a user-adaptive stereo dipole, using a Microsoft Kinect 1 and a pair of stereo speakers to give user adaptive crosstalk cancellation audio. It was the start of further developments toward user-adaptive spatial audio.

At that time, I had just graduated with a PhD in personal audio, specifically experimenting with using loudspeaker arrays to create boosted audio zones to increase speech intelligibility for hard of hearing TV viewers. Filippo Fazi was the principal investigator of the S3A project at the Institute of Sound and Vibration Research (ISVR) at the University of Southampton.

The ISVR had a long tradition with the development of crosstalk cancellation systems, dating back to the early 1990s. Filippo and I were both fascinated by the breathtaking spatial audio performance that a good crosstalk cancellation system could provide. We were driven by the idea of a soundbar that could provide full 3D audio as a method of bringing spatial audio to the people. A soundbar was something more appealing to put in the living room than a full surround setup, and this idea had been tried previously commercially by companies such as Opsodis and Comhear.

Until then, crosstalk cancellation systems were designed to work with a single, fixed, sweet spot. This was a real problem as they were only working effectively if a user was placed exactly in front of them. This was not ideal, as consumers would not widely accept such a system and therefore crosstalk cancellation systems were only used in academic and niche, enthusiast applications.

One of the main research methods of the S3A program was to use computer vision to locate users and then steer the sweet spot of audio systems accordingly. This is what I did with the crosstalk cancellation technology, by combining a Microsoft Kinect with a loudspeaker array, I built one of the first user adaptive spatial audio arrays (the Sound Virtualizer, which was featured in one of the first issues of *The Audio Voice* newsletter that year).

One of the main advances achieved by the technology developed by me and Filippo was simplifying the signal processing required to implement a multichannel crosstalk cancellation using a minimum amount of processing power, making implementation of the technology in small embedded systems possible.

First Investor Meetings

In 2016, we patented the technology with the help of the University of Southampton. Back then, I was increasingly motivated by the idea of finding a way to take the technology outside of the laboratory and implement it as an accessible consumer product.

One of the things I wanted to avoid is having the invention just be published as a Journal Paper and then forgotten, collecting dust at the University labs, as happens with much academic research. And then I saw an opportunity when we were given £35,000 by the UK government to explore the commercialization of the technology and conduct market research.

Back then I built a very early prototype by myself at the University workshops. It traveled the whole world and we were able to show it to the global audio industry. But even though it was still a prototype, it was very encouraging to do demos at CES and hear visitors saying that it was the best demo they had seen at the show...

Coming back from CES in 2017 with such good feedback gave us the required motivation to continue developing the technology. It was sometime later that year when I was introduced to Dr. Lee Thornton (who later became Non Executive Director of Audioscenic). Lee Thornton saw great potential in the technology and started digging into the possibility of funding Audioscenic to take the technology out of the University and form a startup company.

The company was named "Soton Audio Labs," in reference to the long tradition of the development of crosstalk cancellation systems at Southampton. At that time I was getting a great deal of help from university interns who built some early prototypes. We experimented a lot, building both smaller and larger prototypes with a focus on improving the overall sound quality. We built a small 7-channel array using 2" Celestion speakers, and that was the basis of our gaming soundbar platform. The system was built by an intern called Charlie House, and so it was named the "Charlie" Array.

This was also a major breakthrough, because it was the first time that we made our technology run inside a Raspberry Pi and so it became our minimum valuable prototype. Andreas Franck, later our software lead, helped me quite a lot with the

The Sound Virtualizer was the first array built to do user-adaptive crosstalk cancellation as the S3A project. It uses is no less than 28 1" speakers because "let's see what we can do with loads of loudspeakers and then reduce them down!" Microsoft Kinect 1 and 2 were used to track users.

As the first prototype was "too big," a second was built to be demoed around the world. Marcos Simón and Filippo Fazi pose with the prototype at the Future Worlds stand CES 2017. Future Worlds is the University of Southampton's business incubator, which helped the founders in their early days.

real-time audio implementation. Charlie House and I also built a two-way 28-source array that we named "Pata Negra," which is a Spanish synonym for the best quality ham one could obtain. This was a real achievement as we were able to improve the sound quality of our early university made prototypes—which in turn began opening doors for us in the audio industry.

These developments led us to have the first prototype of our gaming soundbar design, which we built for CES 2018. My cousin helped me do the industrial design, and the prototype was built using Intel RealSense development kit with an UP board. I remember spending Christmas in my father's garage finishing the electronics and painting the enclosure.

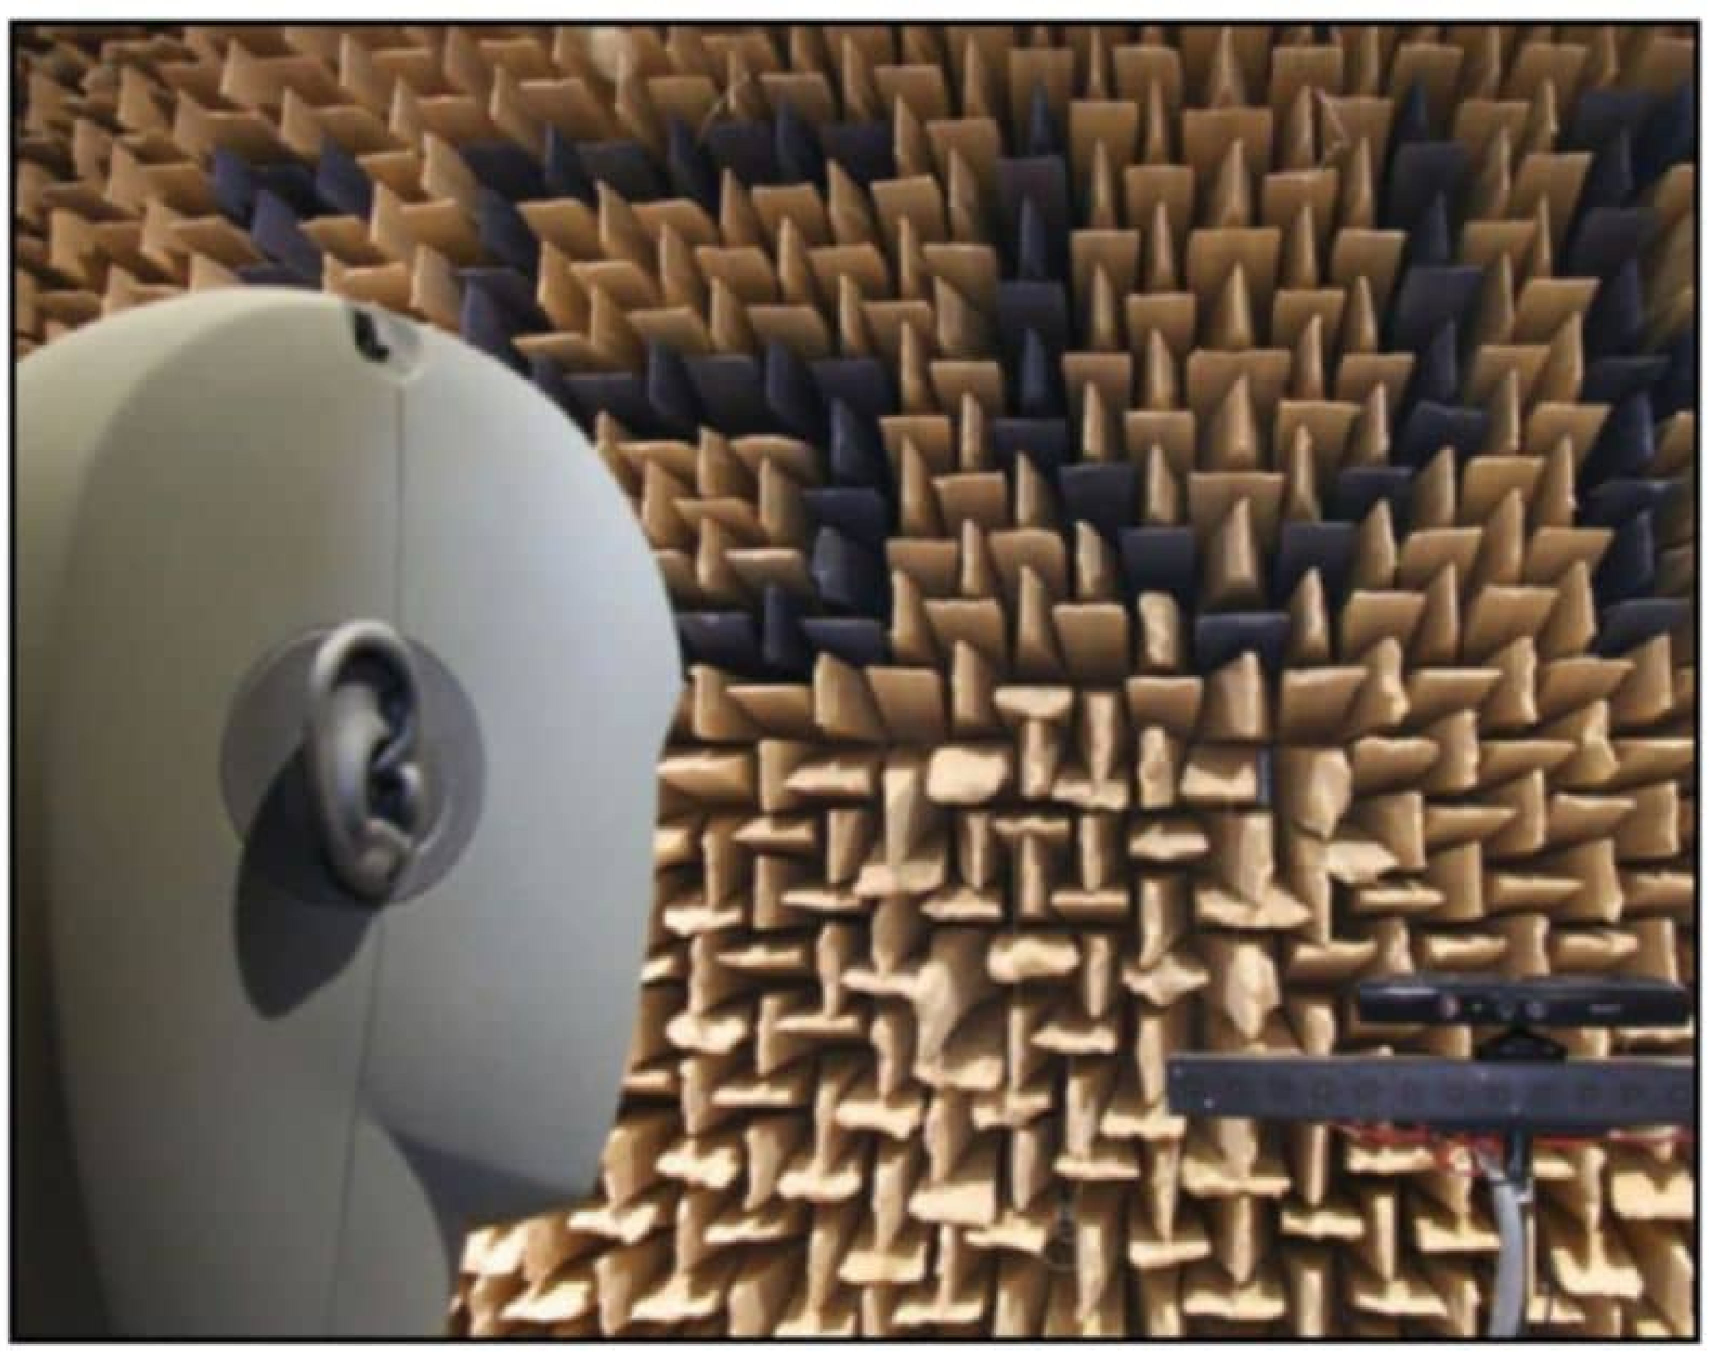

This was the third soundbar that was used to work in user adaptive spatial audio. It is an old prototype that was used in the ISVR and known internally as "Sound Bender," brought by Filippo Fazi from California. The soundbar used a 16-channel loudspeaker array and produced very good crosstalk cancellation in the mid-frequency range. Used to prototype the first iterations of the Audioscenic Virtua technology.

The first gaming prototype using 2" Celestion drivers and an Intel RealSense R200 for user tracking, back then still under the "Soton Audio Labs" name.

Transitioning Audioscenic to be a fully operative company was a slow process and it was hard initially to advance in both technical and commercial terms as we were still developing our team. It was mainly me with the help of some interns, since Filippo was busy with the University. However, in 2018, we were introduced to David Monteith, who was to become Audioscenic's CEO and helped us win a pitch that secured the first round of funding for the company. David Monteith had previously founded and successfully exited a pair of companies that were key to the development of 3D audio in the 2000s: Sensaura and Sonaptic.

The First Year

The team, which now also included Senior Engineer Tyler Ward, traveled again to CES 2019, this time with a prototype that was a small soundbar built as a laptop accessory. It was a long trip and a long show where we were still looking for a way to take the technology toward a product. Luckily, the last meeting of the last day of CES connected us to the technical leads of a Shenzhen, China-based ODM... paving the road for our first commercially oriented prototype.

In 2019, Audioscenic was formally founded and soon it grew to six people, with the addition of software engineers Andreas

Filippo Fazi, Marcos Simón, David Monteith, and Lee Thornton after they pitched to the IPGroup investment committee and secured the first funding that would start Audioscenic (September 2018). Thornton would later be highly involved in the company as a Non Executive Director of Audioscenic.

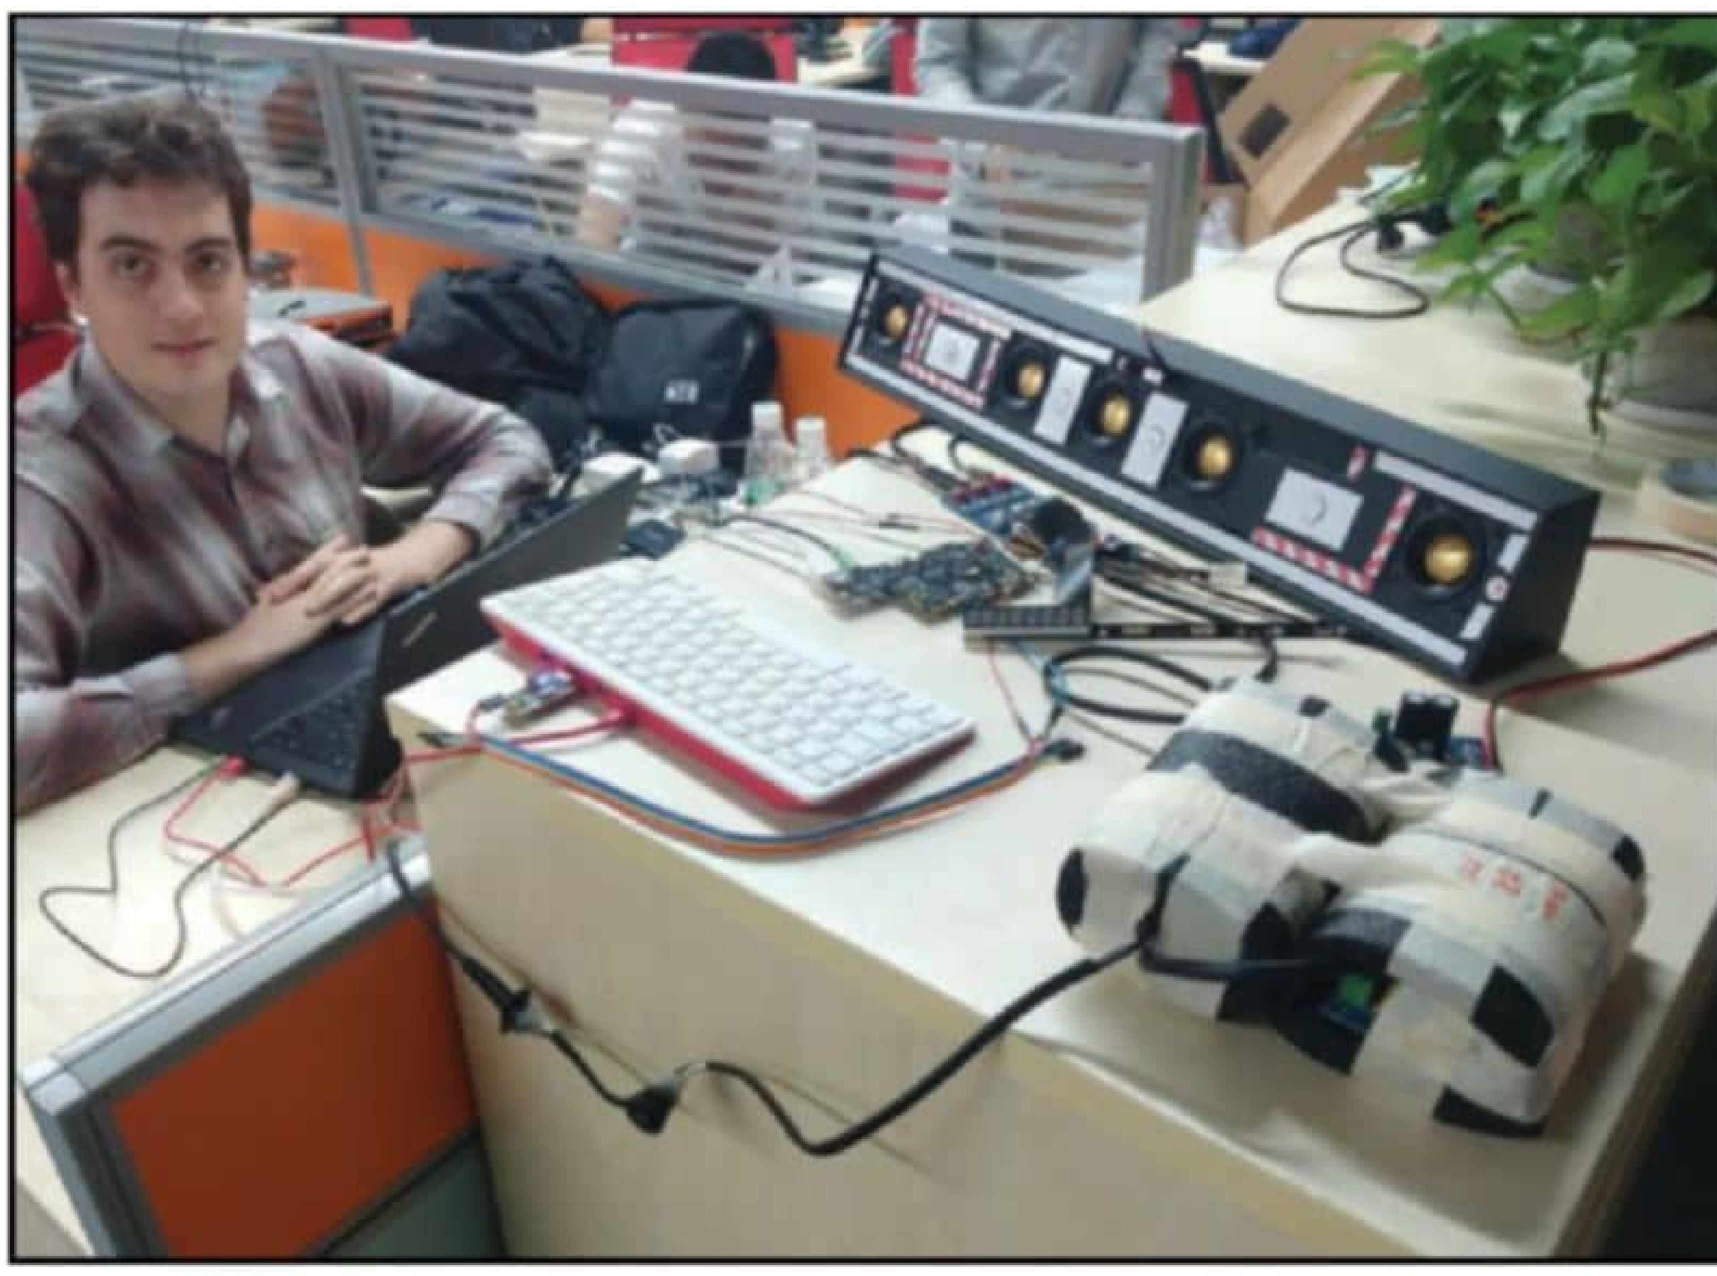

Tyler Ward and Marcos Simón in Shenzhen around Christmas 2019 testing the prototype before CES.

Franck and Joey Laghidze. We built some prototypes of the gaming soundbar and sent it to the ODM so that they could start showing it to their customer base. The year progressed and by Autumn we were working on a reference design for CES 2020, which forced me and Tyler Ward to spend the last two weeks of December in Shenzhen completing the prototype to be presented at CES. The team worked 12-hour days and weekends to make sure the first prototype was fit for a demo!

Although the previously developed prototypes worked well in the lab, obtaining a similar level of performance in a fully production-ready embedded system proved to be tricky! The team was finally successful, and when we arrived at CES, the demo was very well perceived by everyone who visited the ODM booth!

Company Growth in Pandemic Times

We returned from CES 2020 knowing that we had a great show, we had a product with a customer lined up, and we were preparing a pitch to raise more money. Happy days. Little did we know that a pandemic was coming...

It was March 2020 and the UK along with the rest of the world was going into lockdown, causing a stagnation that really complicated our technical and commercial developments. During

Tyler Ward, David Monteith, Marcos Simón, and Filippo Fazi after their last meeting at CES 2019—a demo that connected them with the ODM that would later make the company's first product.

The Audioscenic hardware development kits (HDKs) were built with off-the-shelf components and allowed the company to ship them and virtually give demos during the pandemic.

those days we struggled, given that our technology is one that you have to hear to be convinced of its performance and it wasn't easy to do demos remotely! Apart from that, we needed to raise more money to continue funding the business, which we did in October 2020 thanks to an investment from Foresight Williams Technologies to Audioscenic.

To try to ease the pandemic effect, we hired a commercial person to help us in China, and we also developed our hardware development kits (HDKs), which we could ship and loan to customers and collaborators, so they could remotely listen to our technology.

During the pandemic, there were tough times but we continued developing the technology, always with the vision of "creating systems that will automatically adapt to the user's position." We were successful in our funding round and added more roles to the team: Daniel Wallace, as R&D lead; Roberto Grilli and Davi Carvalho as R&D engineers; Ria Atara as a third software developer; and Kim Davies as office manager.

From the beginning, the vision of Audioscenic has been to develop audio systems that can give an extraordinary spatial audio experience and that require nothing or little setup by the user. User experience is the primary driver of what we do, and of course, all of this is based on solid theory building on our pedigree from the ISVR. During the pandemic, we kept faithful to this vision so that innovative companies were made aware of our technology.

The hard work invested in the company by all its team members paid off, with the first commercial product based on Audioscenic technology being released at CES2023: The Razer Leviathan V2 Pro. The soundbar, developed in partnership with Razer and THX was the world's first PC soundbar with head-tracking AI and adaptive beamforming. We consider this to be a great achievement.

Not only for the company, but for the whole spatial audio industry in general: a speaker-based system capable of rendering high realism spatial audio; and not only providing an alternative to headphones, but able to provide an immersion and externalization far superior to headphones.

The Razer Leviathan V2 Pro obtained 12 best in show awards at CES 2023 and Audioscenic was given the Best in Show tech mention by *audioXpress*!

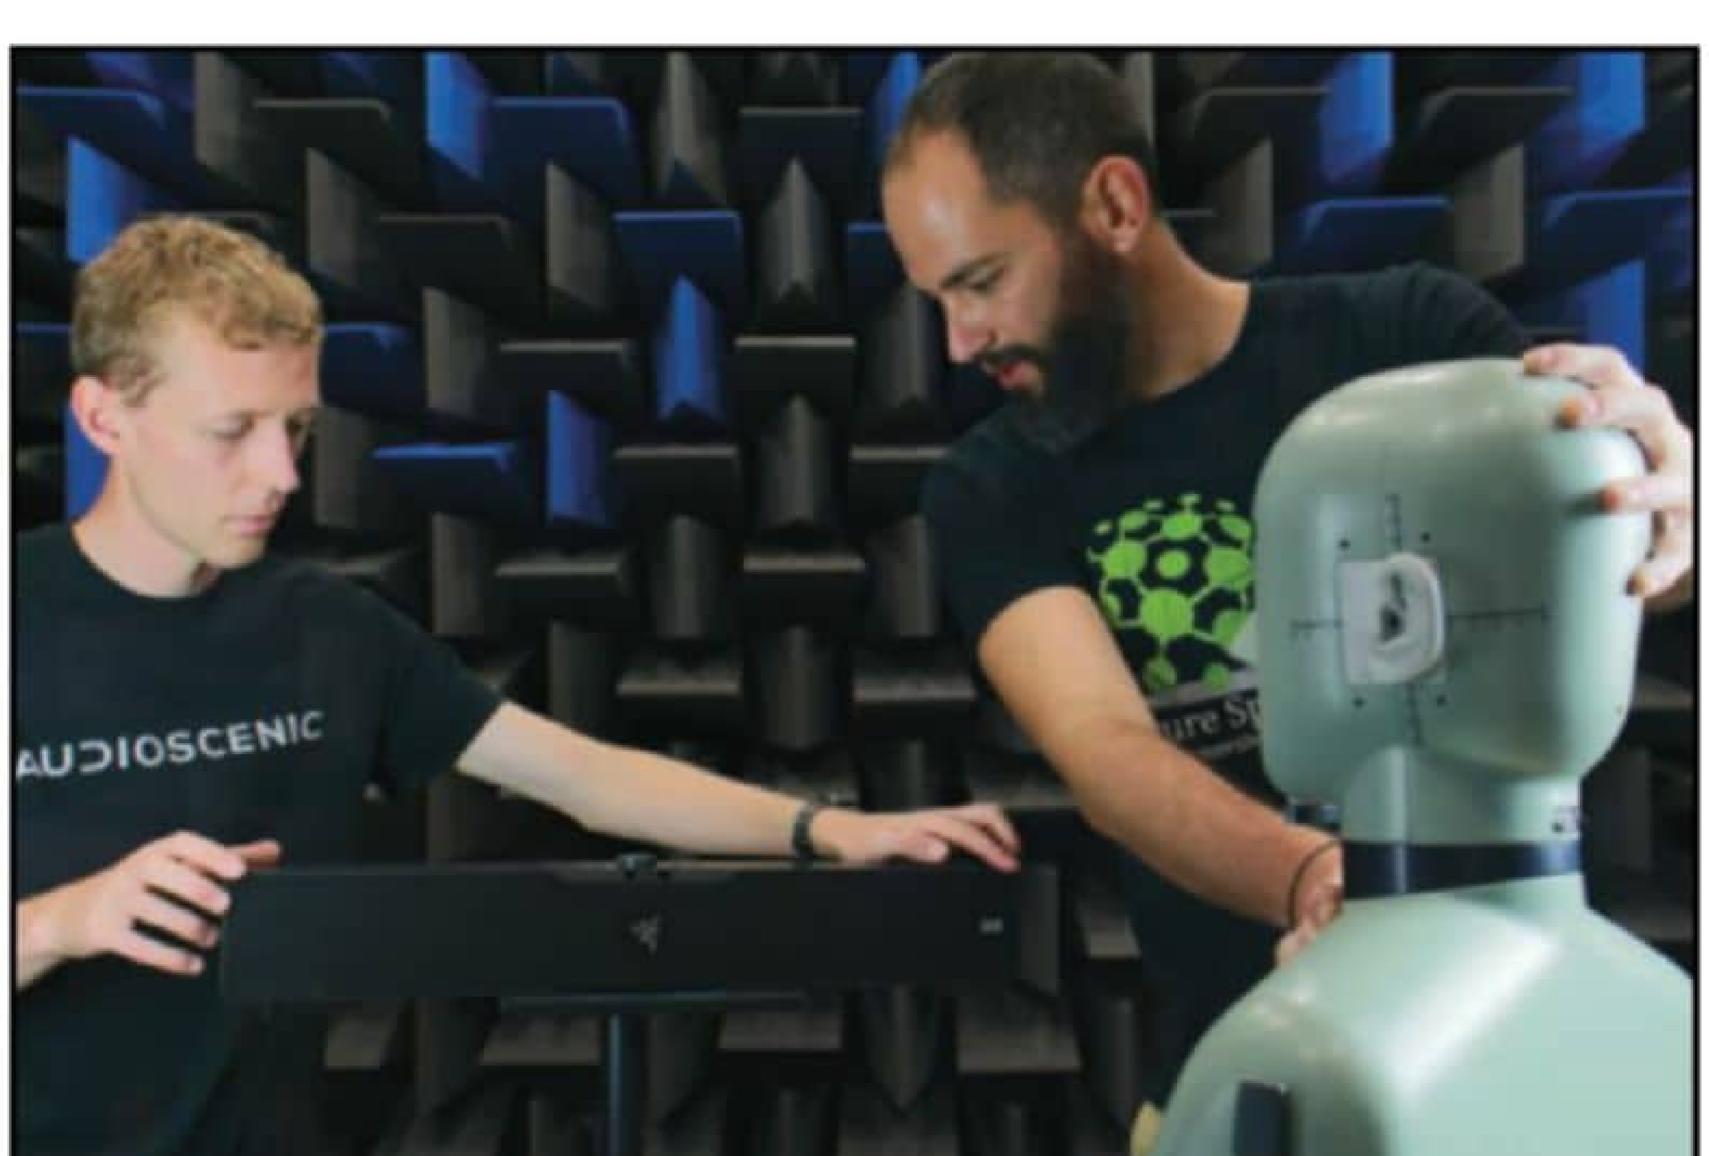

Daniel Wallace and Marcos Simón at the ISVR Anechoic chamber while working on the development of the Razer Leviathan V2 Pro soundbar.

About the Author

Marcos Simón is the Co-Founder and CTO of Audioscenic limited. He is a world expert in sound field control and array signal processing who became an entrepreneur to change the way people currently consume spatial audio. Marcos graduated in 2010 from the Technical University of Madrid with a B.Sc. in telecommunications. In 2011, he joined the Institute of Sound and Vibration Research, where he worked with loudspeaker arrays for sound field control and 3D audio rendering and also in the modelling of cochlear mechanics. He obtained his Ph.D. title in 2014, and between 2014 and 2019 he was part of the S3A Research Programme "Future Spatial Audio for an Immersive Listening Experience at Home." In 2019 he co-founded Audioscenic for the commercialization of innovative listener-adaptive audio technologies, where he currently works as Chief Technical Officer. Since the creation of the company, Marcos has been leading the vision of Audioscenic and established himself as a nexus between the commercial and technical world for the start-up; making sure that the technology is continually evolving and that the customers understand exactly what the technology makes possible. Marcos combines his innovation passion with love for cooking, sailing and DJing electronic music.

Beamforming, Automotive, and The Future

It has taken much time and hard work since Filippo and I started developing the technology in 2015, but the perseverance has paid off and the first product based on the technology is now on the market, with the aim of bringing spatial audio to people in a practical and user-friendly way. A vision that, despite a global pandemic and wars, Audioscenic has managed to push forward as a company.

The technology that Audioscenic has developed allows for any kind of user-adaptive DSP using combinations of speakers, so it can also be applied to other uses apart from crosstalk cancellation. For example, beamforming or the creation of independent listening zones for automotive, both areas that are currently being explored both technically and commercially and that are also expected to become products in the following years.

Right now, Audioscenic is a growing startup and many others have joined its team including George Howell, Balu Raj, Yury Ushakov, Dimitris Zantalis, Jose Cappelleto, and Jago Reed-Jones. This is creating a great future for Audioscenic, where we see ourselves providing technology for a variety of use cases that does not involve wearing headphones. At Audioscenic, we want to believe that "the future of sound is spatial and that it will be consumed either with wearables or with an Audioscenic enabled system." As a company, we are now a happy family with a presence on various continents that will fight and strive to continue growing and spreading its vision!

The Audioscenic team now in 2023.

Resources

Audioscenic, www.audioscenic.com

"Audioscenic Partners with Razer and THX to Launch Desktop Soundbar with 3D Beamforming Technology and Head-Tracking AI," https://audioxpress.com/news/audioscenic-partners-with-razer-and-thx-to-launch-desktop-soundbar-with-3d-beamforming-technology-and-head-tracking-ai

J. Martins, "Audio Innovation Never Stops," *The Audio Voice*, Issue 109, December 2016.

The S3A audio project, https://cvssp.org/projects/s3a

Sound Control

An Exploration of *Rock and Pop Venues*

Acoustic and Architectural Design

By
Richard Honeycutt

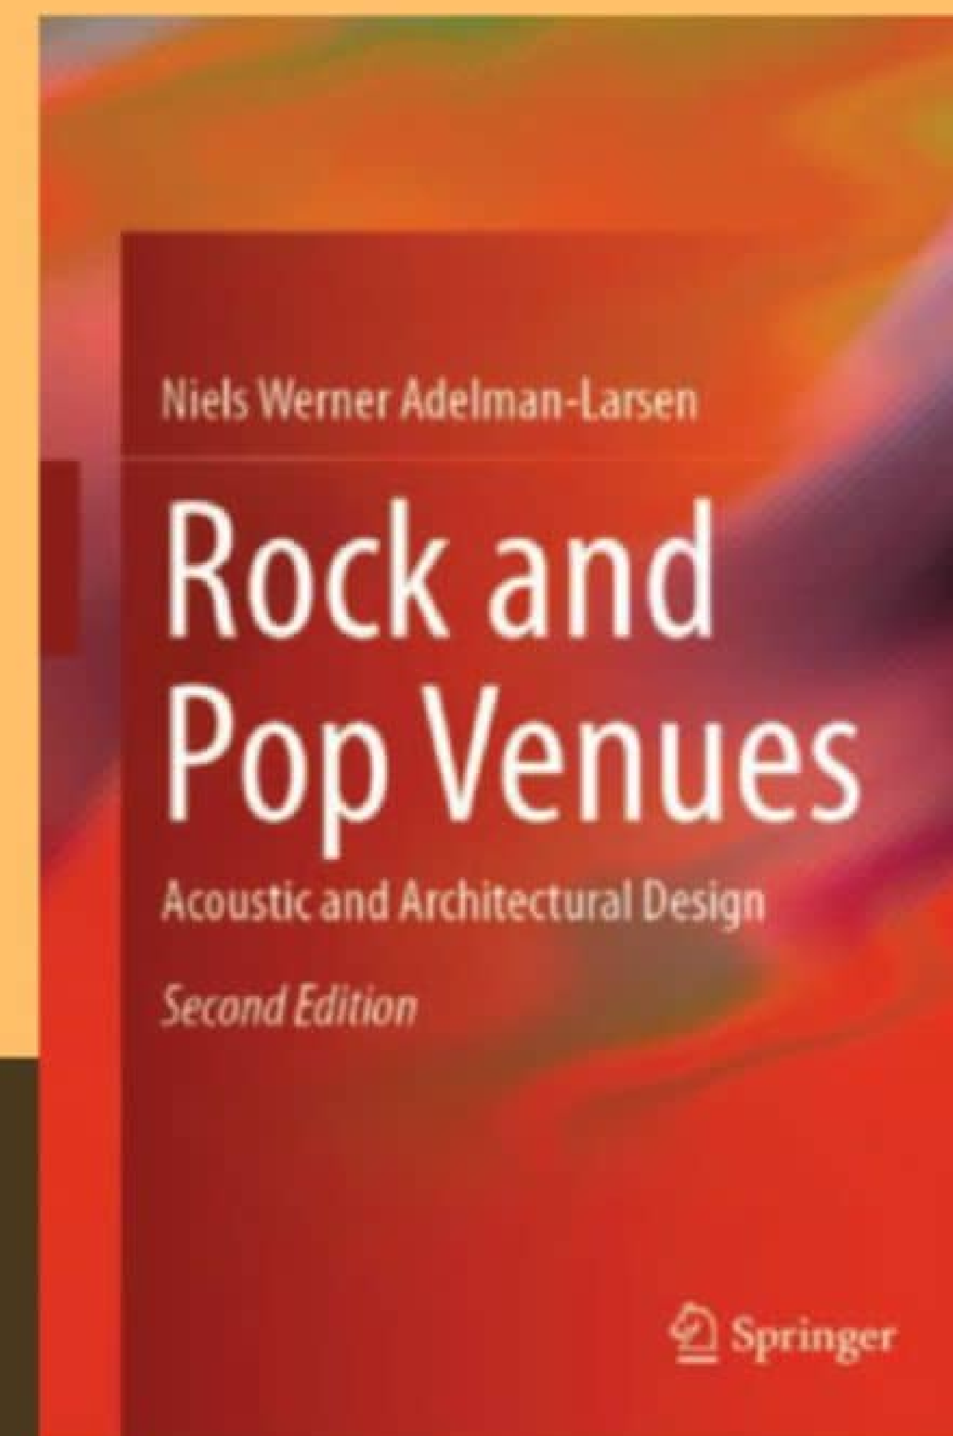

For this month's article, Richard Honeycutt reviews the book *Rock and Pop Venues: Acoustic and Architectural Design*, by Niels Werner Adelman-Larsen, recently updated with a Second Edition. This book was referenced in previous Sound Control articles.

Optimal acoustics for concert halls have been studied since the early 1900s, but mainly with an emphasis on supporting Baroque, Classical, Romantic, and later forms of "serious" music. By the latter half of the 20th century, differences among the acoustical requirements for these types of music had been investigated.

However, little attention had been given to halls for jazz performance, partially because jazz was often thought of as dance music, and partially because good acoustics for jazz performances differed little from good acoustics for most types of "serious" music. One exception was the percussive Latin American forms of jazz, which required better clarity and definition in a listening venue.

Rock and pop music began developing a significant concert presence in the 1960s and 1970s, first outdoors, then later in sports arenas and coliseums. The electronic/electroacoustic amplification of rock and pop music added new dimensions to concert acoustics, but for years, these dimensions were ignored, especially in the "sports" venues often used for rock and pop performances, at least in the US.

If an architect or acoustician wanted to design a truly good performance venue for rock and pop concerts, (s)he would have been severely hampered by lack of rigorously determined objectives for early decay time (EDT), reverberation time (RT) in the various octave bands, clarity, definition, intimacy, lateral reflections (LF), room gain (G), bass ratio, warmth, and stage support/ensemble.

Basing a design upon the preferred values of these parameters for other music genres certainly would—and did—result in suboptimum rock and pop concert experiences. Filling this void in our knowledge of audiences' acoustical preferences for rock and pop music venues has been the focus of Adelman-Larsen's work for more than a decade, and hence, is the focus of this book.

Rock and Pop Venues:
Acoustic and Architectural Design, Edition 2
By Niels Werner Adelman-Larsen
Cham, Switzerland; Springer, 2021
ISBN: 978-3-030-62320-3

Book Overview

This book was first published in 2014. This review covers the second edition, published in 2021, which is available in e-book, soft-cover, and hard-cover. The Table of Contents of the second edition differs from that of the first edition mainly in the elimination of the "Endorsements" from second edition. Some changes in wording can be found within the chapter text, but the organization and coverage remain as excellent as in the first edition.

In his introduction, the author succinctly summarizes the challenge that his research addressed:

"Pop and rock concerts are unique in that they depend heavily on amplification of the sound that the band produces. A significant ingredient in the music is often a highly amplified and well-articulated bass line supported by a more or less syncopated staccato bass drum rhythm. The precise timing of both, down to few hundredths of a second, is what professional orchestras depend on in order to make hard grooving music not only satisfactory to the orchestra itself but also, most notably, to the audience.

Many other instruments are active in the bass frequency domain. The room's acoustics should enable this message to be transmitted to the audience with a high degree of intelligibility. It is evident that if a hall allows the bass tones to last for a long time, then this rhythmic backbone of the type of music we are describing is easily lost. The starting and stopping of bass sounds is converted rather into a long, dull, legato continuum that can hardly be identified as "rhythmic" by performers or audience. What was expected to be a musical experience is transformed into sound with little sense, sometimes close to noise." [1]

The book begins with an introduction to acoustics and hearing. Readers who have learned about acoustics primarily from high-school or undergraduate college courses in general science or physics will gain new insights as Adelman-Larsen addresses the production, characteristics, and radiation of human vocal sound; the human voice usually contributes the most important components to rock and pop music.

Another particularly valuable part of the first chapter is the discussion of absorption by porous, diaphragmatic (membrane), or resonant elements, and by air. Many of my architect clients could benefit from a study of this discussion, as could many people who design their own home studios.

The Mathematics

The concept of critical distance (the distance from the sound source at which the direct and reverberant sound levels are equal) is well known to most sound engineers and acousticians. But musicians, facilities managers, and others reading this book may not be familiar with it.

Yet understanding the parameters that determine critical distance, and its importance in the listener experience in a hall, is of great importance in understanding concert hall acoustics. The author presents this information using enough math to elucidate the concepts, but not so much as to threaten "math-o-phobes."

Speaking of math, I should mention a potential source of confusion on page 23 of Chapter 1. The author correctly gives the equation:

$$L_P = SPL = 10\log\left(\frac{P_{rms}{}^2}{p_{ref}{}^2}\right)$$

where p_{rms} is the effective sound pressure in pascals, and P_{ref} is the reference sound pressure of 20µPa.

However, many, if not most acoustics texts give these two equations:

$$L_P = 20\log\left(\frac{P_{rms}}{p_{ref}}\right)$$

$$PWL = 10\log\left(\frac{W}{W_0}\right)$$

where PWL is the sound power expressed in decibels (dB), W is the sound power in watts, and W_0 is the reference sound power of 1 picowatt.

Actually, these equations are equivalent, as you can understand if you know that sound power is proportional to sound pressure squared, which you can infer from the fact that the fraction in the upper equation is squared. Thus, if you double the sound pressure (twice as many pascals: an increase of 3dB) in the upper equation, you're taking the log of a value four times as large (an increase of 6dB). Often we say that doubling the sound pressure is an increase of 6dB, but doubling the sound power is an increase of 3dB. Without this clarification, the author's use of "10 log" instead of "20 log" for SPL could confuse casual readers.

Photo 1: The Concertgebouw Hall in Amsterdam is one of the finest traditional concert halls in the world.

Concepts and Terminology

In Chapter 2, the author discusses the concepts and terminology of auditorium acoustics. In addition to RT, he mentions the important but less commonly known parameters of EDT, reverberance, liveliness, clarity definition, early reflections, intimacy, lateral fraction, envelopment, G (strength or room gain), bass ratio, warmth, and ensemble. These are presented clearly and in appropriate detail.

Next comes a presentation of loudspeaker types and their applications to the reinforcement of the music of rock bands. The chapter begins with a characterization of the typical rock-music audio signal. In particular, the author discusses both the levels and the articulation of the bass. In doing so, he lays the groundwork for recommendations presented in later chapters, calling for a different

Photo 2: The Amager Bio in Copenhagen is one is the best pop/rock halls in Europe.

design approach regarding bass ratios. Leo Beranek defined bass ratio as:

$BR = (RT_{125} + RT_{250})/R(T_{500} + R\ T_{1k})$

He considered bass ratio to be the objective parameter corresponding to subjective "warmth." Adelman-Larsen redefines the bass ratio:

$BR = (RT_{63} + RT_{125})/(RT_{500} + RT_{1k} + RT_{2k})$

He finds that this definition corresponds well with subjectively optimum RT for rock/pop venues in the lower versus middle octave bands. In this understanding, BR is more closely associated with "bass intelligibility."

Chapter 3 discusses factors affecting the design of speaker systems for amplified music halls. The author describes the operation and performance of single-point-sources, virtual point sources (clusters, "exploded clusters," etc.), directional subwoofer arrays, and line arrays. Notably, the treatment of split mono vs. stereo discusses performance issues not always appreciated by system designers. This chapter will be helpful to musicians, sound technicians, architects, and facilities managers.

In 2005, the author made objective measurements and collected subjective data on 20 halls used for rock and pop music. For each of these halls, a photo, a scaled plan drawing, and the data are presented in Chapter 4. In Chapter 5, the author develops recommendations based upon the objective and subjective data. The February 2023 issue of *audioXpress* mentioned journal articles in which Adelman-Larsen recommended a substantial decrease in RT in the 63Hz and (most importantly) the 125Hz octave bands for rock/pop halls as compared to more traditional concert halls. On pages 117 and 118, the author summarizes:

Some textbooks on room acoustics recommend that "halls for music" have an increase of reverberation time at frequencies below 250 Hz (despite the fact that many of the very best rated halls for symphonic music don't show this trait without an audience). Surely, this brings "warmth" to the sound. But this is true only for (unamplified) classical music. For amplified pop/rock music, as shown in this chapter, it is enemy number one!

….

Some consultants have chosen almost anechoic acoustics for amplified music halls, assuming that the only focus point was freeing the audience from undefined reverberant sound. And that hypothesis has in some cases been taken to the extreme, sometimes even without enough focus on absorbing low frequencies. But as shown later, this is not a correct path to pursue. In this chapter, we will see what values of reverberation time are recommended for a given hall volume; it must be relatively short, but not too short, and within limits it can vary with frequency.

The author goes on to discuss his method of investigation and how the results informed his recommendations, which specify optimal RT versus frequency (in octave bands) for halls having cubic volumes ranging from $500m^3$ to $10{,}000m^3$. As indicated earlier in this review, such scientifically-derived recommendations were unavailable in the acoustical literature prior to Adelman-Larsen's journal articles. [2,3].

Conclusion

The recommendations are augmented by a comprehensive chapter of design principles, making them easier for an acoustician or architect to use. These design principles include a wide range of architectural factors, not only—or even primarily—acoustical parameters. The designer who adheres to these principles will be well on the way to achieving a good music hall design. An extensive gallery of halls used as rock and pop venues, plus four appendices, rounds out this 480-page tome.

I highly recommend Adelman-Larsen's book to anyone involved in the design of halls used as rock or pop music venues, as well as the musicians and sound technicians who work in these venues.

About the Author

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

References

[1] N. W. Adelman-Larsen, *Rock and Pop Venues: Acoustic and Architectural Design*, Edition 2, Cham, Switzerland, Springer, 2021, p. viii.

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

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

TAKE YOUR MEASUREMENTS ON THE ROAD

with AudioConnect 2™

Introducing the first laptop-powered portable audio interface for field measurements of car infotainment systems, headphones, and other devices. And it's great for production line measurement too!

- Low cost
- Compact, lightweight, and portable
- USB-C power and data transfer
- Sample rates up to 192kHz and 24 bits
- 2 TEDS-compatible inputs for microphones, head and torso simulators and test heads
- 2 output channels for driving powered speakers, amplifiers etc.
- Low distortion, low noise headphone amplifier
- Error-proof operation with exclusive software control, input protection, and BNC connectors

Listen

www.listeninc.com/ac2

The AP-2 Amplifier

With Analog Signal Processor for the ACH-01 Accelerometer

This project is a welcome upgrade of the original ACH-01 accelerometer system used for measuring and studying the vibration behavior of a loudspeaker, cabinet, stands, and more. The updated AP-2 amplifier received new useful features, some of which were suggested by users of the first design.

Photo 1: The updated AP-2 amplifier with an analog signal processor is designed to work with the ACH-01 Accelerometer.

By
George Ntanavaras

More than 10 years ago, *audioXpress* published my article describing the design and construction of a system for testing a loudspeaker with an accelerometer. The system consisted of the ACH-01 accelerometer and an amplifier with an integrated analog signal processor (see Resources).

The cost of the system was and is still very reasonable and the unit proved to be very popular among the amateur and professional loudspeaker builders who wanted to investigate vibration issues, particularly in loudspeaker enclosures. I made the PCB and the front panel of this amplifier available to DIY builders and a lot of amplifiers have been built around the world.

During all these years, I received several requests for modifications and additions to the amplifier. A couple of years ago, I decided to gather and look more carefully at these suggestions, and I started redesigning the amplifier.

The result is the AP-2 preamplifier, shown in **Photo 1**, which I will describe in this article. The new version includes more and higher Gain selections (from 0 to 60dB), high-pass and low-pass filters with selectable cut-off frequencies, a notch filter at 50Hz or 60Hz to reduce the mains interference that is usually picked up by the ACH-01, and a clip circuit with a LED indicator. The amplifier provides at its output, the acceleration, the velocity, and the excursion signal.

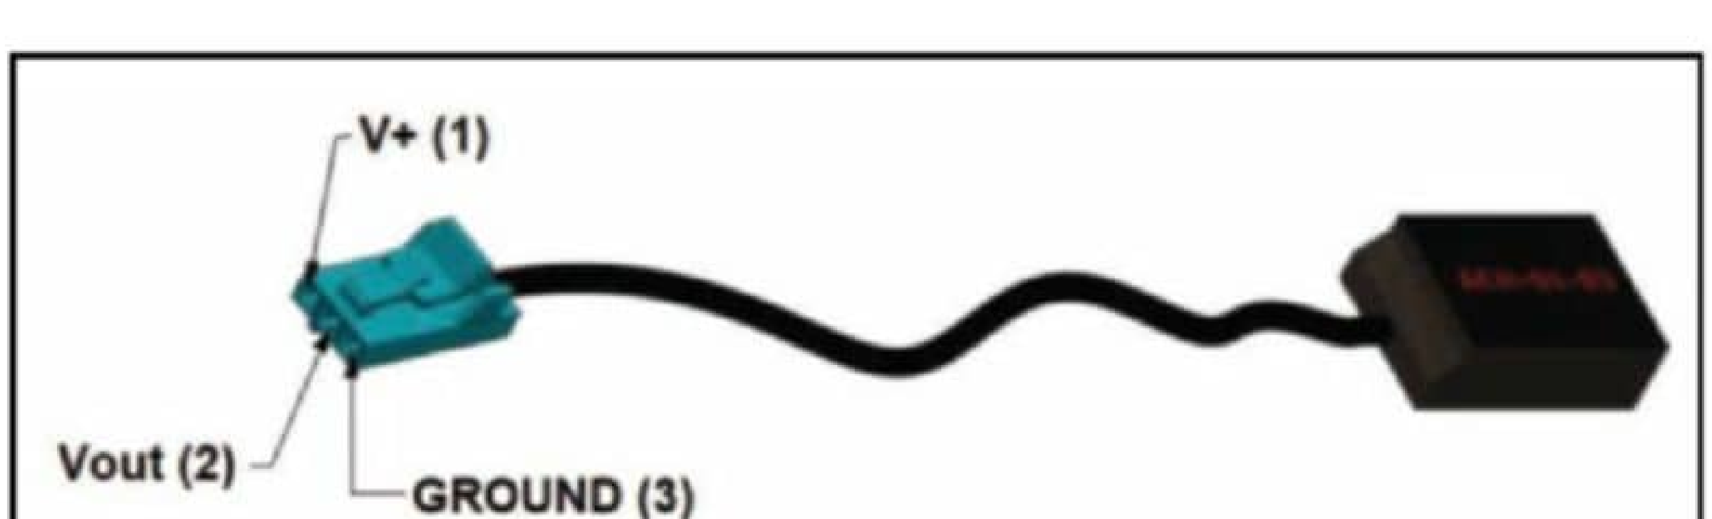

Photo 2: The ACH-01-03 version of the accelerometer is supplied with a shielded cable and a convenient connector. Here is the accelerometer with its connector pin-out.

The ACH-01 Accelerometer

The ACH-01, manufactured by TE Connectivity, is an inexpensive, general-purpose, linear, single-axis accelerometer, housed in a small, rugged, flat package with dimensions 13X19X6 mm.

According to the datasheet, it is internally buffered for lower output impedance and has a very wide frequency response, typically from 2Hz to 20kHz with the worst case -3dB response from 5Hz to 10kHz. It can measure up to 150g and has a high resonance frequency at 35kHz. The linearity is typically 0.1% with a maximum deviation of 1%. Its output can provide an analog voltage directly analogous to the acceleration that it measures.

The ACH-01-03 version of the accelerometer is supplied with a shielded cable and a convenient connector. **Photo 2** shows the accelerometer with its connector pin-out.

Acceleration, Velocity, and Excursion

To understand the operation of the system, it will be very helpful to refresh the memory with the basics of the simple harmonic motion theory. This was included in the previously mentioned article (see Sources) and will not be repeated here.

The conclusion is that if the acceleration is known, the velocity can be calculated by the integration of the acceleration signal and the excursion by the integration of the velocity signal over time.

A very simple electronic circuit can be used to perform this mathematical integration based on an op-amp, a capacitor and a resistor, similar to an active low pass filter with a very low cut-off frequency. The analog processor of the amplifier is based on two integrator circuits, one for the calculation of the velocity and the second for the excursion.

The Block Diagram

The block diagram of the AP-2 amplifier is shown in **Figure 1**. The first stage buffers the signal from the accelerometer and the current source bias the accelerometer's internal JFET. Using a rotary switch it can provide different gains from 0dB to 60dB.

Next, the signal goes to a 50Hz or 60Hz notch filter to suppress the mains interference. This notch filter can be bypassed with a switch accessible from the front panel. Then the signal continues to the second-order high-pass and low-pass filters, which have user selectable cut-off frequencies.

The first integrator follows to calculate the velocity of the signal while an additional integrator calculates the excursion (or displacement). Each output signal can be adjusted separately with multi-turn trimmers to calibrate the amplifier in accordance with the specific sensitivity of the ACH accelerometer. Finally, the signal is buffered before it goes to the output, which is monitored by the Clip indication circuit to prevent any possible overload of the amplifier that could happen if excessive gain is used.

The Electronic Diagram

The detailed electronic schematics of the amplifier are shown in **Figures 2–5**. A description of the circuit operation follows:

At the input of the circuit, a three pins connector exists for the connection of the ACH accelerometer. One pin is used to supply the voltage, the middle pin is the output of the accelerometer and the third is the Ground pin.

FB1 is a ferrite bead inductor offering electromagnetic induction (EMI) protection from several megahertz (MHz) up to gigahertz (GHz) by suppressing any noise coming to the amplifier from the long cable of ACH accelerometer.

The supply voltage for the ACH-01 should have very low ripple and noise and the RC low pass filter consisting of the R13-C1 filters additionally the regulated voltage.

Transistor Q1 with the R23, R28 and LED bias the internal JFET of the ACH with a constant current of about 26μA and can be disconnected by removing the jumper "Adj". This is useful during the calibration, as we will see later.

Capacitor C3 filters additionally the High Frequencies to reduce noise and C4 decouples the DC from entering the amplifier input. Resistor R21 provides a high impedance load for the ACH and bias the IC1A.

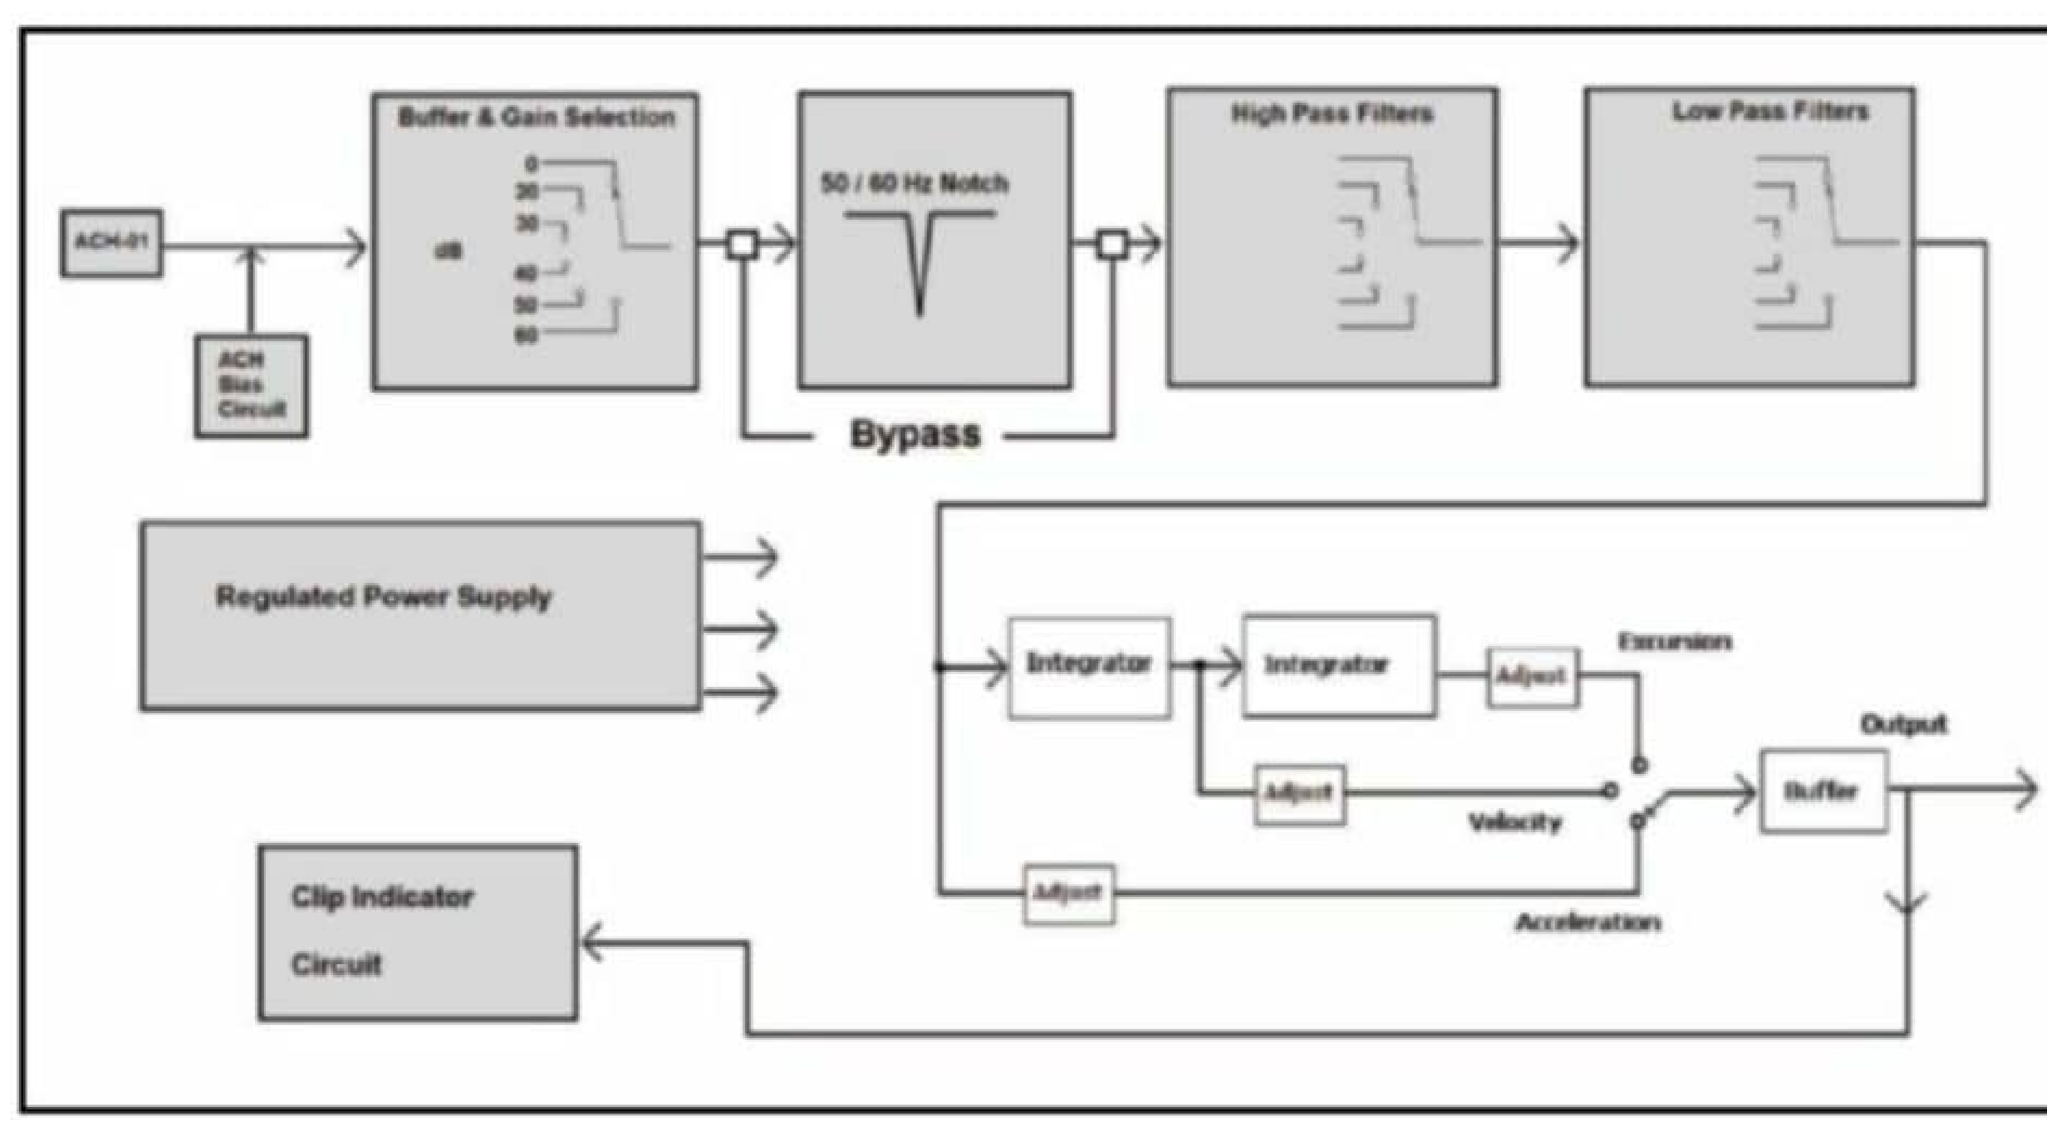

Figure 1: This is the block diagram of the AP-2 amplifier.

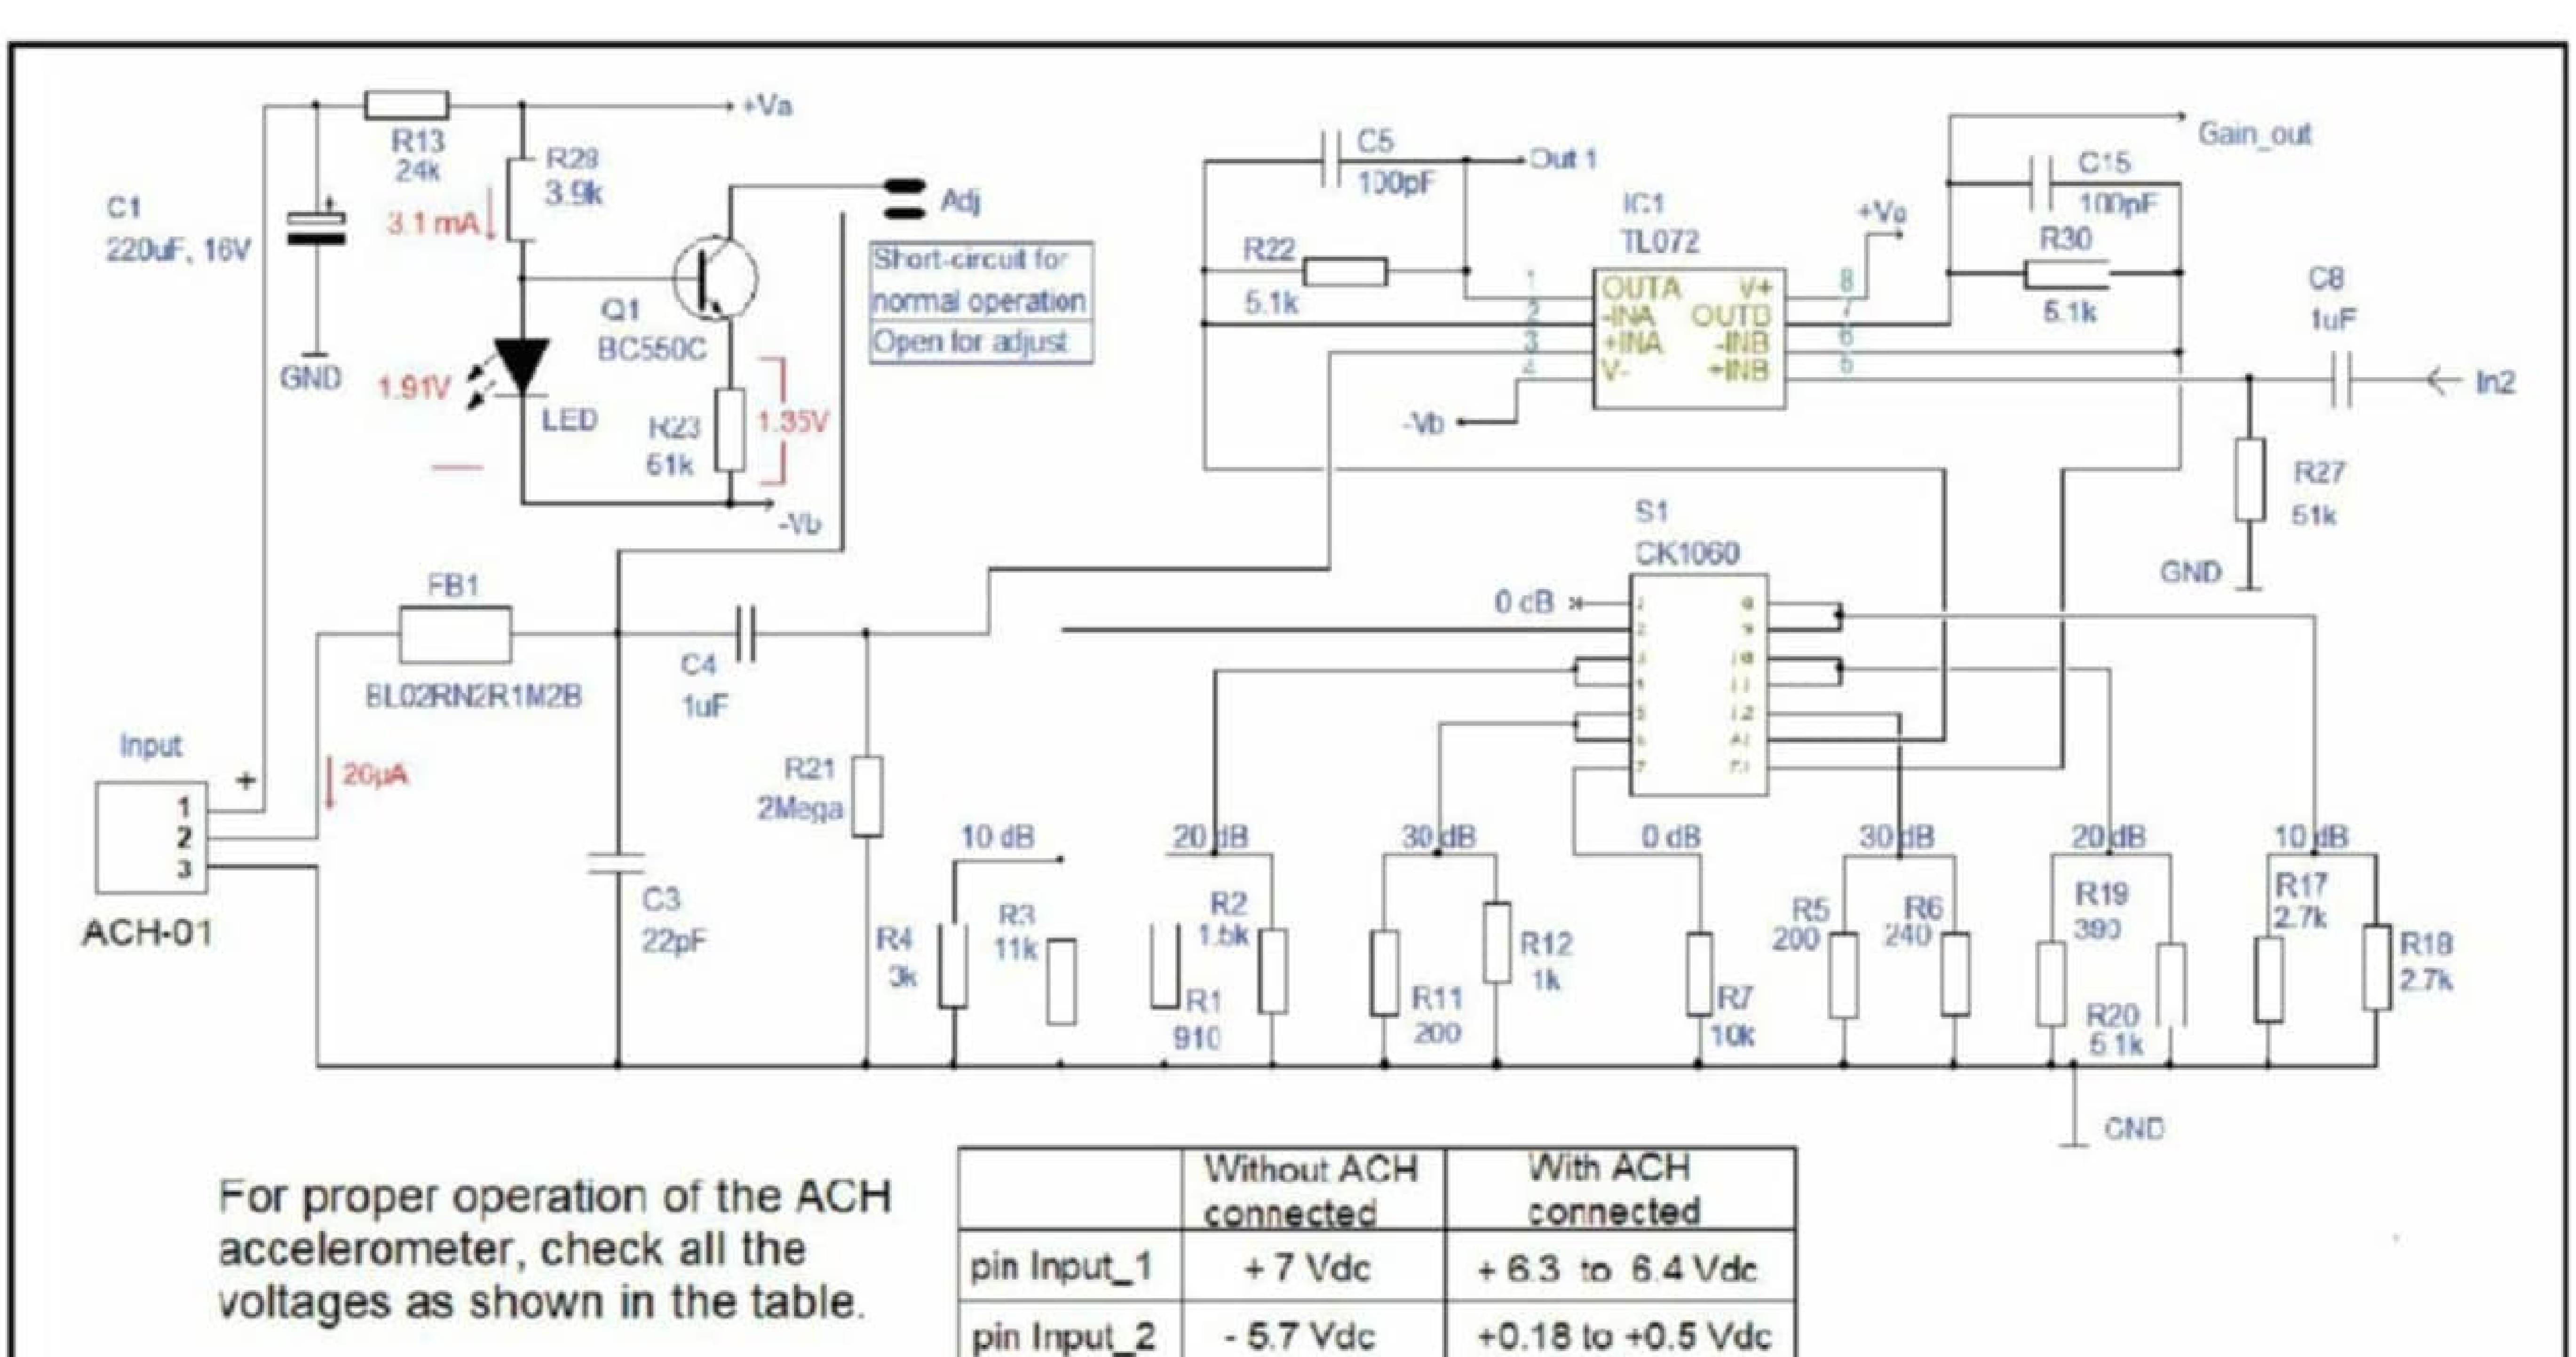

	Without ACH connected	With ACH connected
pin Input_1	+ 7 Vdc	+ 6.3 to 6.4 Vdc
pin Input_2	- 5.7 Vdc	+0.18 to +0.5 Vdc

Figure 2: Here are the input and gain circuits.

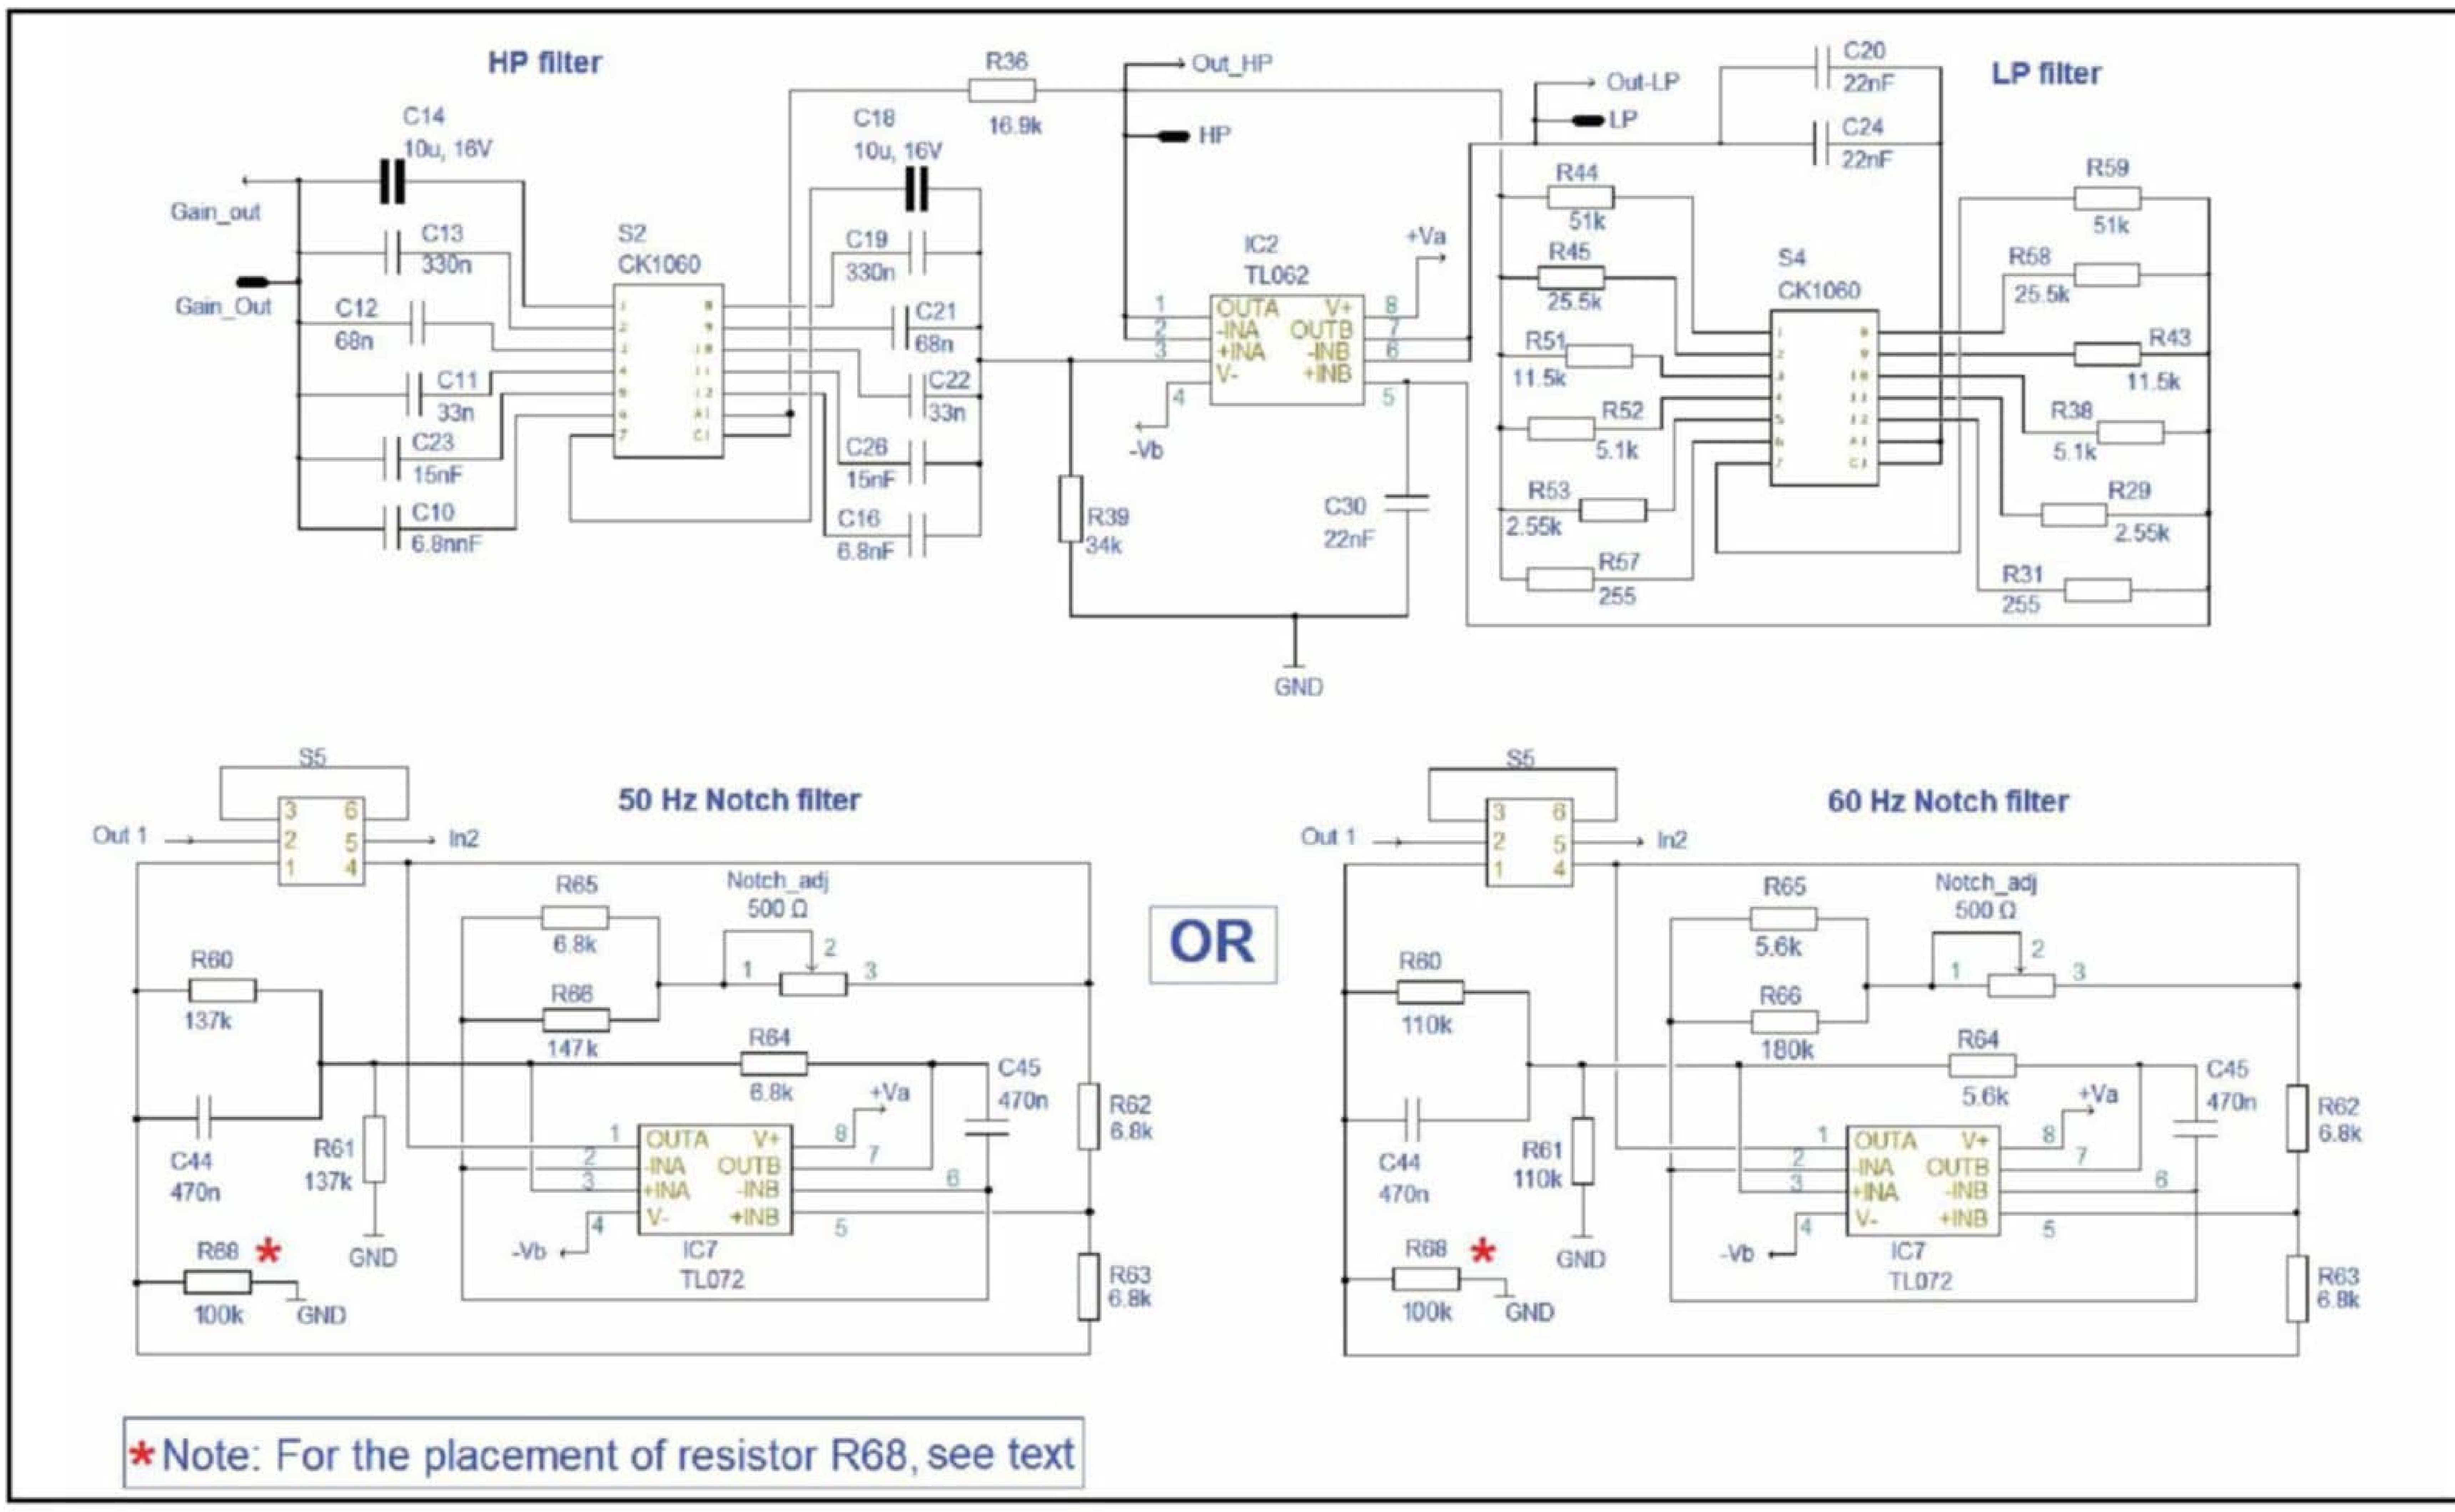

Figure 3: This diagram shows the filters.

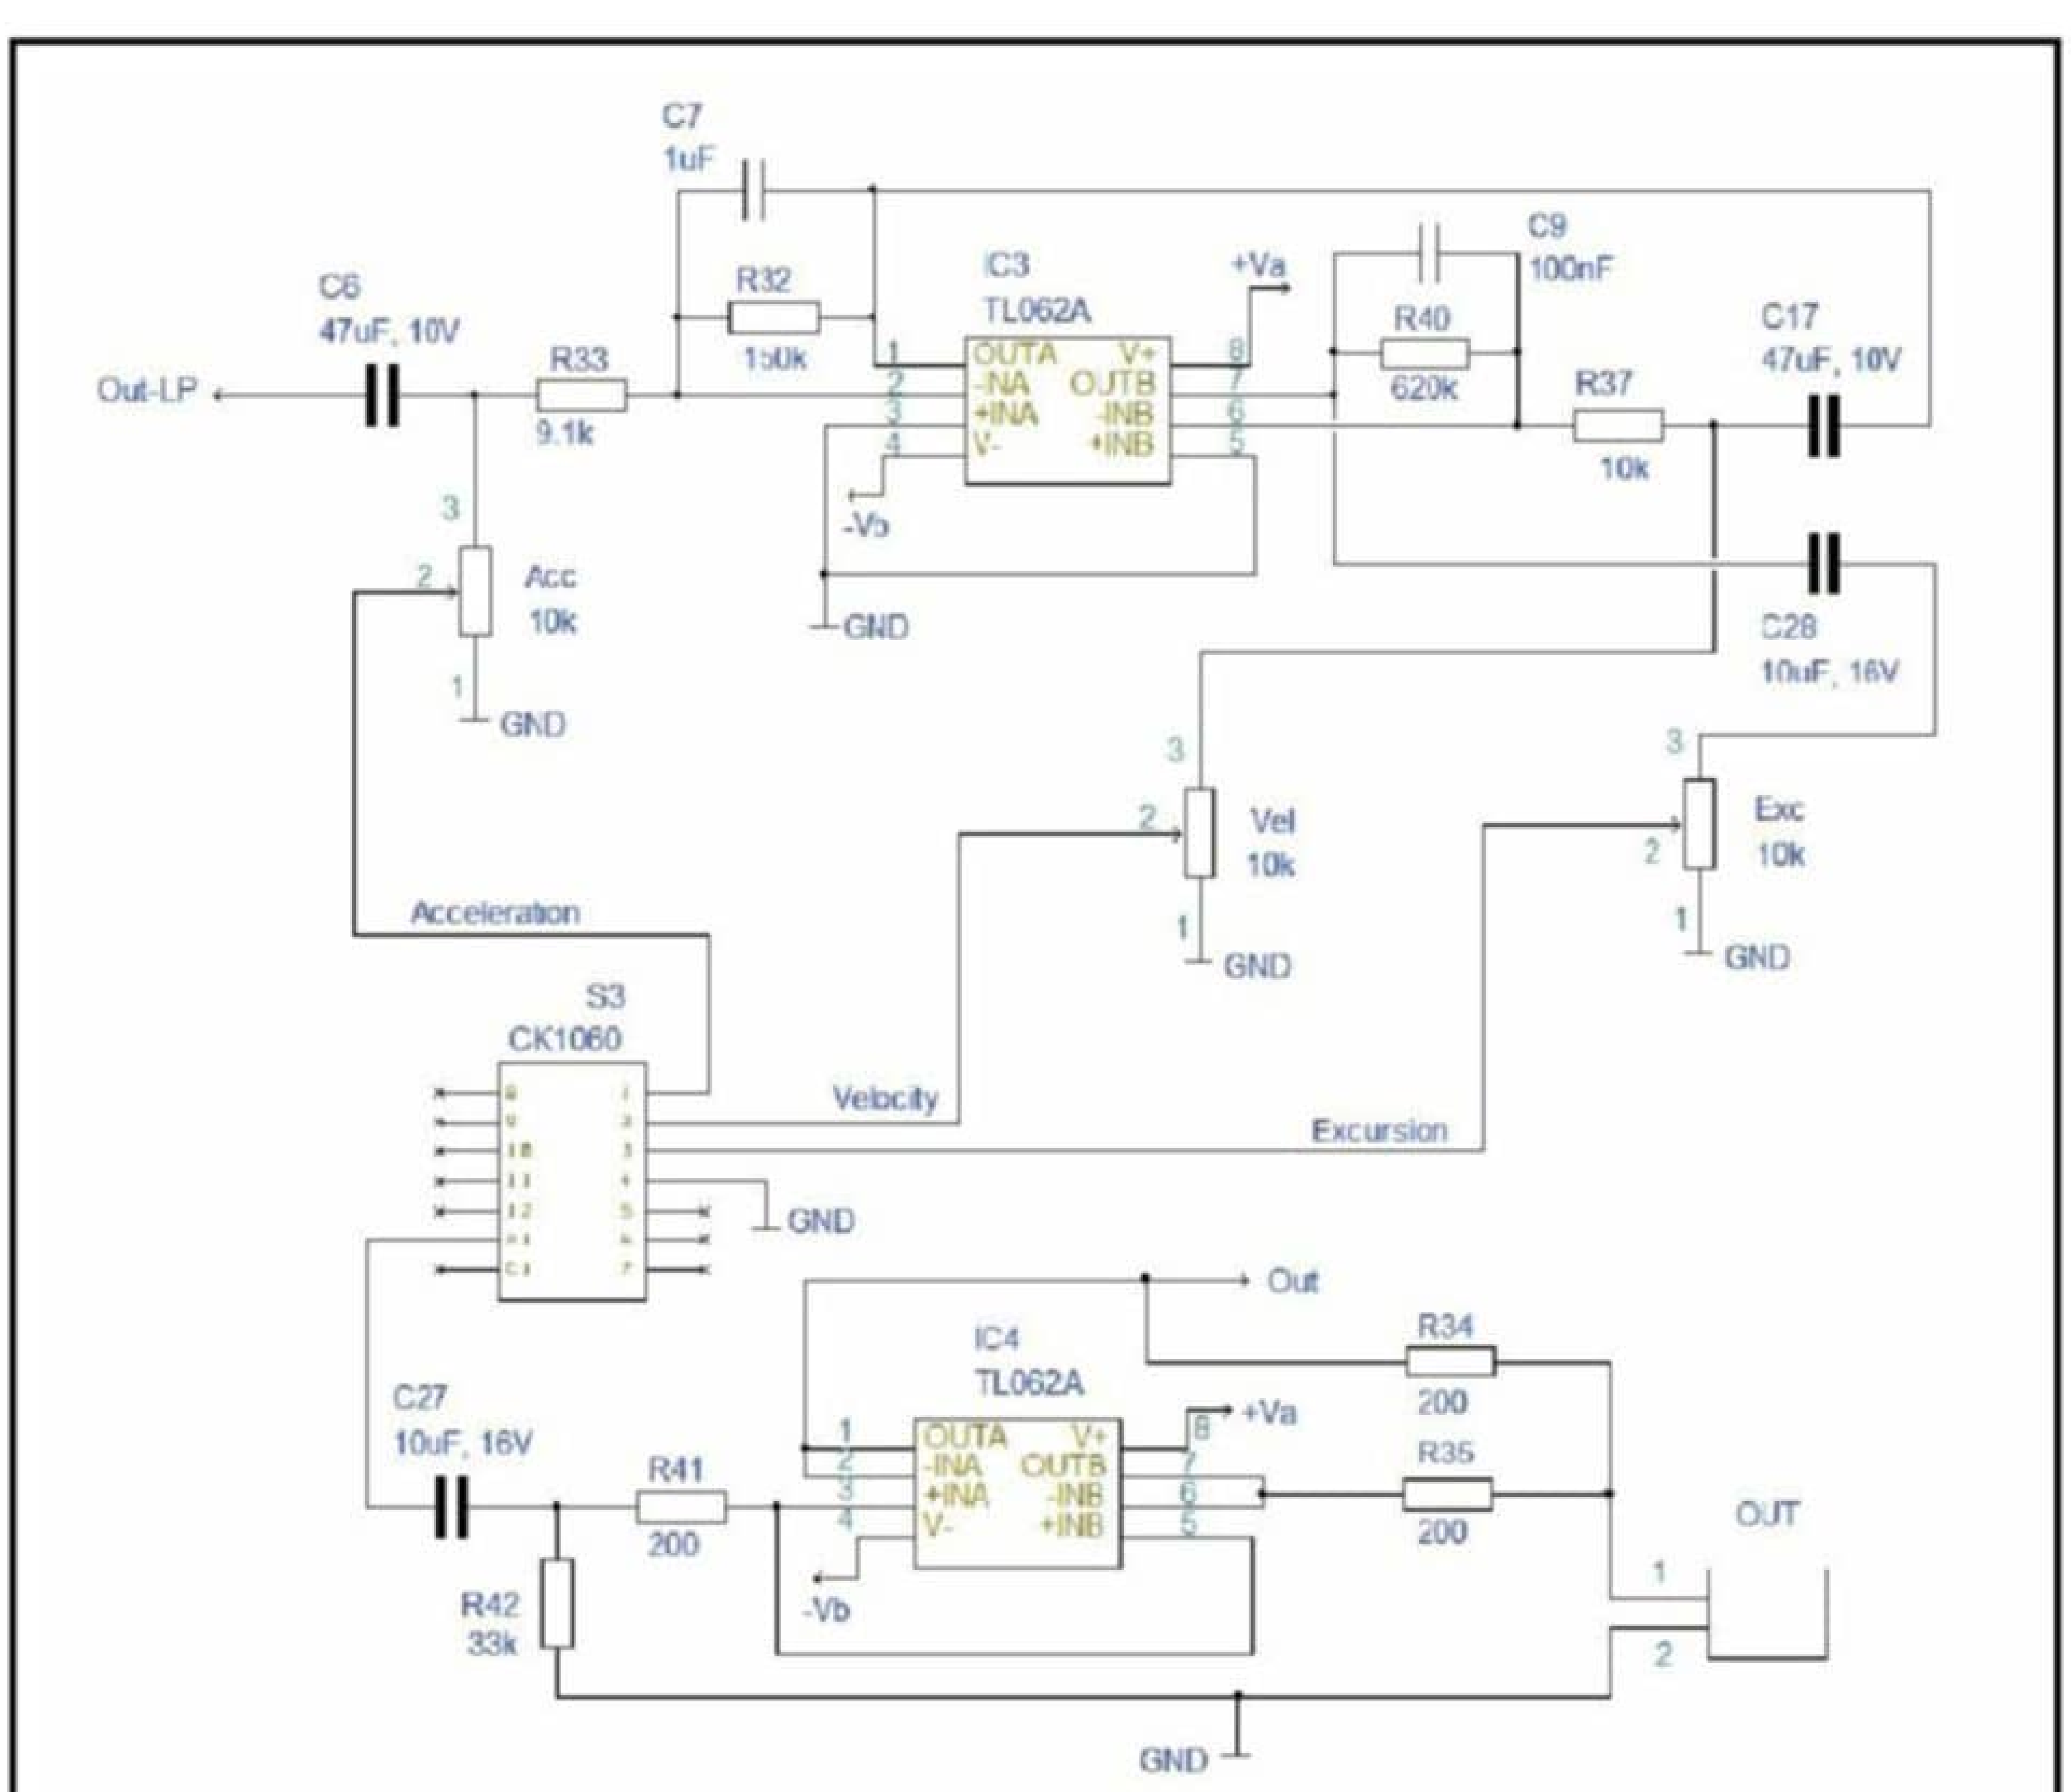

Figure 4: Here is a schematic for the integrators.

The op-amp IC1 amplifies the accelerometer signal and switch S1 selects the gain from 0 to 60dB. The higher gains will be used when the output of the accelerometer is lower, for example when the vibration of a loudspeaker panel is measured.

For achieving 60dB with adequate bandwidth, the amplifier uses equal gains provided by the two op-amps of IC1 that are connected in series.

The first stage could have one of the following six gains: 0dB / 10dB / 20dB / 20dB / 30dB / 30dB while the second stage the following: 0 / 10dB / 10dB / 20dB / 20dB / 30dB.

Combining both stages, the total gain becomes 0 / 20dB / 30dB / 40dB / 50dB / 60dB. The resistors R3, R4, …. R17, R18 that are connected to S1 adjust the different gains of the two op-amps.

The notch filter around IC7 can be used to reduce the mains interference that is usually picked up by the ACH-01. It is shown in the bottom of Figure 3 and can be inserted between the first and the second gain stage using the switch S5. More details will be later in the article.

The components around op-amp IC2A form the second-order high-pass (HP) filter. Switch S2 selects the cut-off frequency between 3Hz, 20Hz, 100Hz, 200Hz, 450Hz and 1kHz by changing the values of both capacitors of the filter.

The components around op-amp IC2B form the second-order low-pass (LP) filter. Switch S4 selects the cut-off frequency between 100Hz, 200Hz, 450Hz, 1kHz, 2kHz, or 20kHz by changing the values of both resistors of the filter.

The maximum bandwidth of the amplifier is restricted to 20kHz in order to reduce the additional noise due to higher frequencies.

The op-amp IC3A with R32, R33 and C7 form the first integrator, which gives at its output the velocity.

The op-amp IC3B with R37, R40 and C9 form the second integrator that gives at its output the excursion (or displacement) of the acceleration signal.

The three trimmers "Acc", "Vel", and "Exc" are used for the calibration of the output signal: "Acc" for the acceleration, "Vel" for the velocity, and "Exc" for the excursion.

Switch S3 selects what will be connected to the output of the amplifier and IC4 buffers the output signal. The two op-amps of IC4 are connected in parallel to increase the drive capability of the amplifier in case it is needed.

For the supply of the amplifier, I used an external power supply pack with a nominal output voltage of 9V AC/0.5A. The consumption of the amplifier is very low, typically at ±20mA, and the actual output voltage goes up to about 10V AC to 11V AC. Voltage higher than 12V AC at the input power supply connector "V AC" should be avoided because otherwise the regulators could be overheated.

FB2 is a ferrite bead inductor offering electromagnetic induction (EMI) protection from several megahertz up to gigahertz by suppressing any noise coming from the power supply. F1 is a slow-blow fuse, mounted on the PCB to protect the external power supply from short circuits within the device.

The AC voltage is rectified by diodes D3 and D4 and through the RC filters formed by R50-C38, R48-C41, R56-C42 and R46-C39, the positive voltage goes to IC5 (LM317L) and the negative to the IC6 (LM337L). Their output voltage is set at ±7V DC by the resistors R47-R54 and R49-R55.

Resistor R15 and Zener Z1 drive the "ON" LED, which indicates that the amplifier is powered on. The LED will only illuminate if the voltage is above a certain threshold (defined mainly by the Zener voltage), indicating that both the regulated voltages are good.

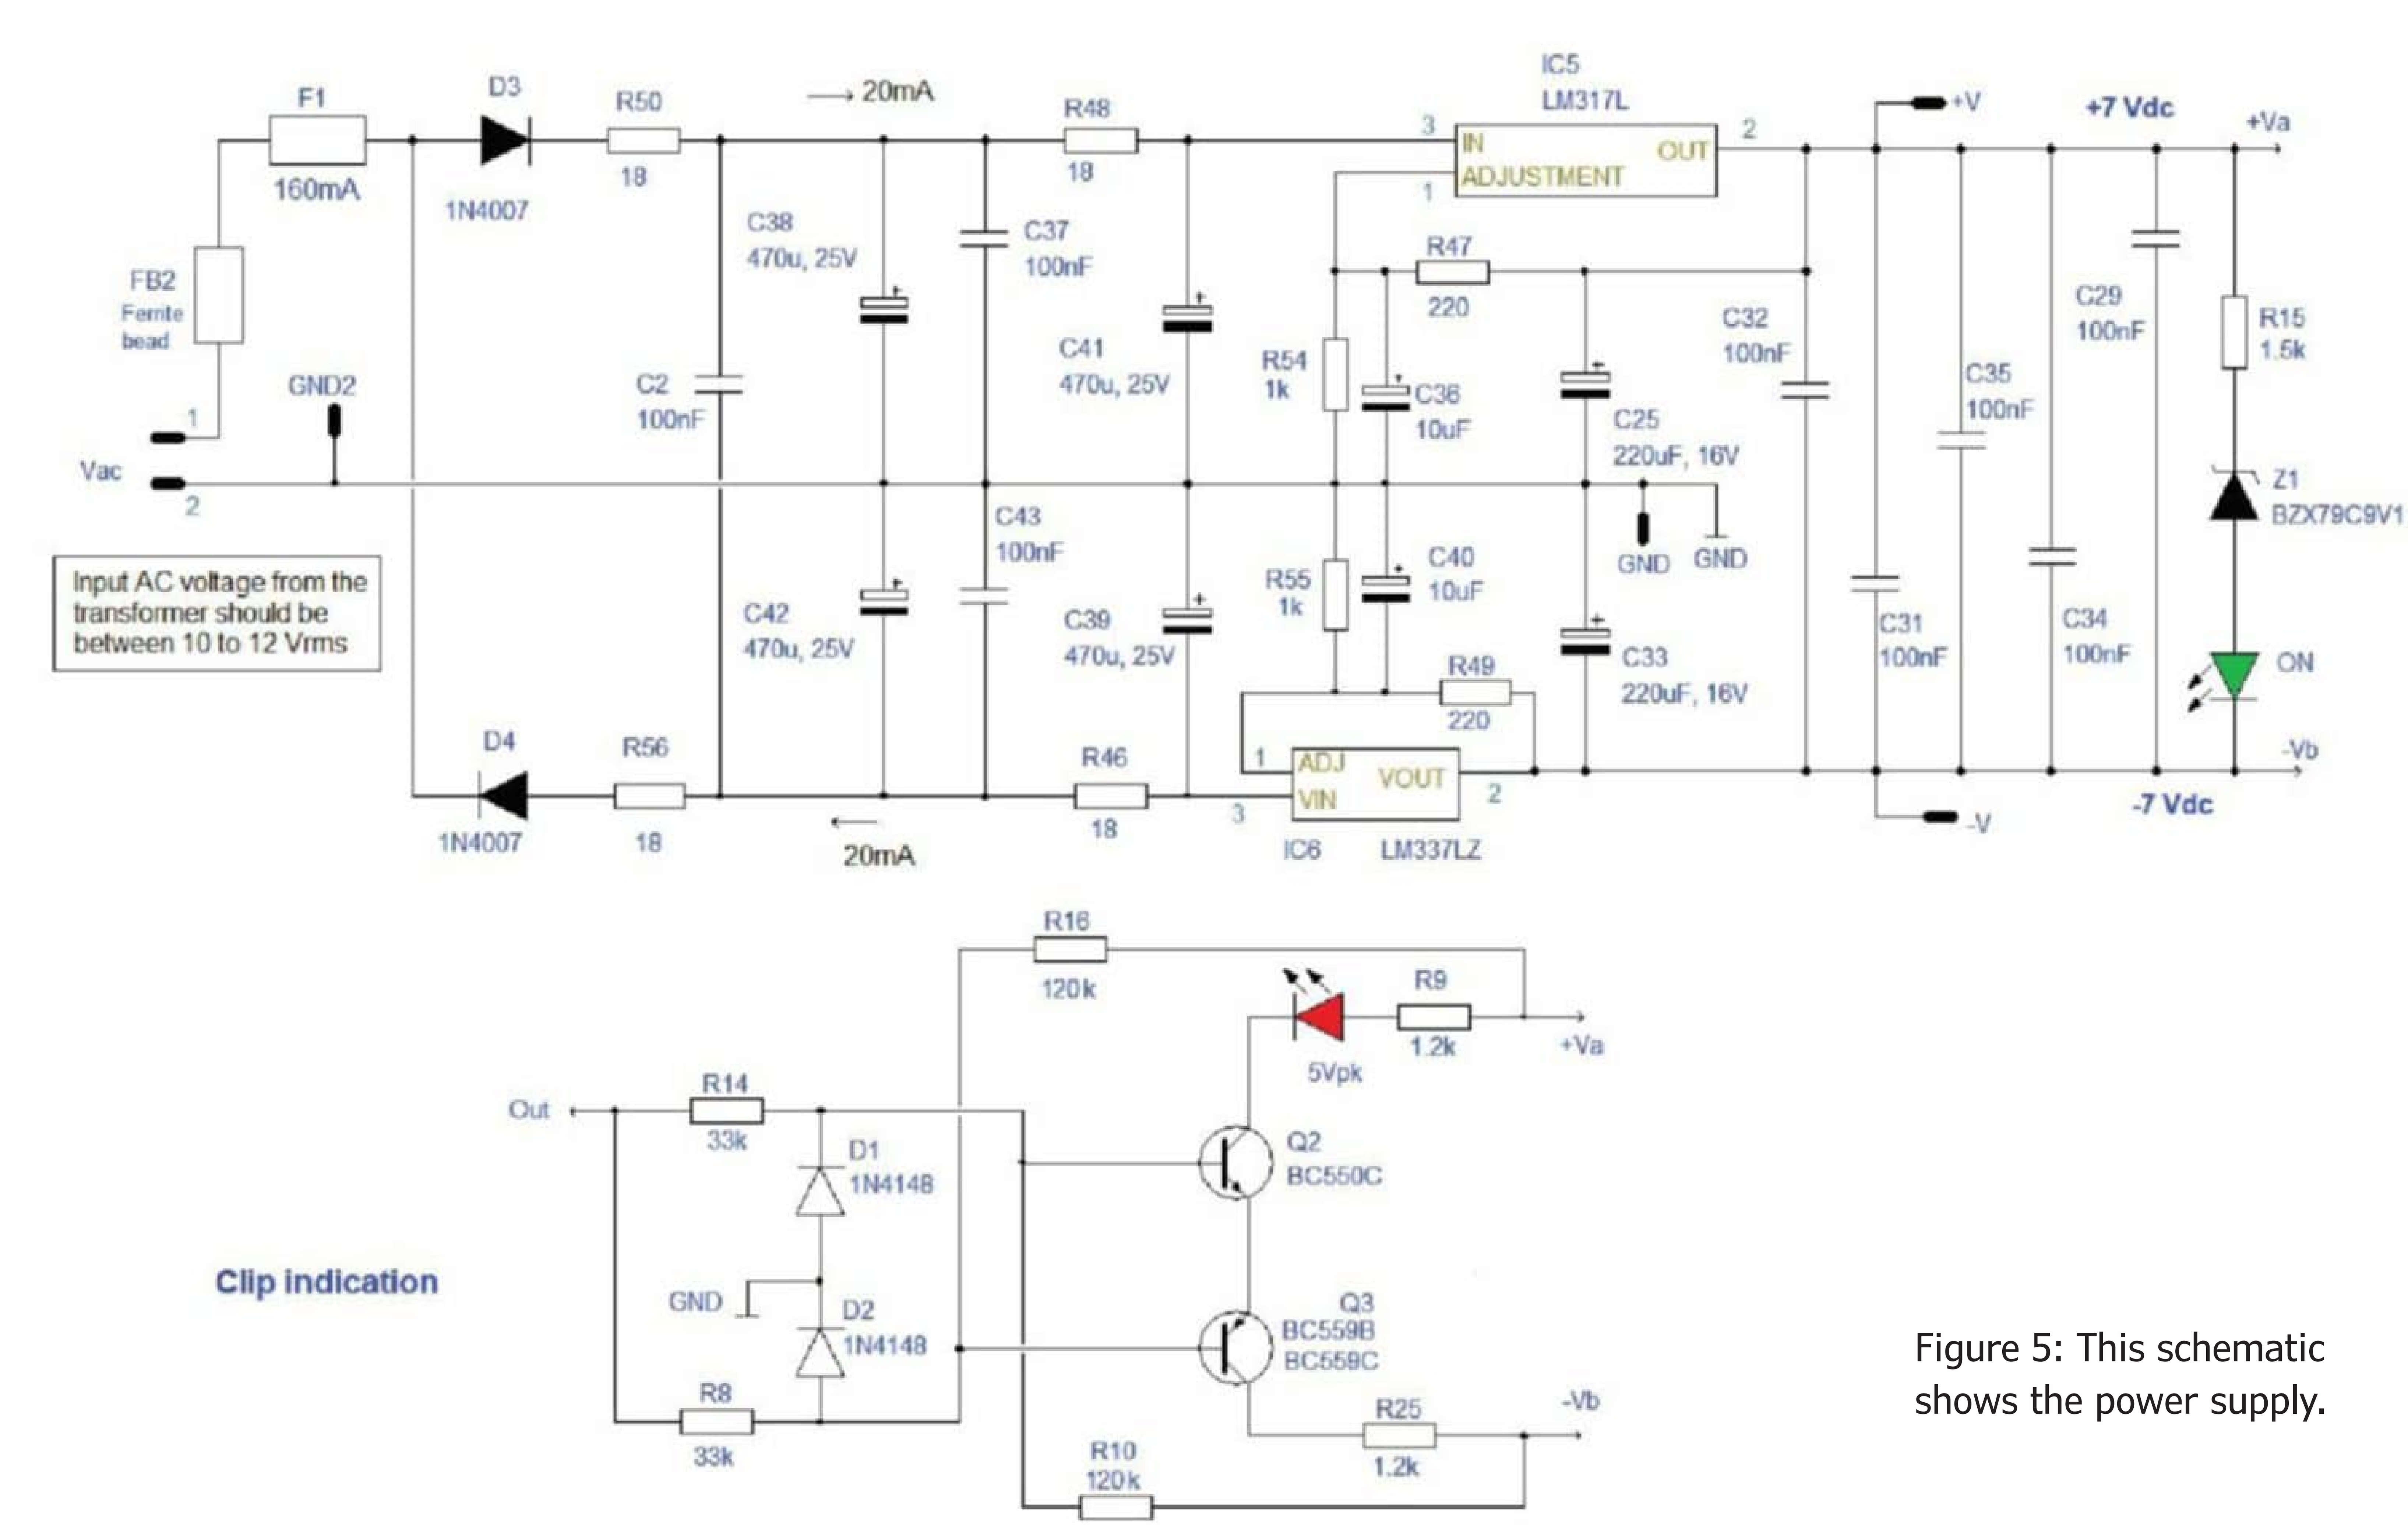

Figure 5: This schematic shows the power supply.

Finally, the circuit around Q2 and Q3 form a clip indicator, which is connected directly at the output of the amplifier. The "5Vpk" LED will illuminate when the output voltage is close to ±5V peak. The circuit was presented in an article by C.W. Russel in *Audio Amateur* (see Resources).

I used the TL072A op-amp for IC1 and IC7 because of its lower noise. In all the other places I used the lower consumption TL062A op-amp because its higher noise is not critical for the operation of the circuit. I recommend following this suggestion, otherwise the TO-92 regulators might overload and overheat since their heat dissipation is limited.

The PCB Assembly and the Construction of the Amplifier

I used the Design Spark PCB software to design the PCB. The outcome was a 100mm × 150mm double-sided PCB with ground planes on the top and bottom layers. The Parts List shows the components used for the project. The placement of the components is shown in **Figure 6** and the assembled PCB is shown in **Photo 3**.

To make the assembly of the amplifier easier and give it a more professional touch, I also designed an engraved front panel, as shown in **Photo 4**, using the program "Front Panel Designer" (see Resources).

The PCB contains all the components of the amplifier. There is no position for the resistor R68 in the PCB, but **Photo 5** shows how it is placed.

The front panel replaces the top cover of an aluminum Hammond box, which is used for the housing of the PCB offering protection from the external EM fields. The four switches are fastened directly on the front panel and support the whole PCB. The output connector and the supply connector are placed directly on the front panel.

Parts List

Reference Name	Component	Description
D1, D2	1N4148	Signal diode, DO-35
D3, D4	1N4007	Rectifiers
Z1	BZX79C9V1	Zener diode, 9.1V
F1	Fuse	160mA, slow blow, PCB mount
Q1, Q2	BC550C	NPN Bipolar transistor
Q3	BC559B or BC559C	PNP Bipolar transistor
FB1, FB2	Ferrite bead	EMI suppression
C1, C25, C33	220µF, 16V	Electrolytic capacitor, D=8, H=11.5
C2, C29, C31, C32,		
C34, C35, C37, C43	100nF	Multilayer ceramic capacitors X7R
C3	22pF	Ceramic C0G, 5.0mm
C4, C7, C8	1µF	Film capacitors
C5, C15	100pF	Ceramic C0G, 5.0mm
C6, C17	47µF, 10V	Bipolar electrolytic capacitor
C9	100 nF	Film capacitors
C10, C16	6.8nF	Film capacitors
C11, C22	33nF	Film capacitors
C12, C21	68nF	Film capacitors
C13, C19	330nF	Film capacitors
C14, C18, C27, C28	10µF, 16V	Bipolar electrolytic capacitor
C20, C24, C30	22nF	Film capacitors
C23, C26	15nF	Film capacitors
C36, C40	10µF, 35V	Electrolytic capacitor
C38, C39, C41, C42	470µF, 25V	Electrolytic capacitor, D=10, H=13
C44, C45	470nF	COG capacitors
R1	910Ω	Metal film, 0.25W, 1%
R2, R15	1.5kΩ	Metal film, 0.25W, 1%
R3	11kΩ	Metal film, 0.25W, 1%
R4	3kΩ	Metal film, 0.25W, 1%
R5, R11, R34,		
R35, R41	200Ω	Metal film, 0.25W, 1%
R6	240Ω	Metal film, 0.25W, 1%
R7, R37	10kΩ	Metal film, 0.25W, 1%
R8, R14, R42	33kΩ	Metal film, 0.25W, 1%
R9, R25	1.2kΩ	Metal film, 0.25W, 1%
R10, R16	120kΩ	Metal film, 0.25W, 1%
R12, R54, R55	1kΩ	Metal film, 0.25W, 1%
R13	24kΩ	Metal film, 0.25W, 1%
R17, R18	2.7kΩ	Metal film, 0.25W, 1%
R19	390Ω	Metal film, 0.25W, 1%
R20, R22, R30,		
R38, R52	5.1kΩ	Metal film, 0.25W, 1%
R21	2MΩ	Metal film, 0.25W, 1%
R23, R27, R44, R59	51kΩ	Metal film, 0.25W, 1%
R28	3.9kΩ	Metal film, 0.25W, 1%
R29, R53	2.55kΩ	Metal film, 0.25W, 1%

Reference Name	Component	Description
R31, R57	255Ω	Metal film, 0.25W, 1%
R32	150kΩ	Metal film, 0.25W, 1%
R33	9.1kΩ	Metal film, 0.25W, 1%
R36	16.9kΩ	Metal film, 0.25W, 1%
R39	34kΩ	Metal film, 0.25W, 1%
R40	620kΩ	Metal film, 0.25W, 1%
R43, R51	11.5kΩ	Metal film, 0.25W, 1%
R45, R58	25.5kΩ	Metal film, 0.25W, 1%
R46, R48, R50, R56	18Ω	Metal film, 0.25W, 1%
R47, R49	220Ω	Metal film, 0.25W, 1%
R60, R61	137kΩ for 50Hz 110kΩ for 60Hz	Metal film, 0.25W, 1%
R62, R63	6.8kΩ	Metal film, 0.25W, 1%
R64, R65	6.8kΩ for 50Hz 5.6kΩ for 60Hz	Metal film, 0.25W, 1%
R66	147kΩ for 50Hz 180kΩ for 60Hz	Metal film, 0.25W, 1% Metal film, 0.25W, 1%
R68	100kΩ	Metal film, 0.25W, 1%
Acc, Vel, Exc	10kΩ	3006P Trimmer side adjust
Notch_adj	500Ω	3296Y Trimmer top adjust
IC1, IC7	TL072A	Operational Amplifier
IC2, IC3, IC4,	TL062A	Operational Amplifier
IC5	LM317L	Linear voltage regulator 100mA
IC6	LM337L	Linear voltage regulator, 100mA
Input	3-pin header	provided with ACH-01
Adj	2-pin header	Header 2 position
S1, S2, S3, S4	Rotary Switch	Lorlin CK1060, 6 Position DP Lorlin CK1050, 6 Position DP
Knobs for the switches	4 pieces	Knob with external diameter 15mm or Knob with external diameter 20mm
Cap for the knobs	4 pieces	Matt or glossy
S5	100DP1T1B1M2QEH	Toggle switch DPDT panel mount
5Vpk	3mm LED	Red LED
LED	3mm LED	Green LED
ON	3mm LED	Green LED
Power supply connector	Panel Mount, plug	two-pole, 6.3mm
Output Connector	BNC Socket	Isolated from chassis
Ground terminal		Connection between PCB GND and Metal Box
Washers	Outer Diameter = 16mm Inner Diameter = 10mm Thickness = 1.5mm	S1, S2, S3, S4 (see instructions for details)
Metal Box	Enclosure	Hammond

After the PCB was assembled, I connected the power supply and I measured the voltages at several points, as shown in the electronic diagrams. Also, without the ICs in their sockets, I checked that the voltages at the supply pins of all ICs were correct.

It is very important to measure the voltages at the "Input" connector of the ACH with the accelerometer connected and verify that they are as shown in Figure 2. Only by doing so, is it assured that the ACH is biased and measures correctly.

It is very easy to put the ACH-01 cable connector upside down to its PCB connector. I recommend that when you have measured the input voltages and found they are correct, to put some indication at both connectors in order to avoid possible mistakes the next time. For better shielding, the ground point of the amplifier should be connected to the chassis with a short wire as shown in the bottom left side of Photo 3.

The Notch Filter

The notch is a Fliege filter topology with a fixed gain of 1 and Q equal to 10 based on a circuit presented in the book "Op-amps for everyone" (see Resources). Its great advantage is that the notch frequency can be adjusted using only one trimmer, the "Notch_adj" trimmer in our circuit, which is accessible from the front panel.

In accordance with the mains frequency, the center frequency of the notch filter can be set at 50Hz or 60Hz, using different components values, as shown in the bottom of Figure 3.

For accuracy in the notch frequency and maximum notch depth, it is necessary to match, as close as possible, the values of capacitors C44 and C45 and of the resistors R60 and R61. The matching should be done not only between the values of the two components but also as close as possible to their nominal values shown in the schematic. Matching resistors R62 and R63 is also critical but their absolute values are not.

COG capacitors are the best choice for C44 and C45. They have a dielectric that exhibits no change in capacitance with respect to time and voltage and a negligible change in capacitance with reference to ambient temperature making them suitable for resonant circuit applications where Q and stability of capacitance are required.

It is difficult and very expensive to find COG capacitors with tight tolerance (e.g., ±2.5%). Instead, their typical tolerance is about 10%, so probably the value of C44 and C45 will not be very close to its nominal value.

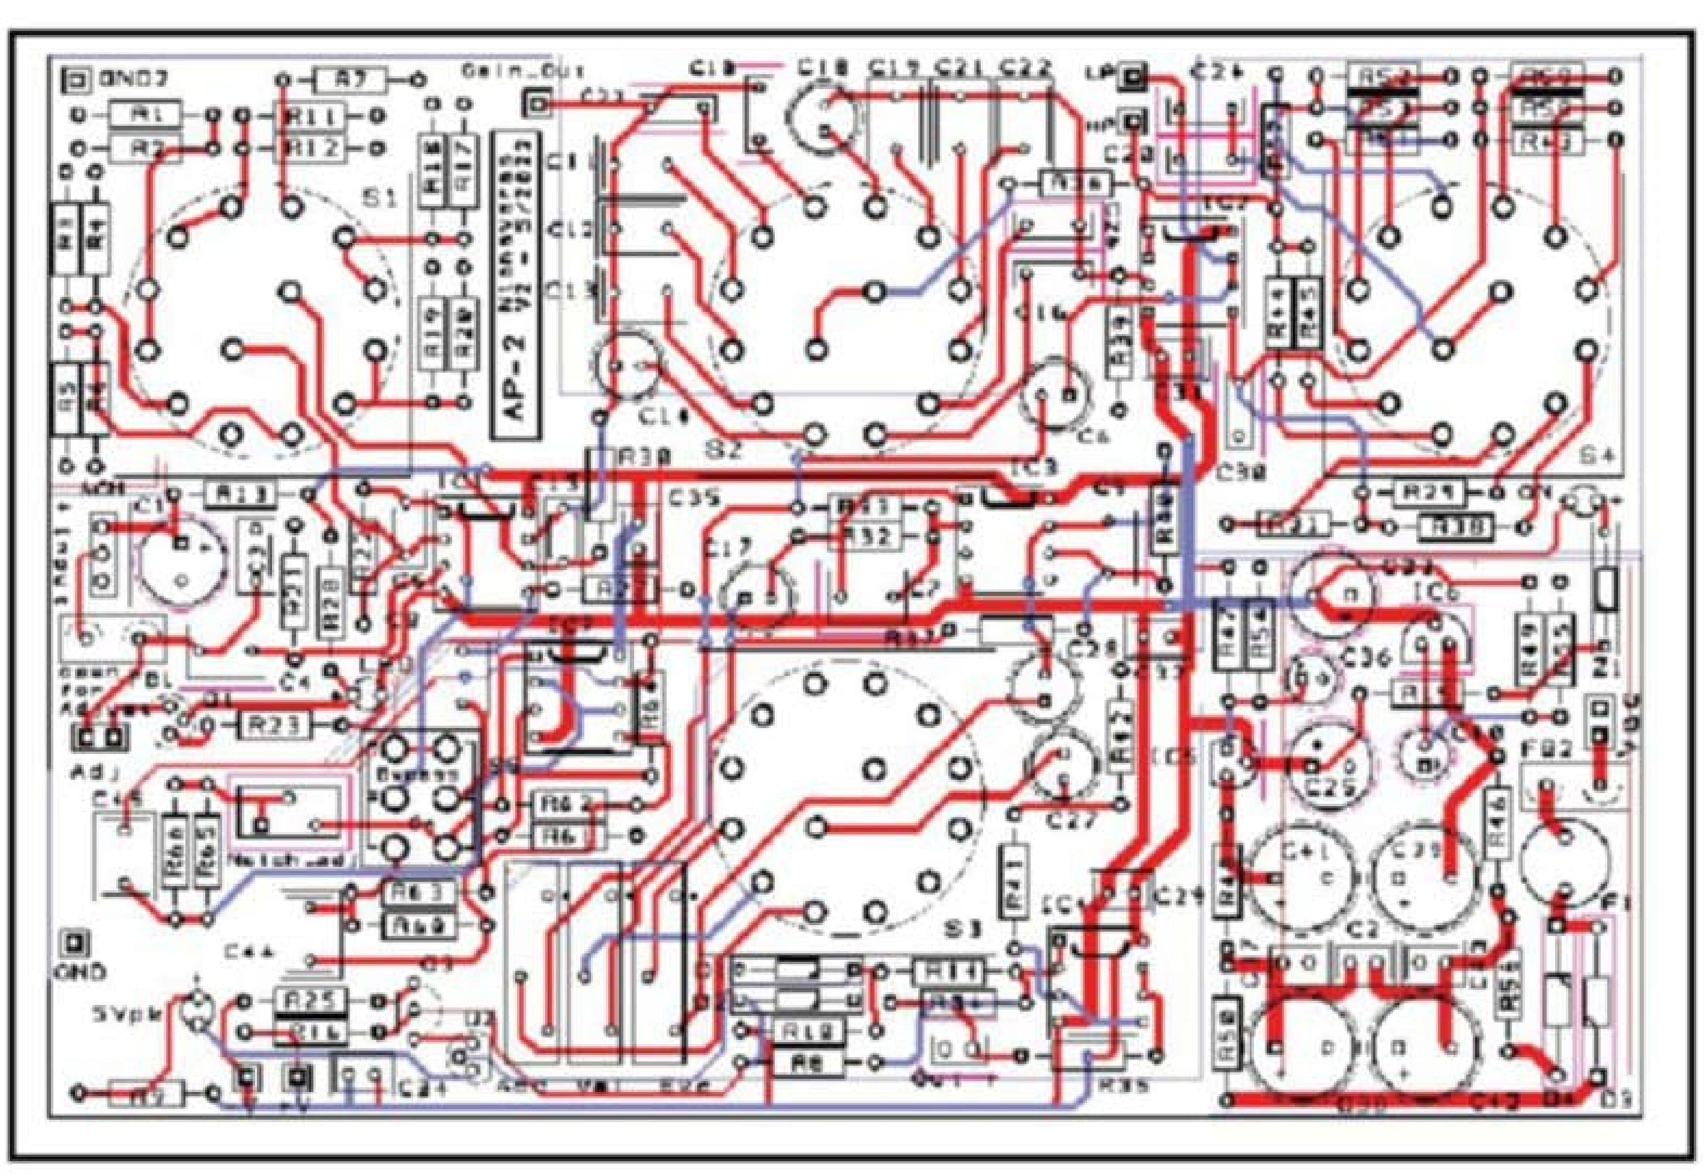

Figure 6: This image shows the PCB of the AP-2 amplifier and the placement of the components.

Photo 4: Here is the engraved front panel for the AP-2 amplifier.

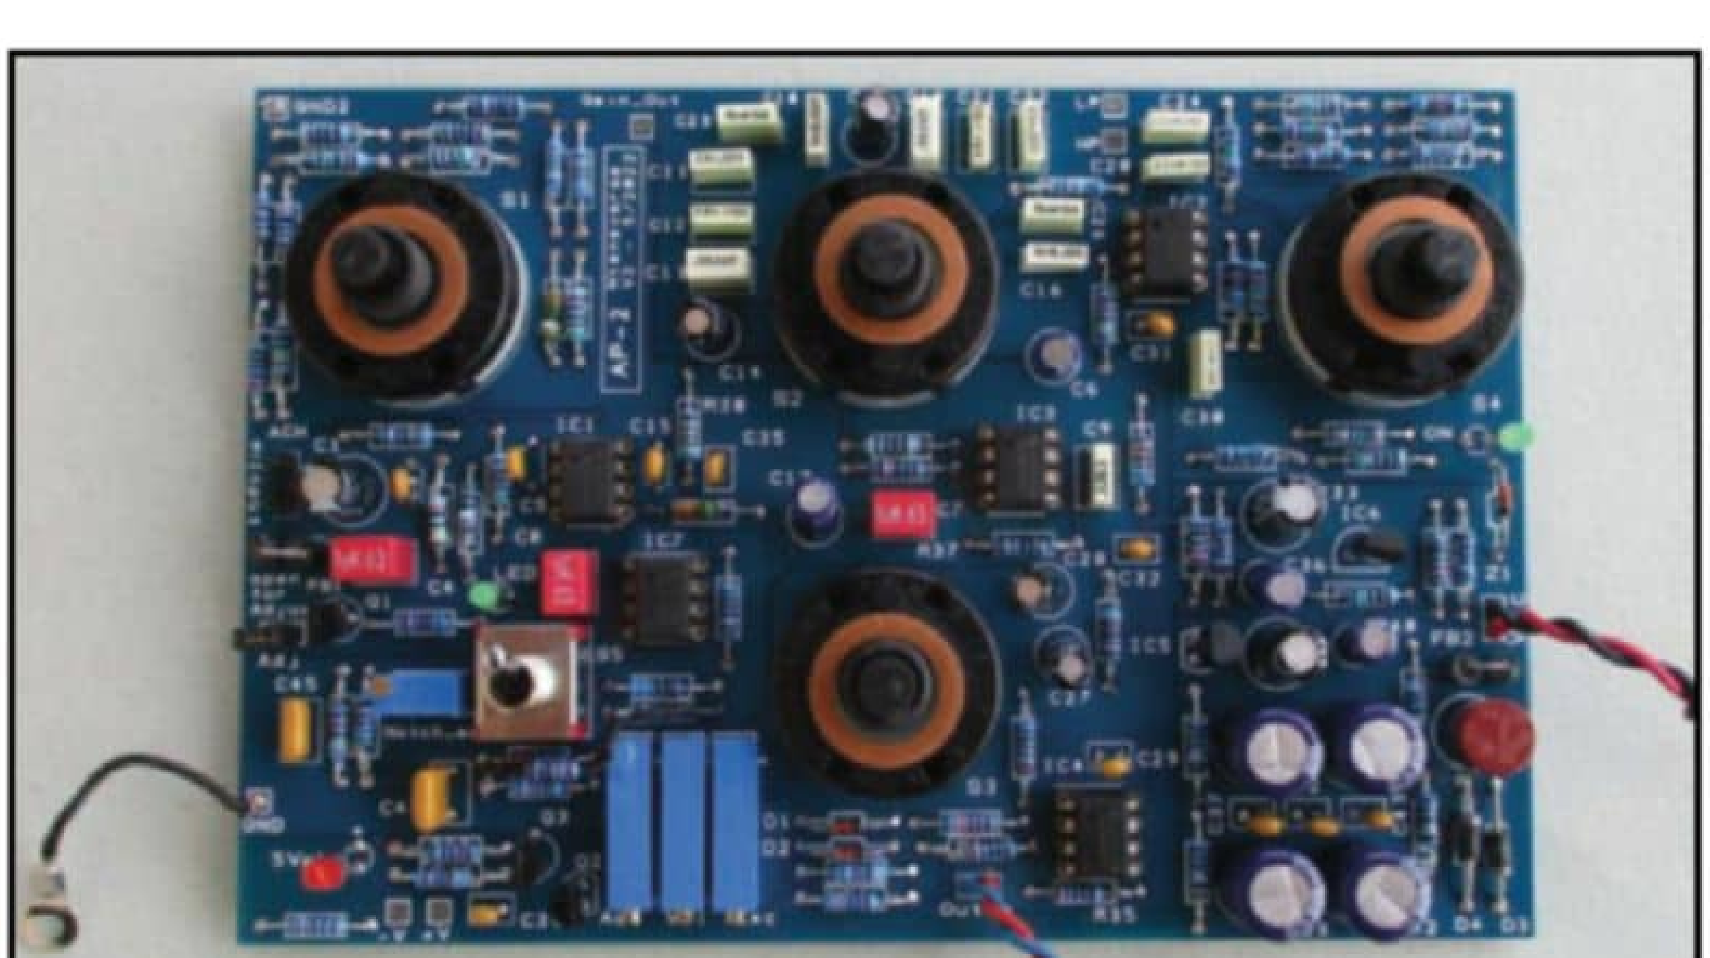

Photo 3: The assembled PCB is shown here.

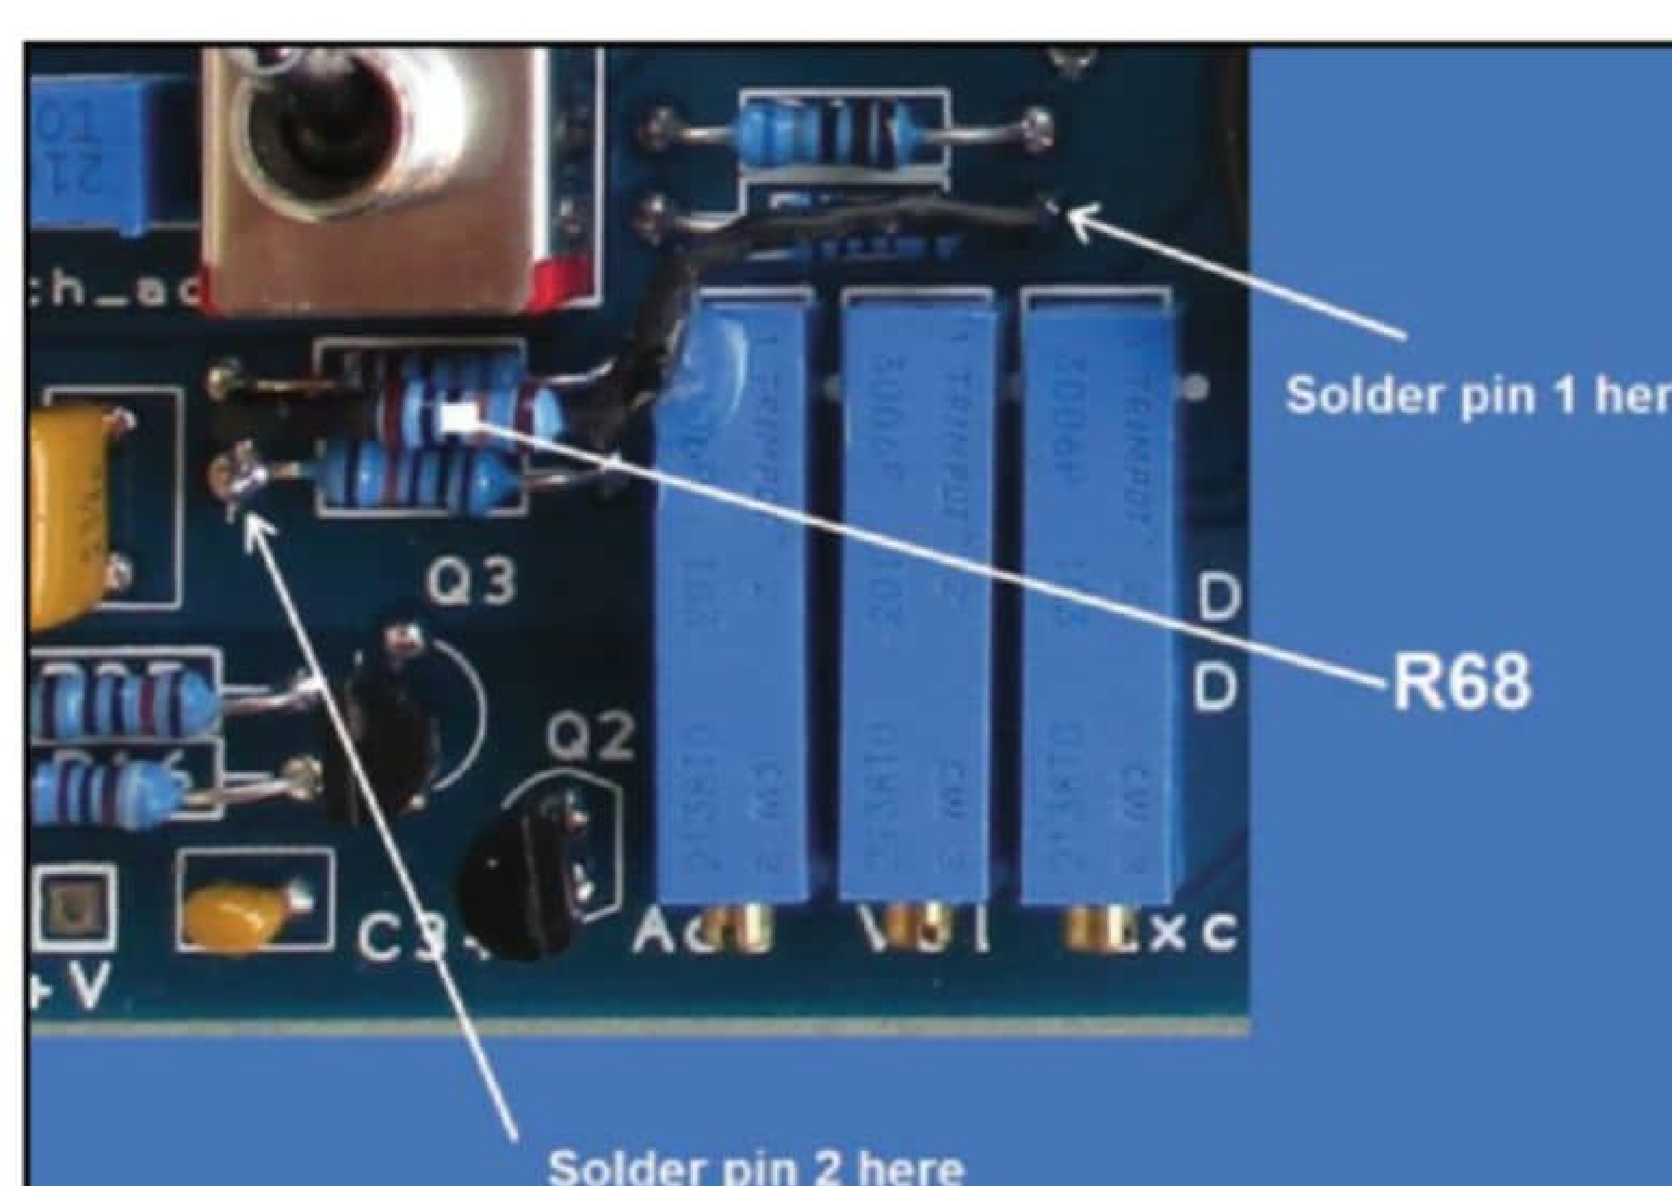

Photo 5: There is no position for the resistor R68 in the PCB, but this image shows how it should be placed.

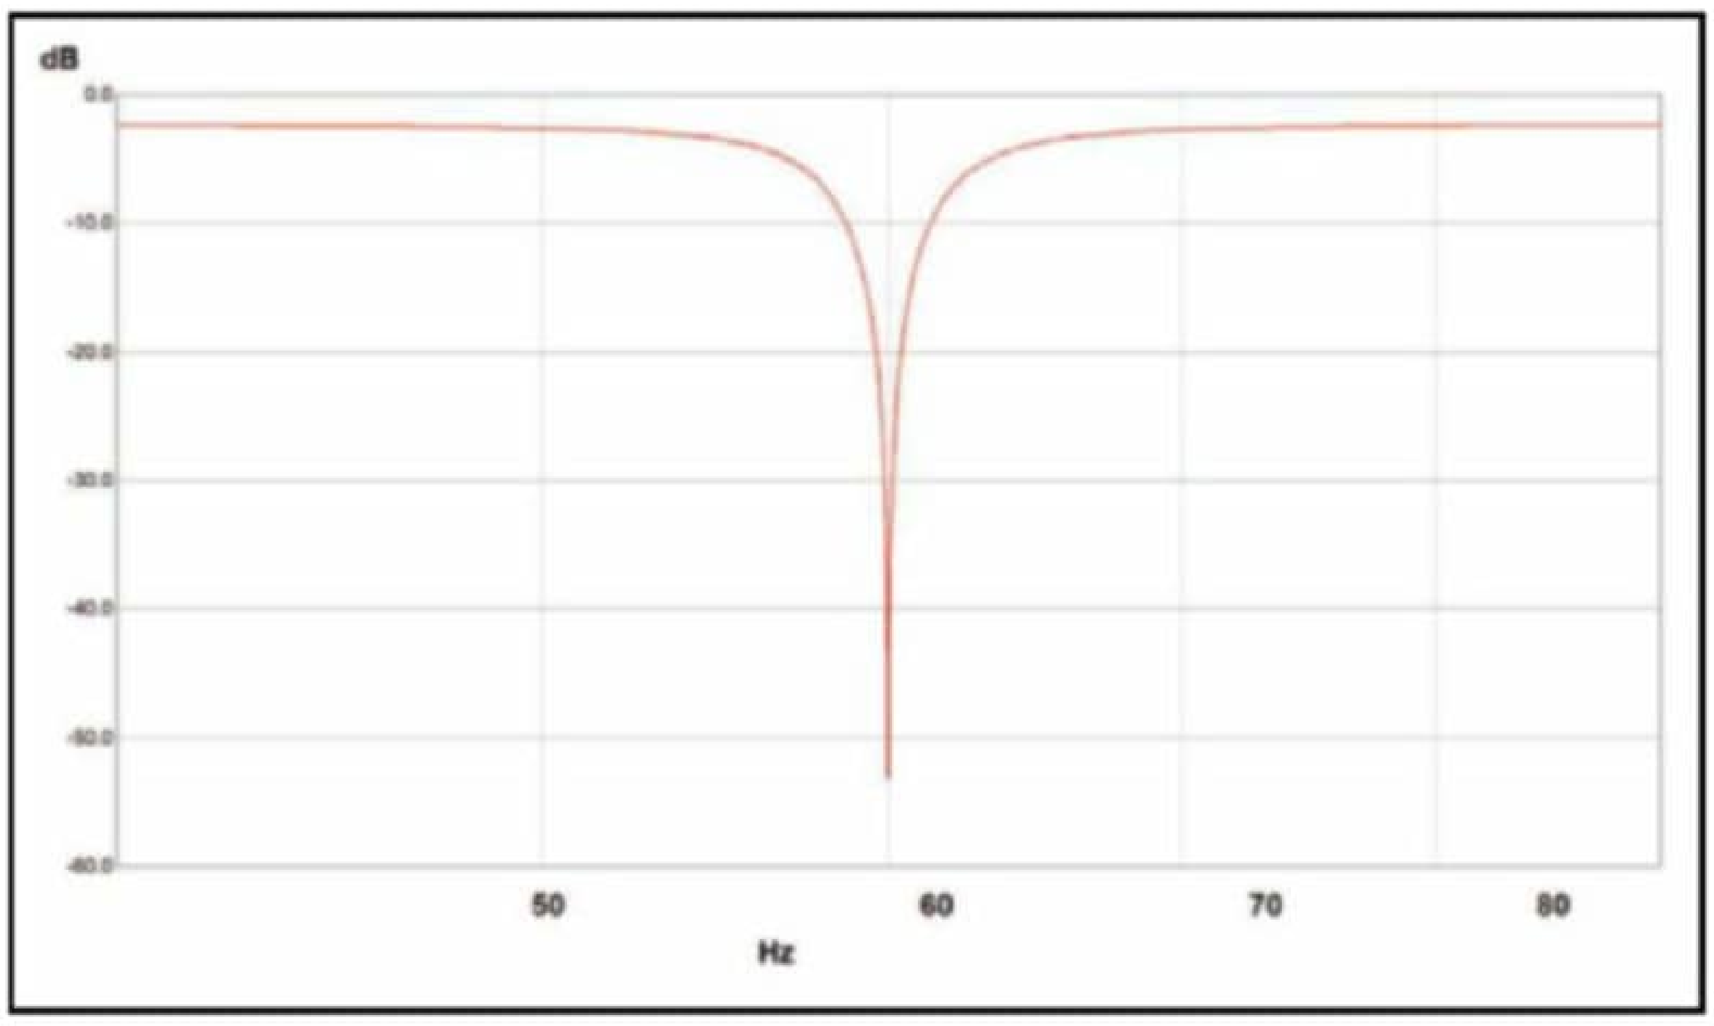

Figure 7: This is a typical response of the notch filter at 60Hz.

The trimmer can adjust the center notch frequency for about ±2%. Although I could have chosen a trimmer with a larger adjustment range, I decided to include the resistor R66 and adjust it accordingly by choosing a lower or larger value. When the total value of resistors R65 in parallel with R66 plus the value of "Notch_Adj" trimmer is increased, the notch frequency decreases and vice versa. For example, if the capacitors C44, C45 have a value of 497nF, then the resistor R66 should be decreased to 36kΩ for a notch frequency of 60Hz.

The procedure to adjust the notch frequency is as follows:

- Open the jumper "Adj" to remove the input bias that is used for the ACH accelerometer.
- Set the Gain at 30dB, the high-pass filter at 20Hz, the LP filter at 100Hz, the output switch at "ACC" and the "Notch" switch to the "Bypass" position (which is the one toward the bottom of the front panel).
- Connect a generator producing a sine wave of 50Hz or 60Hz at the pin 2 (+) and pin 3 (-) of the "Input" connector. It is very important to have a very accurate and stable frequency, exactly at 50.00Hz or 60.00Hz. Otherwise the notch, which is very sharp, will not be adjusted correctly. I found out it is much easier to use a PC program as the generator of the input signal. I recommend the program "Soundcard Scope," which can be downloaded from: https://zeitnitz.eu/scope_en. The program has also a very convenient oscilloscope, which is used to monitor the output signal of the amplifier.
- Adjust the level of the generator signal so that the output level of the AP-2 amplifier is about 4V peak.
- Now set the "Notch" switch to the "On" position.
- Monitor the output to the oscilloscope and adjust the trimmer "Notch_adj" until you measure the minimum output voltage.
- The notch should provide an attenuation at the center frequency of typically 35dB to 50dB. The COG capacitors for C44 and C45 will give a larger attenuation at the notch than the film capacitors. This depends on the accuracy and the matching of the resistors and the capacitors of the filter as I described earlier.

A typical response of the notch filter at 60Hz is shown in **Figure 7**, where the notch attenuation is almost 50dB. The benefit of using the Notch filter is that the level of the 50Hz interference drops to about the same as its harmonics.

Calibration of the Amplifier

The amplifier can operate without calibration, if there is no need for accurate calibrated outputs. Otherwise, it can be calibrated for the specified sensitivity of an ACH-01 accelerometer as measured by its manufacturer and noted on the package.

The calculations that are required were given in my previous article (see Resources) and I will not repeat them here. The calibration is done by connected at the input of the amplifier a sinus frequency with a specific voltage level and the corresponding trimmers, that is "Acc" for the acceleration, "Vel" for the velocity and "Exc" for the Excursion, are adjusted until the calculated voltage is measured at the output.

For the calibration, the jumper "Adj" should be opened before the connection of the external generator. The front panel switches should be set as follows: the Gain at "0dB", the HP at "3Hz", the LP at "20k" while the output switch at "Acc", "Vel", or "Exc", accordingly.

Mounting Methods of the Accelerometer

The mounting of the ACH-01 plays a critical role to its overall performance since an improperly mounted accelerometer can give very erroneous results.

The manufacturer recommends using a quick setting viscous methyl cyano-acrylate adhesive or any epoxy adhesive. This should be applied sparingly to the clean surface and should be allowed to set. This is an excellent method for permanent placement but not at all practical when someone wants to test different points.

Resources

"Accelerometer Testing of Loudspeaker Drivers," diyAudio forum, 2015, www.diyaudio.com/community/threads/accelerometer-testing-of-loudspeaker-drivers.267734

B. Carter and R. Mancini, *Op-amps for Everyone*, Newnes; 5th edition, July 28, 2017.

DesignSpark, www.rs-online.com/designspark/pcb-software

Front Panel Designer, www.schaeffer-ag.de

"Notch Filter Calculator," https://booksite.elsevier.com/9780123914958/content/Notch_Filter_Calculator/notch_2.htm

G. Ntanavaras, "Accelerometer Testing of Loudspeaker Drivers," *audioxpress*, June 2011, www.audioxpress.com/article/Test-your-speakers-performance-with-this-do-it-yourself-measurement-system.html

C. W. Russell, "A Bilateral Clipping Indicator," *Audio Amateur*, Issue 3, 1975.

Although the manufacturer recommends avoiding soft adhesives, such as double-sided tape since they can adversely affect the ACH performance, this is the only method I have found for measurements on loudspeaker walls or speaker cones. My preference is a standard thin double-sided tape from Tesa, which is strong enough to hold the ACH in good contact with the area where it is mounted but not too strong so it can be removed carefully after the measurement. This is something that requires some experience. For example, some loudspeaker cones are made from paper or paper-like material and when the ACH is removed, some part of the cone material could also be removed.

Another very critical point is that the part of the cable close to the accelerometer should not be allowed to move at all and should be adhered firmly to the surface that it is tested with tape.

More Measurements

I performed a lot of measurements to verify the operation of the amplifier and check its accuracy, but I will only present some of them.

Figure 8 shows the frequency response of the high-pass filters of the amplifier while **Figure 9** shows the response of the low-pass filters. **Figure 10** shows a few interesting band pass responses, which can be realized by using a combination of low-pass and high-pass frequencies. For example, the first curve is formed when the cut-off frequencies of both the low-pass and the high-pass filters are set at 200Hz, the second when they are set at 450Hz, and the third one at 1kHz. Having a band pass response helps to measure at specific frequency bands and identify easier a vibration problem.

I measured the distortion of the amplifier at all gains with the output level at ±4V peak at a load of 2kΩ. The distortion was better than 0.02% below 1kHz and 0.1% at 8kHz. When the load decreased to 600Ω, the distortion increased to 0.035% at 1kHz.

Conclusion

The updated AP-2 amplifier with an analog signal processor along with the ACH-01 Accelerometer is a very useful measurement system that offers more functionality than its previous version. It is easy to use, and the cost-to-build is also very reasonable.

If you are interested in vibration measurements and you don't have such an equipment in your lab, here is a good opportunity to build one.

Author's Note: I have PCBs and Front Panels available for the construction of the accelerometer amplifier. If someone is interested send an email to: gntanavaras@gmail.com

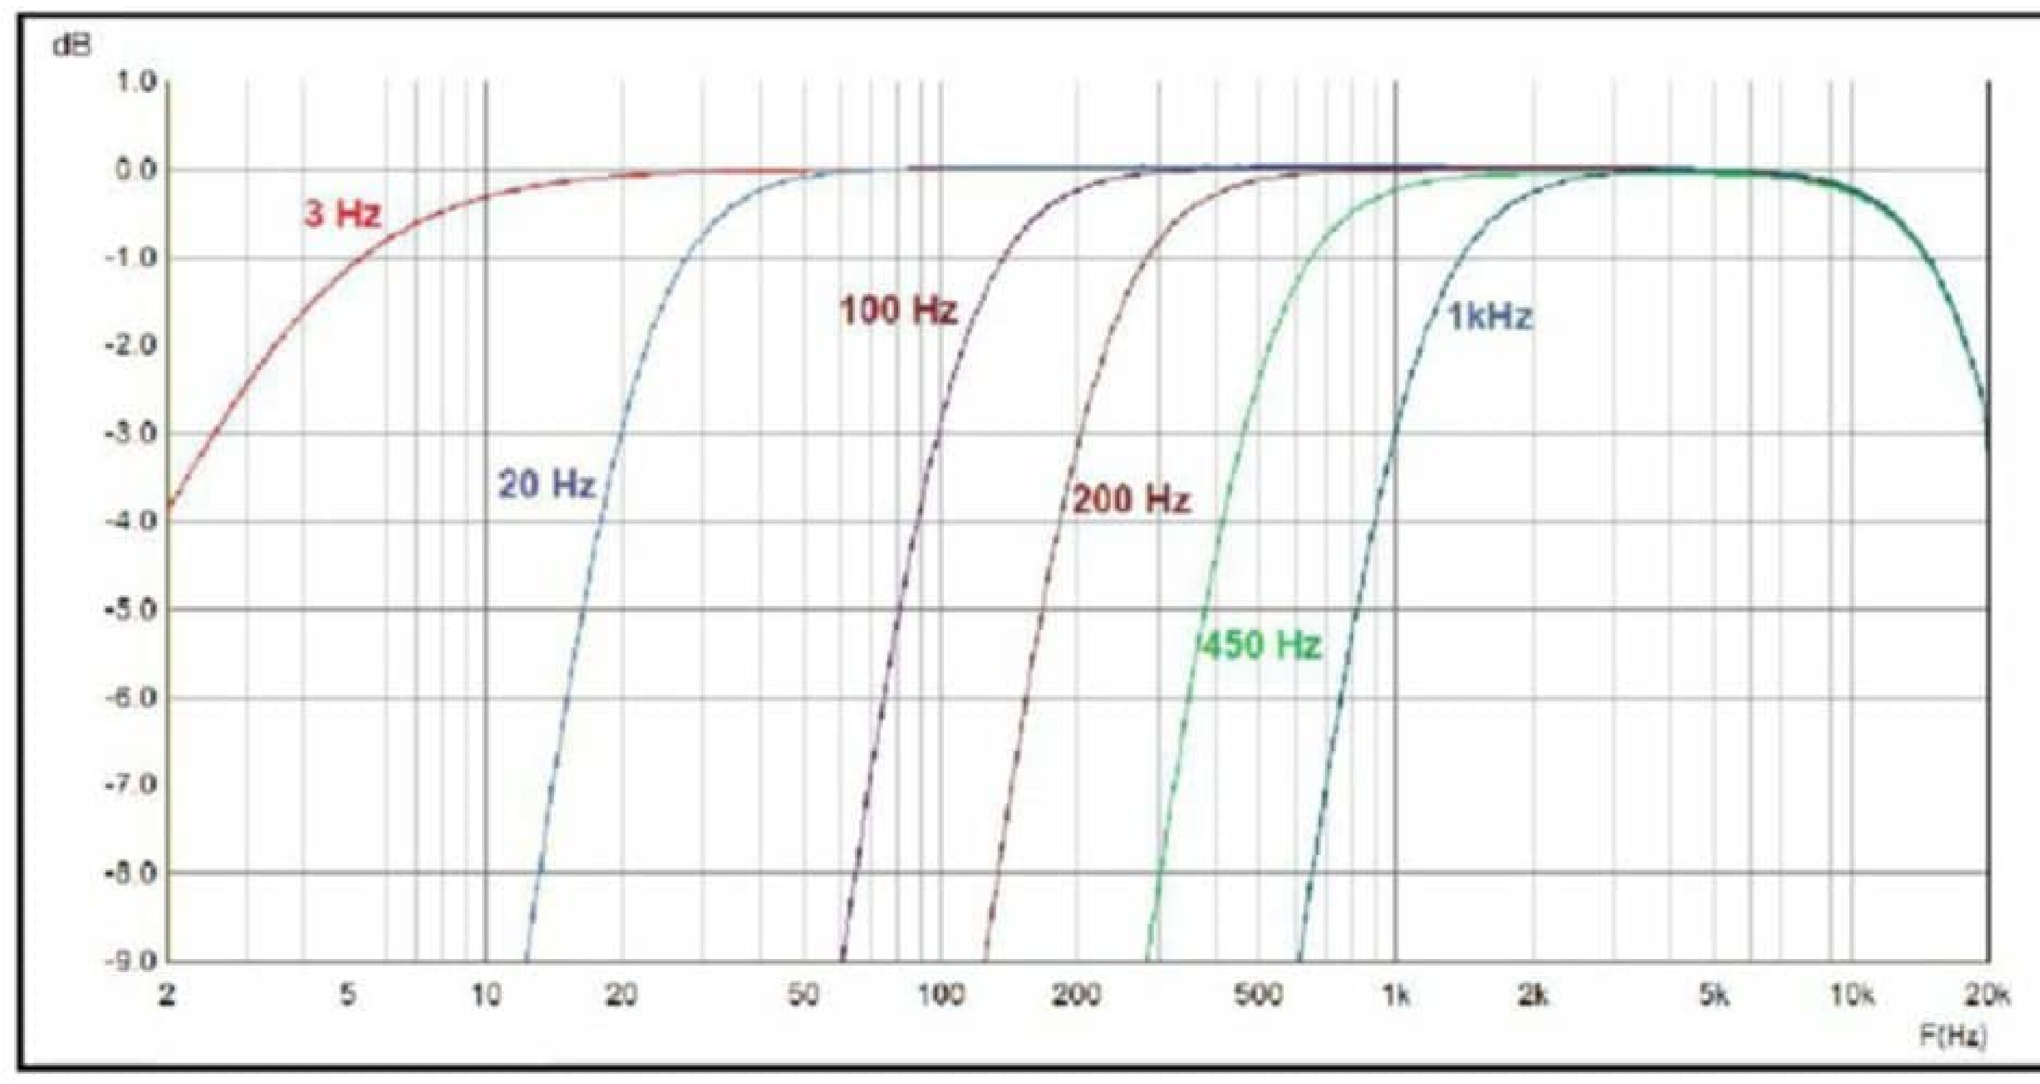

Figure 8: The graph shows the frequency response of the high-pass filters of the amplifier.

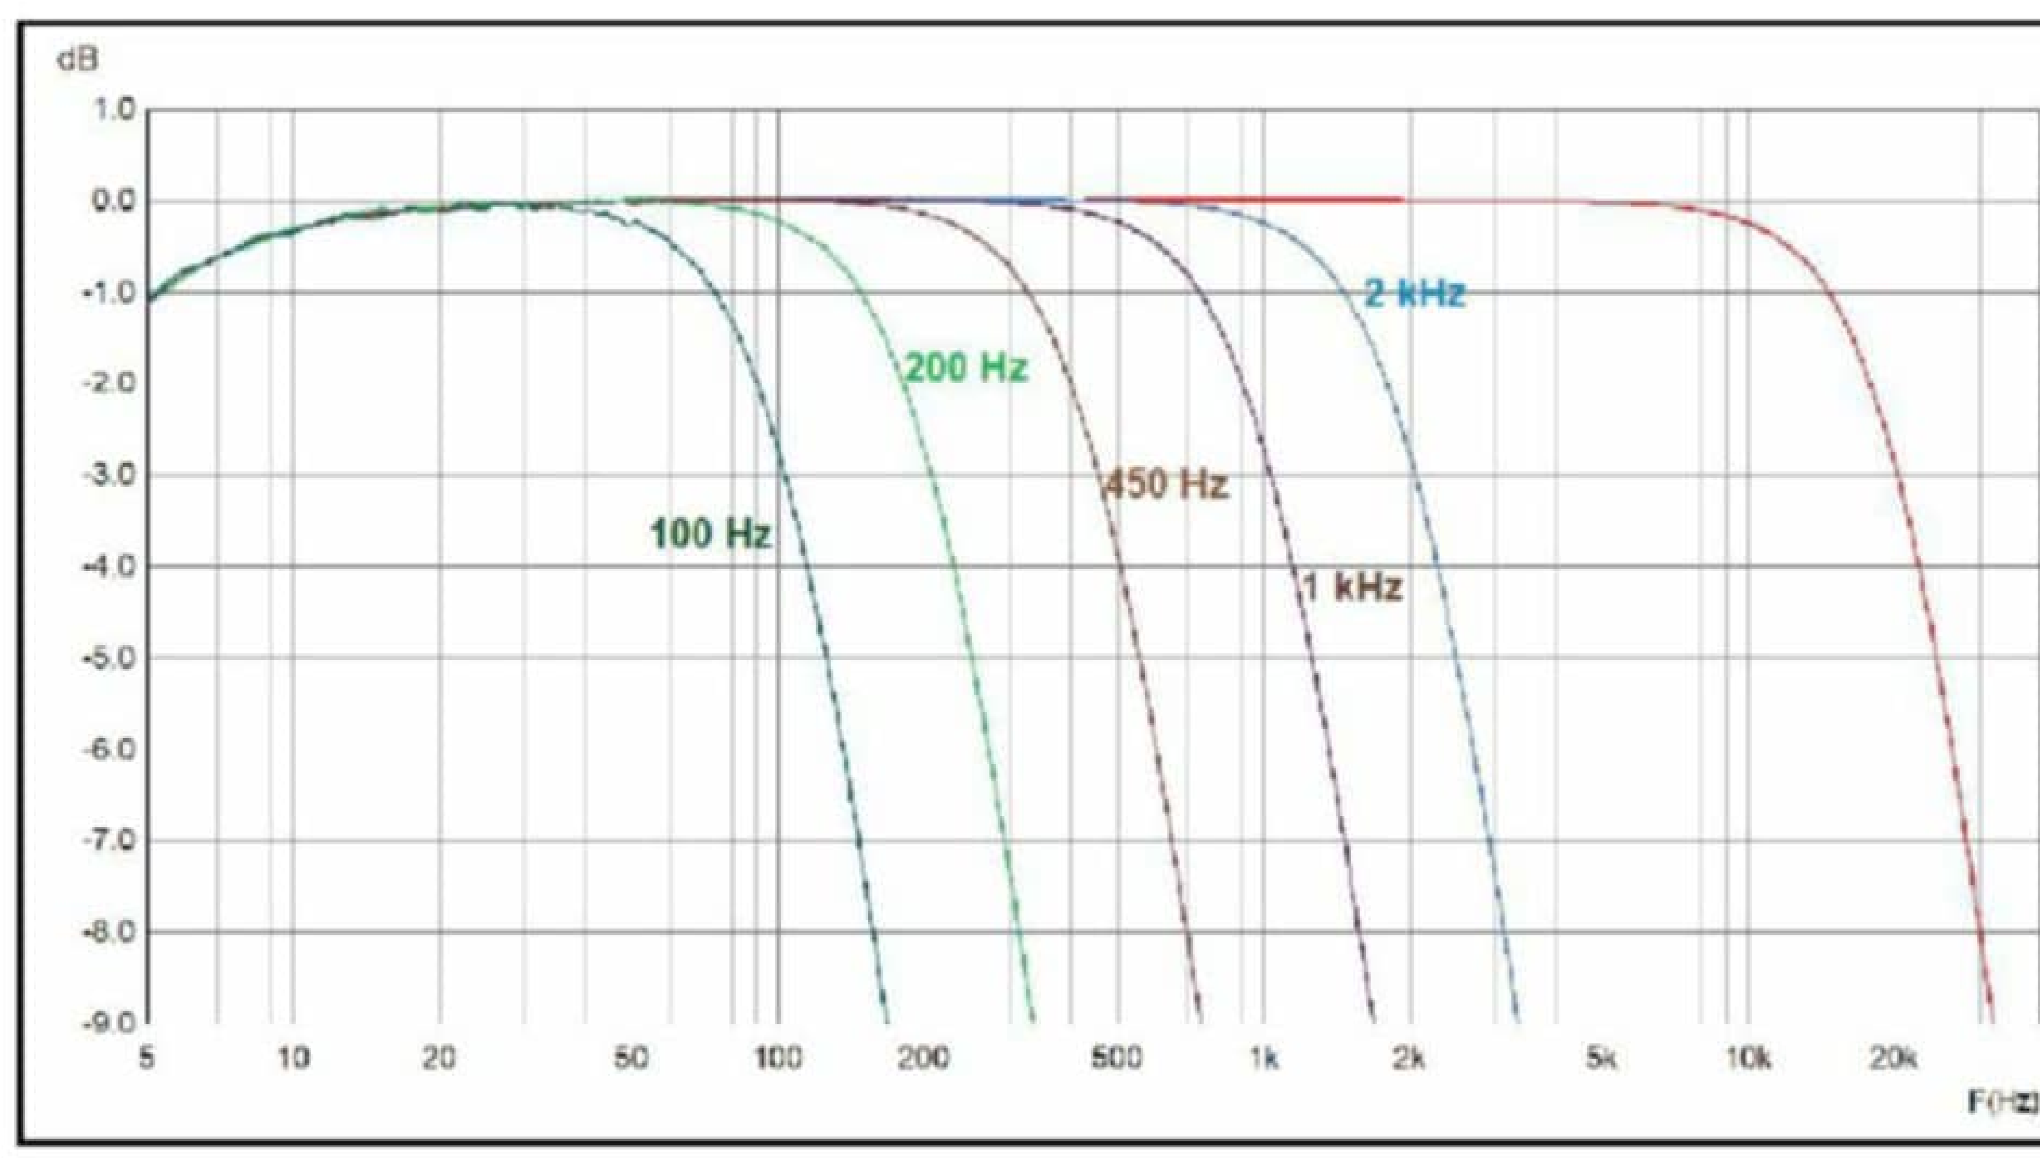

Figure 9: This graph shows the frequency response of the low-pass filters.

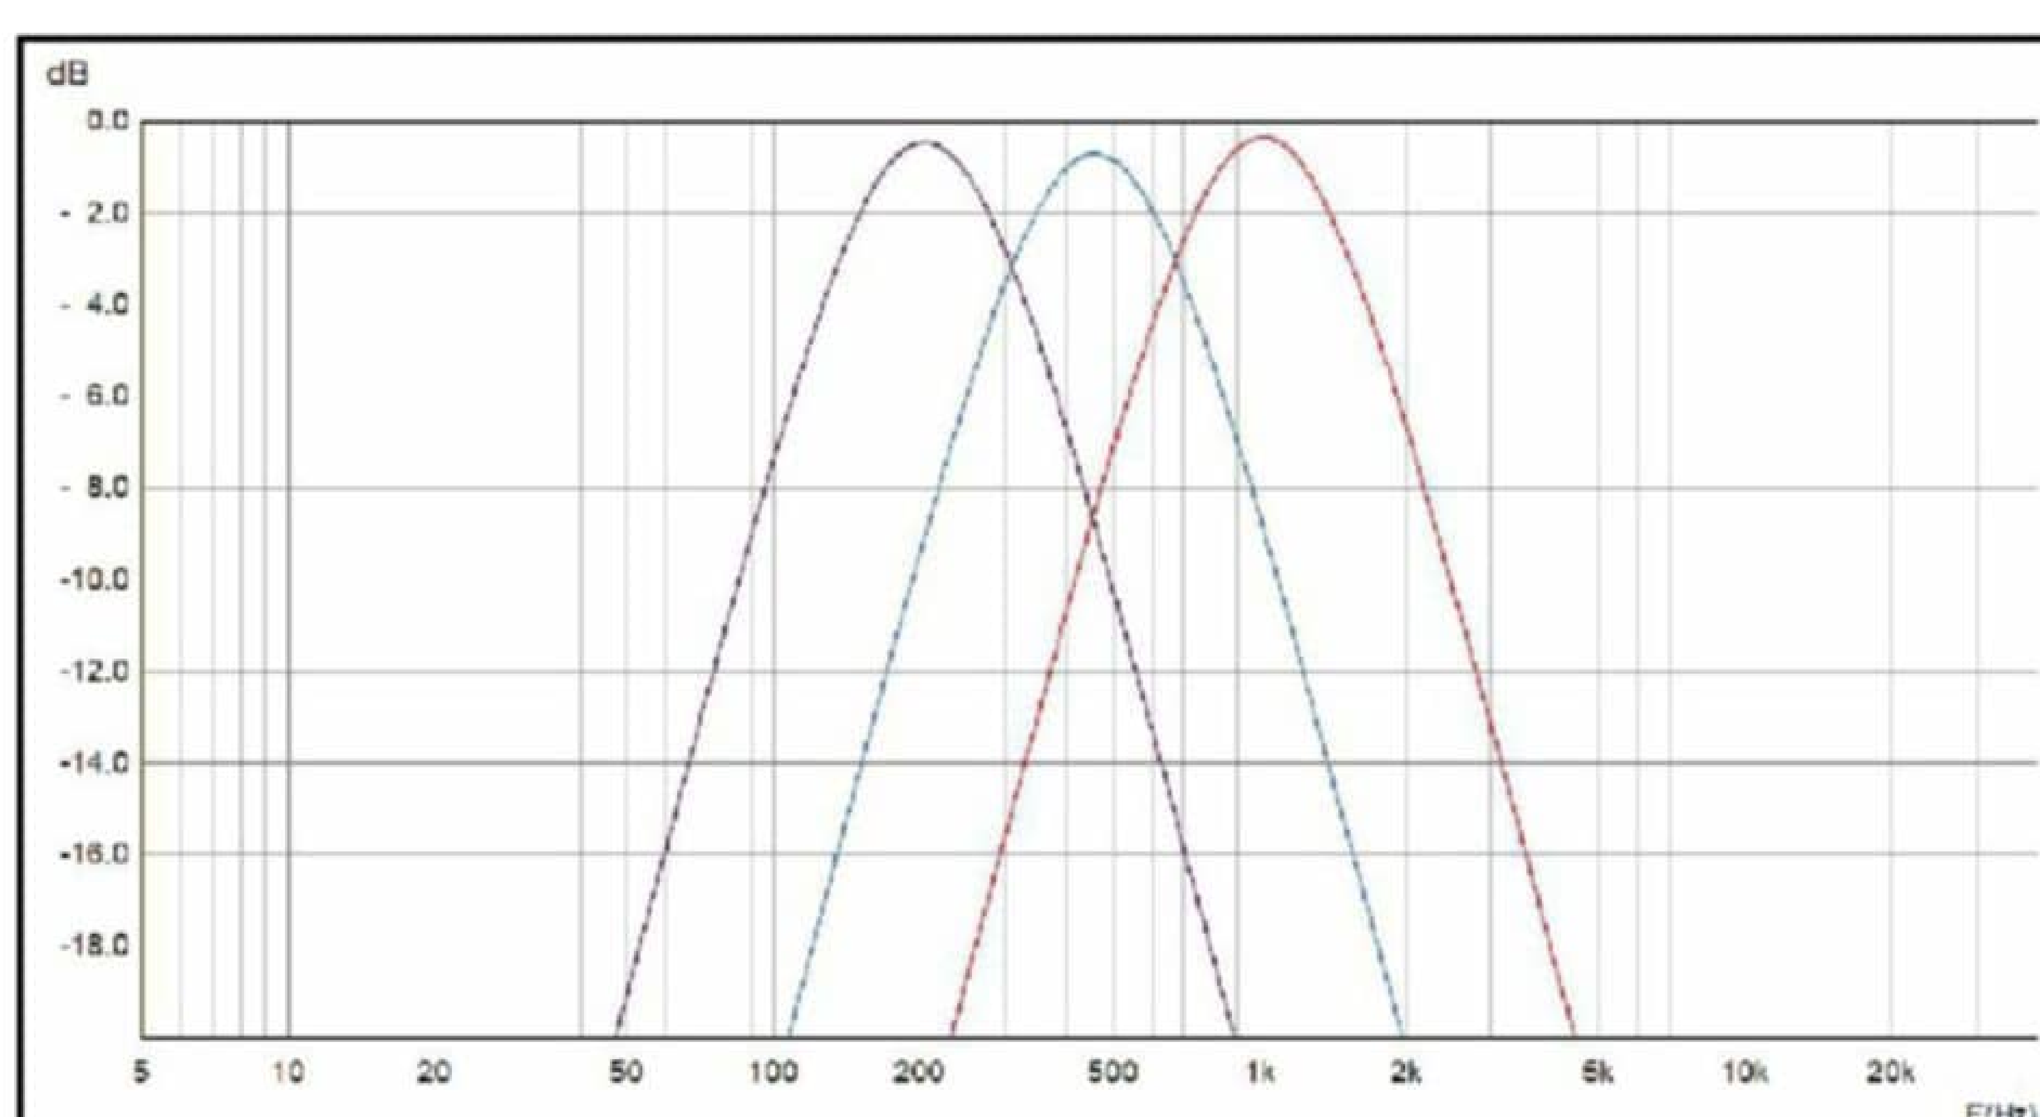

Figure 10: I also obtained a few interesting band pass responses, using a combination of low-pass and high-pass frequencies.

About the Author

George Ntanavaras graduated from the National Technical University in Athens, Greece, in 1986 with a degree in Electronic Engineering. He currently works in the Development Department for a Greek electronics company. He is interested in the design of preamplifiers, active crossovers, power amplifiers, and most loudspeakers. He also enjoys listening to classical music.

Misconceptions in the Audio Industry

Because there is no simulation without a sound theory foundation, this article explores some apparently simple examples of misconceptions, such as clarifying what actually happens when a loudspeaker moves.

By
René Christensen

In the audio and hi-fi communities, some people have certain beliefs about how loudspeakers work that are demonstrably false. This article targets three of them.

"The group delay/phase delay is the 'actual' delay in a system."

Group delay and phase delay are quantities often discussed on forums in the context of loudspeaker and subwoofer behavior, and it is evident that neither are well understood. A variation of the misconception is that "phase and time are the same thing." Let us first look at a system and its transfer function, and then work our way toward the delays. A transfer function describes system behavior as a complex function $H(s)$ as function of the complex frequency s [1]. It is the ratio of the output phasor $Y(s)$ and the input phasor $X(s)$, which are the Laplace transforms of the time signal output $y(t)$ and the time signal input $x(t)$, respectively. For physical systems described via linear differential equations, the transfer function is typically formulated as a rational function, meaning a function that is a fraction with a polynomial in the numerator and a polynomial in the denominator.

$$H(s)=\frac{Y(s)}{X(s)}=K\frac{s^m+b_{m-1}s^{m-1}+\ldots+b_1s+b_0}{s^n+a_{n-1}s^{n-1}+\ldots+a_1s+a_0}$$

The transfer function directly tells us a particular behavior of the system, namely its steady-state behavior. There is an associated frequency response, which consists of a magnitude response and a phase response. These responses focus on the subset of frequencies that fall on our "standard" frequency line, where the sinusoids reside. While the general transfer function also directly gives us the response for exponentially damped or increasing sinusoids, the sinusoidal response will typically suffice to describe the system, as any relevant music can be decomposed into such sinusoids. The magnitude response is not important for the delays as they are calculated from the phase response, which can be expressed as:

$$\phi(\omega)=\angle H(s)=\arg H(s)=\angle K+\angle\left(\frac{i\omega}{z_1}+1\right)+\angle\left(\frac{i\omega}{z_2}+1\right)+\ldots$$
$$+\angle\left(\frac{i\omega}{z_3}+1\right)-\angle\left(\frac{i\omega}{p_1}+1\right)-\angle\left(\frac{i\omega}{p_2}+1\right)-\ldots-\angle\left(\frac{i\omega}{p_3}+1\right)$$

The zs indicate the zeros of the transfer function, meaning the solutions to the denominator polynomials, and the ps are the poles; the solutions to the numerator polynomial. When the phase response has been established, two relevant delays can be calculated, the phase delay t_p and the group delay t_g:

$$t_p=-\frac{\phi(\omega)}{\omega}$$
$$t_g=-\frac{d\phi(\omega)}{d\omega}$$

Before going into too many details, it can be informative to illustrate how to think about these delays. The phase delay can be found by drawing a straight line from any frequency in question to the (0,0) point and finding the negative slope of said line to arrive at a time value. Similarly, the group delay is found via a negative slope, but here it is the local slope of the phase calculated via the derivative. The major point to note about these delays are that they are defined via the phase, and since phase is defined for steady-state conditions only, the delays are also limited to steady-state behavior. So while, for example, the group delay can be negative in part of the frequency range, that does not indicate any issues with causality in general.

For the special case of a linear phase, the phase delay and the group delay have the same value across all frequencies, and this case is a "pure" time delay, which we will discuss a little later. In general, however, the phase delay will indicate the apparent delay in the system when viewed at steady-state, and the group delay will indicate how the envelope of a group of frequencies around a certain frequency will appear delayed. But there might not be any latency in the system at a frequency, where both phase delay and group delay are non-zero, as we can illustrate with a second-order allpass filter transfer function with Q=1 and ω_0=1000 rad/s (i.e., a frequency of 159Hz):

$$H(s)=\frac{s^2-\frac{\omega_0}{Q}s+\omega_0^2}{s^2+\frac{\omega_0}{Q}s+\omega_0^2}=\frac{s^2-1000s+1000^2}{s^2+1000s+1000^2}$$

The phase of this transfer function is plotted alongside the phase and group delays in **Figure 1**, and at the characteristic frequency ω_0 these delays are clearly non-zero.

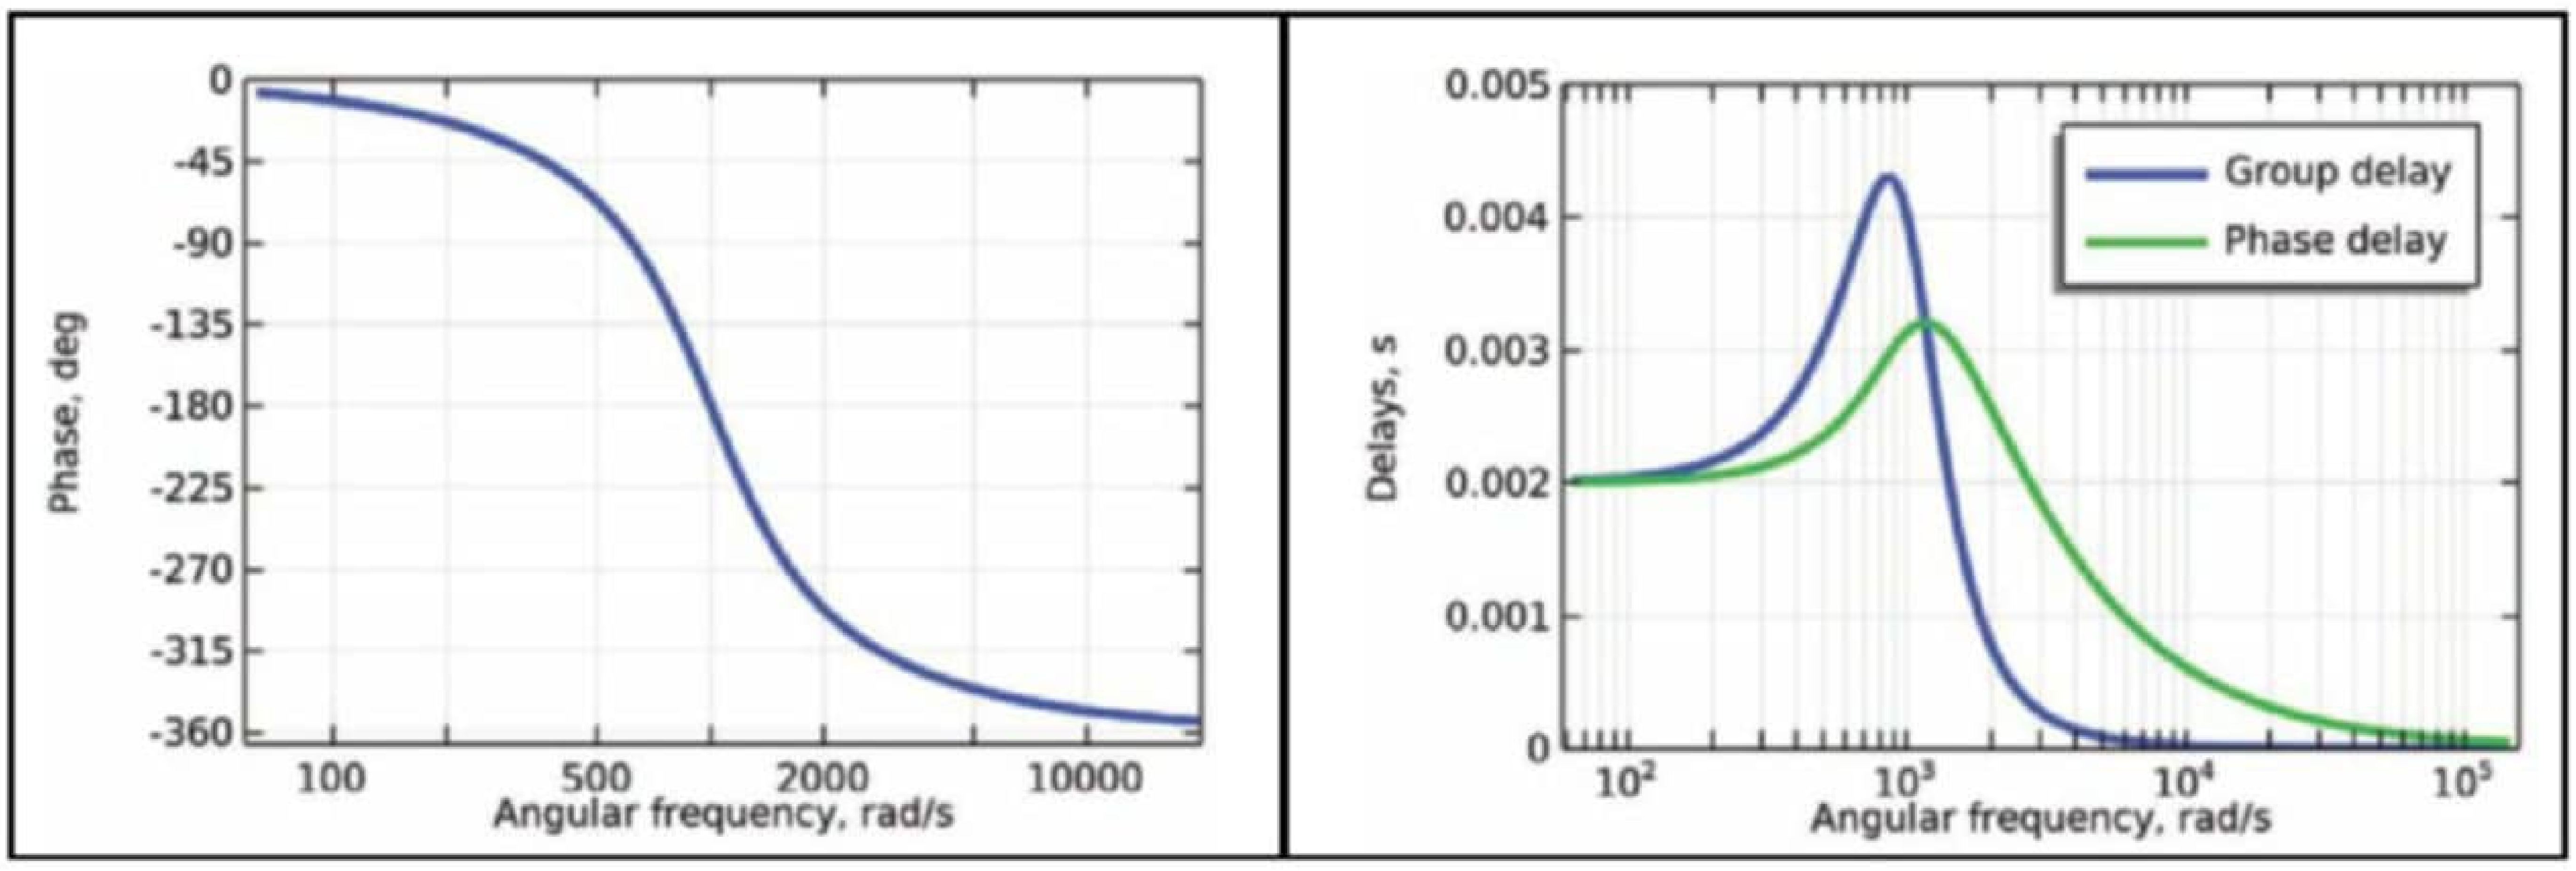

Figure 1: The phase of the allpass filter, and the associated delays

But if we calculate the output including the transient response at this frequency for a sinusoidal input as illustrated in **Figure 2**, we see clearly that there is an immediate output, so there is zero inherent latency in the system at this frequency. The phase delay does indeed match up with the apparent delay seen between the input and the output after steady state is reached, but the group delay only gives an indication as to when the transient response settles into steady state, and its actual value does not indicate any "true" delay in the output. Phase delay and group delay need to be considered as a pair, and as we will see later, their differential is important. They are always be derived from the steady-state phase, but this phase cannot in general be derived from any one of them, and neither one of them generally reveal the transient delay directly without some explicit calculations. Hence, it is also incorrect to say that phase and time are the same thing.

"A loudspeaker driver works by creating a high pressure as it moves outward, and a low pressure as it moves inward."

This is one of the most significant misbeliefs that I see repeated often, as it is completely opposite to how the sound from a loudspeaker is generally generated. Let us first set the stage: We are looking at steady-state conditions where the system has had time to settle and simple oscillation is achieved, which is how we analyze loudspeakers most of the time. For each frequency we look at how the driver is moving and compare it to the resulting pressure. And we are looking at anechoic conditions, so there is no room to consider. This situation can of course be modeled, and the perfect place to start is with the so-called First Rayleigh Integral [2], where the pressure phasor is calculated as:

$$p(P)=\frac{-\omega^2\rho_0}{2\pi}\int_s w(Q)\frac{e^{-ikR}}{R}\,dS$$

The integral gives us exactly the complex pressure p in point P at a distance R from a point Q on a vibrating flat surface with an outward displacement $w(Q)$ possible varying over the piston surface S under free-field conditions, and we can immediately note a minus sign being present on the righthand side of the equation. This indicates that there is a general sign change between displacement (seen as being positive moving outward from equilibrium) and the resulting pressure (seen as positive under compression). We can simplify the expression by assuming a circular piston, a constant displacement across the piston, and that the measurement point P is in the far-field relative to the radius a of the piston with $R>>a$, giving a pressure of:

$$p(P)=\frac{-\omega^2\rho_0 w\pi a^4}{2R}\left(\frac{2J_1(ka\sin\theta)}{ka\sin\theta}\right)e^{-ikR}$$

Furthermore, if the measurement point is located on-axis with θ=0, the expression in the parentheses equals one, and the pressure simplifies to:

$$p(P)=\frac{-\omega^2\rho_0 w\pi a^4}{2R}e^{-ikR}$$

The exponential term accounts for the phase related to the time it takes for the waves to travel from the piston to the measurement point, and that phase is typically ignored in measurements by subtracting it from the total phase to get to the relevant phase of the pressure related to the system behavior. With an inherent vector direction being positive outwards from the piston to the exterior domain in front of it, the pressure phasor is in anti-phase with the displacement phasor, whereas the acceleration phasor of the piston is in-phase with the displacement phasor. As the piston reaches its outermost position achieving its maximum positive displacement, the sound pressure is at its lowest negative value, meaning maximum rarefaction is achieved, as opposed to maximum compression. At this position, the piston acceleration will have reached its lowest negative value for its acceleration phasor, and in general the pressure phase and level will correlate well with the acceleration dittos.

Instead of an analytical expression, we can also do a numerical analysis with the simulation software COMSOL Multiphysics. A circular piston is placed in an infinite baffle, and an enclosure is placed behind the piston as shown in **Figure 3**.

We can see from the color legend that at its most inward position, the piston compresses the air in the enclosure on its rear side, as well as the air on its front, so a positive sound pressure is seen on either side. This goes against the common belief that the exterior sound pressure should be at its highest positive value in the most outward position of the piston.

The enclosure pressure illustrates where this wrong intuition might come from, as here we do indeed see the displacement being in-phase with the pressure due the acoustical impedance

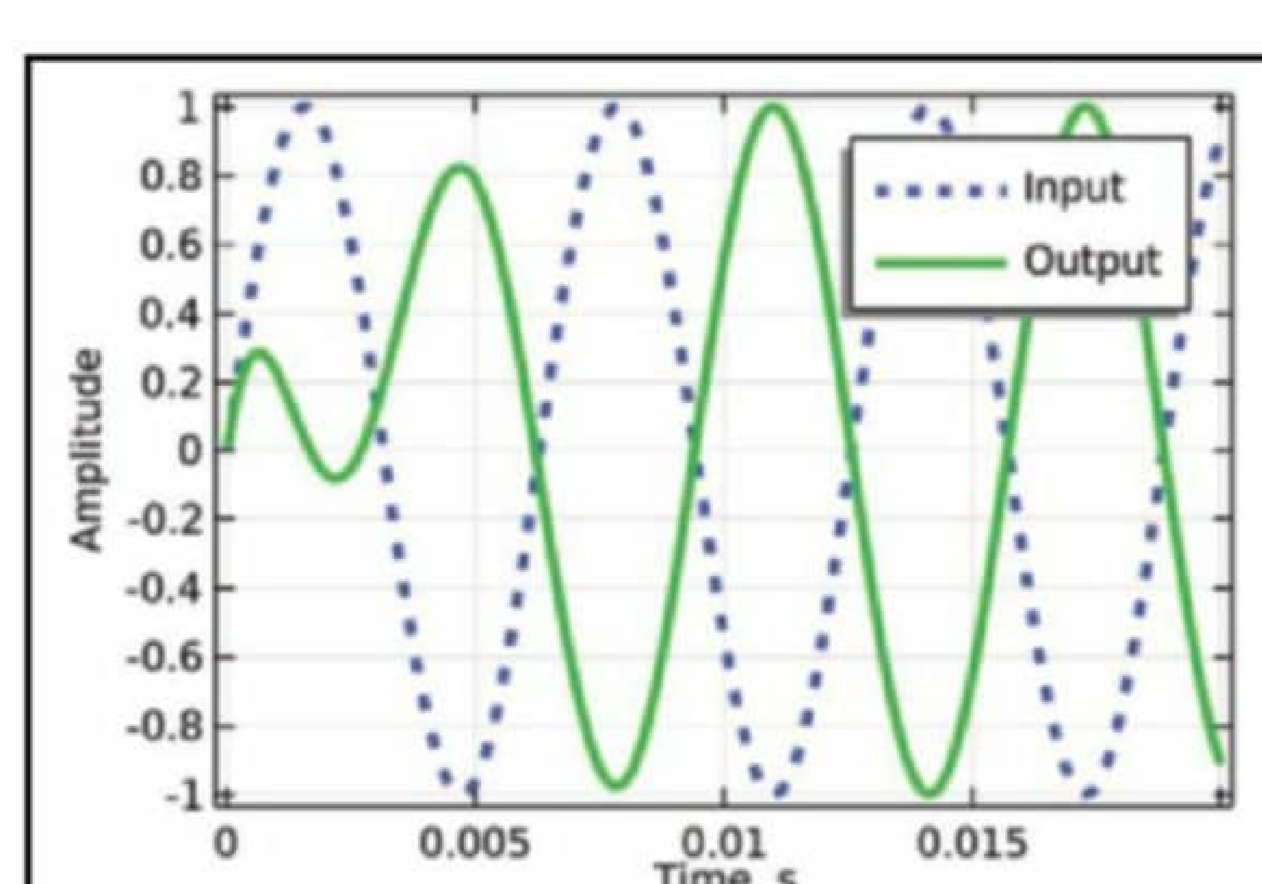

Figure 2: The time response output (green) of the allpass filter to a sinusoidal input (blue)

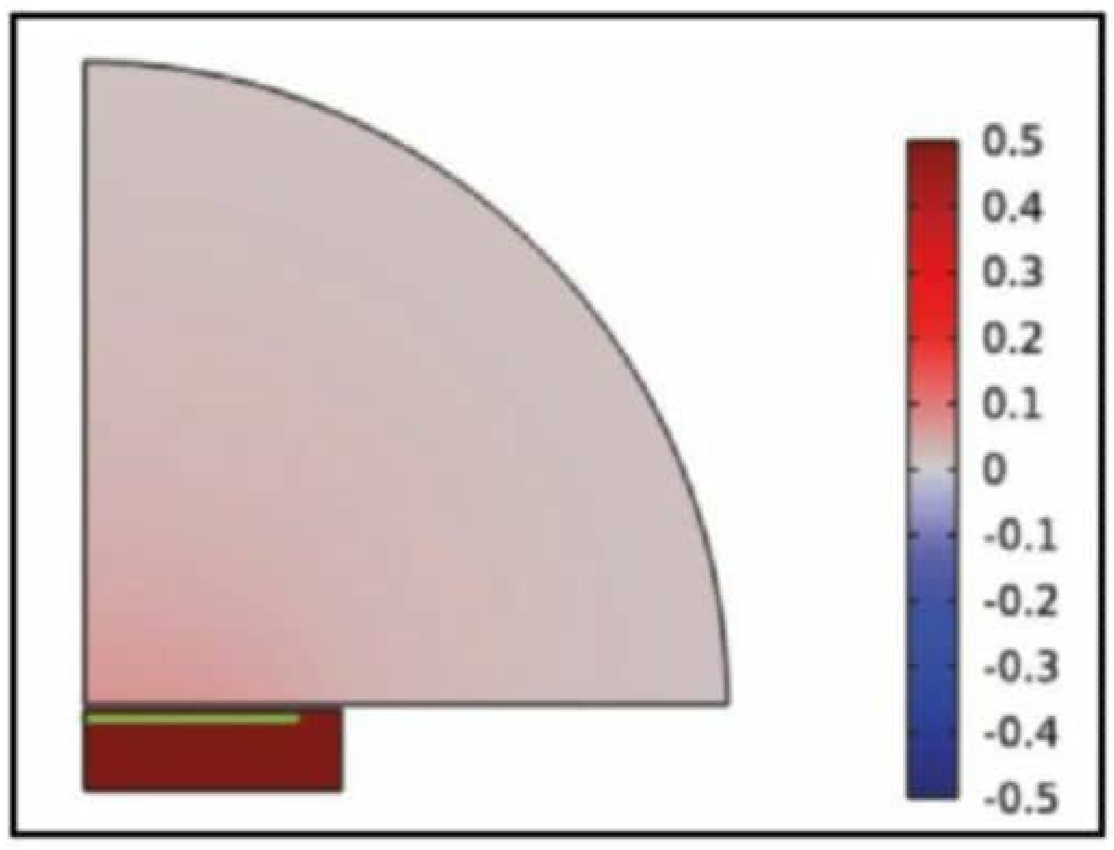

Figure 3: The sound pressure for a piston (green) in an enclosure and infinite baffle radiating into free space

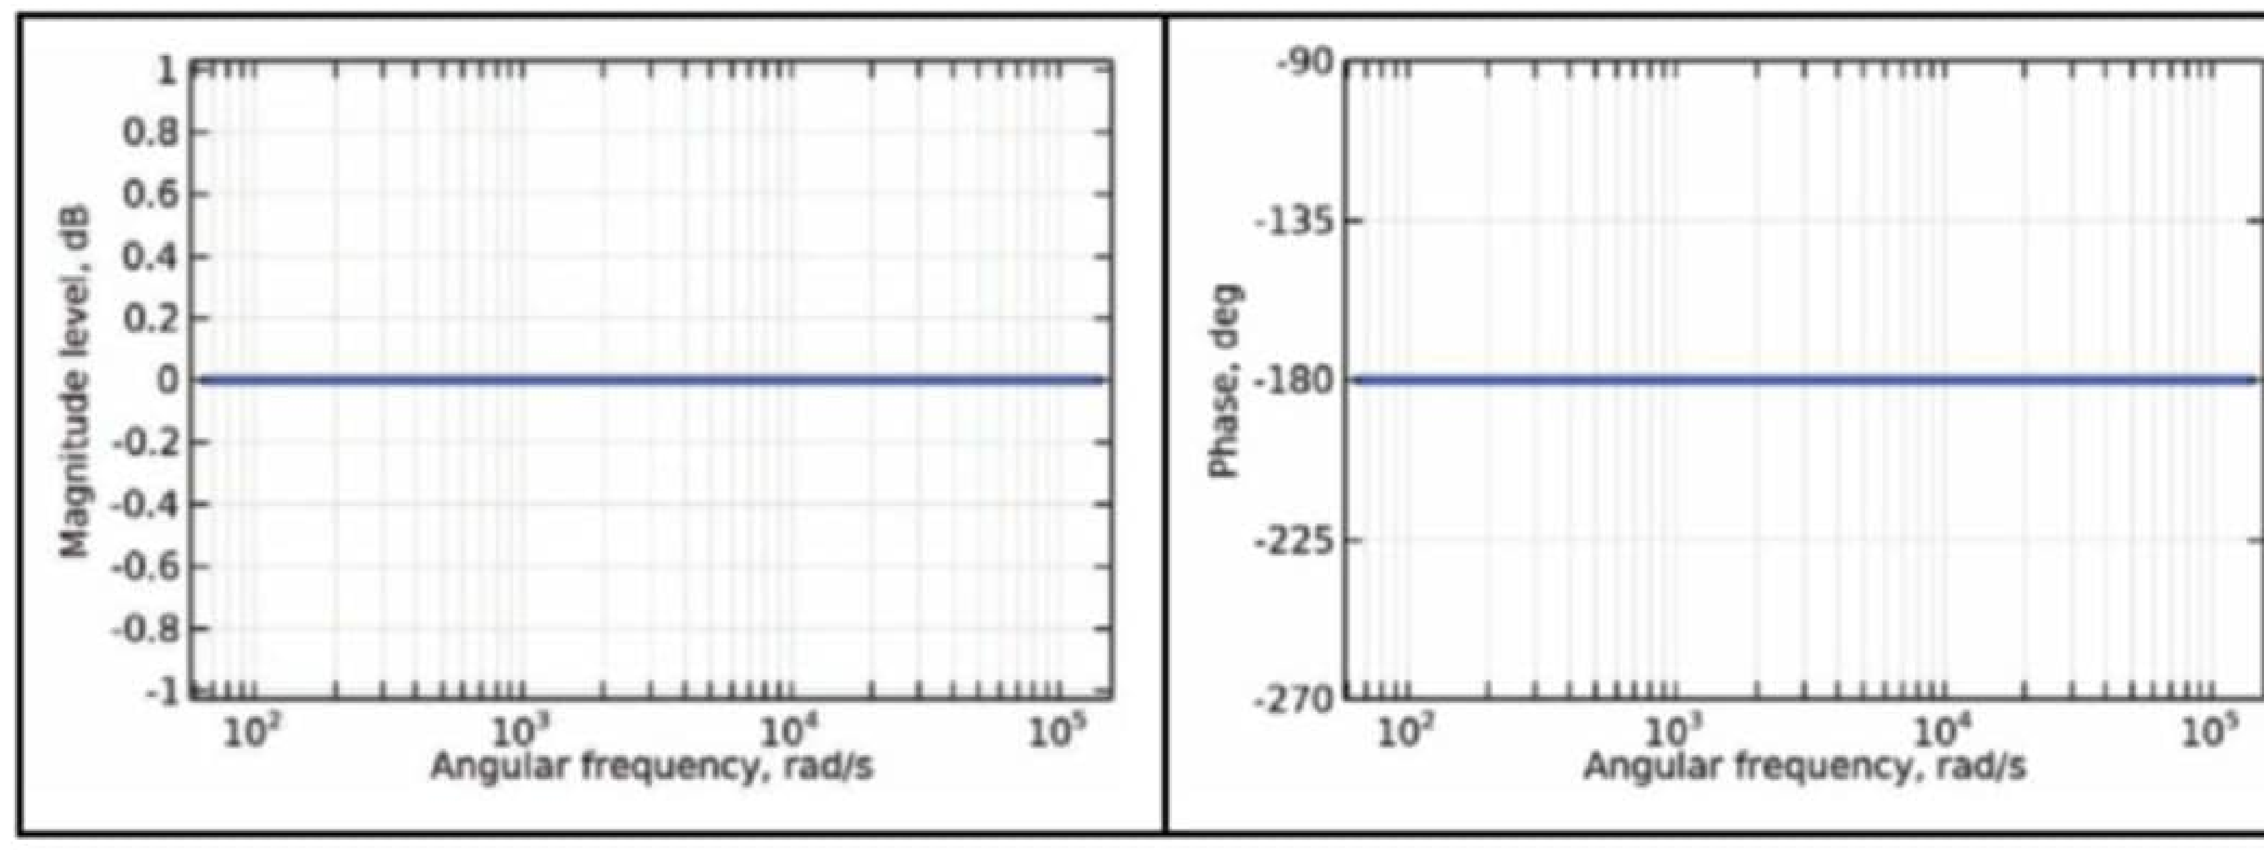

Figure 4: The magnitude level and phase responses for a polarity flip

being a compliance and not free-field conditions. In general, the resulting pressure phase will not only depend on the phase of the piston displacement, but also on the phase of the acoustic impedance that describes the acoustic environment in question.

If a loudspeaker worked by creating compression air as its cone moves outwards from its equilibrium, all our textbooks would have to be rewritten, and our definition of correct polarity would need a flip. Speaking of polarity flip…

"A polarity flip cannot be equivalent to a 180° phase shift, since such a phase shift would entail a time delay."

This is a classic on Hi-Fi forums and has been for years. A polarity flip is what you get when you flip your loudspeaker wires, on both left and right speakers, so that the "red" wire goes to "black" terminal and vice versa. The resulting displacement on the cone will now be opposite of what it would otherwise; an inward displacement now becomes an outwards displacement. We can describe the polarity flip as a system with the following transfer function:

$$H_{pol}(s) = -1$$

A transfer function is generally a complex valued function and here the argument is the complex frequency $s=\sigma+i\omega$. From the transfer function we can derive the system's magnitude and phase responses. The magnitude response is the same as that of a wire, as magnitude is unaffected by a sign change, so we get:

$$\left|H_{pol}(\omega)\right| = 1$$

With the magnitude unaltered, a polarity flip must then change the phase, and indeed it does. Starting out with Euler's identity we can see how a complex exponential with a phase of φ can be written via trigonometric functions as:

$$e^{i\phi} = \cos(\phi) + i\sin(\phi)$$

To arrive at a value of minus one, we set the phase to π radians or 180°:

$$e^{i\pi} = \cos(\pi) + i\sin(\pi) = -1 + i0$$

The 180° might as well be written as -180° as these two phases coincide, and the latter is the choice typically made when looking at transfer functions, ending up with the phasor phase associated with a polarity flip being minus $-\pi$ radians at all frequencies as:

$$\angle H_{pol}(s) = \arg(-1) = \arg\left(e^{-i\pi}\right) = -\pi$$

So, the effect of a polarity flip can be visualized via its frequency response shown in **Figure 4**.

Now, what can be confusing is that other transfer functions can have a 180° phase shift at certain frequencies, and they will have the same steady-state output, assuming of course they share the same magnitude response value, at those frequencies. One such transfer function is that of a pure time delay. All the delay does is… well, delay the output. So any input will "buffered" in some sense before being output, which could be achieved with a digital system or via a transport delay where the signal is taken out of a wave at a distance from some input source. It could be also to some degree be approximated via a high-order Bessel allpass function [3]. The time delay constitutes a so-called "distortionless system" [1], as it has a magnitude response of one and a linear phase leading to no magnitude distortion and no phase distortion. The transfer function for a time delay of t_d seconds is:

$$H_{t_d}(s) = e^{-i\omega t_d}$$

And the magnitude response and phase response are, respectively:

$$\left|H_{t_d}(\omega)\right| = 1$$

And:

$$\angle H_{t_d}(s) = \arg\left(e^{-i\omega t_d}\right) = -\omega t_d$$

The responses are plotted in **Figure 5**.

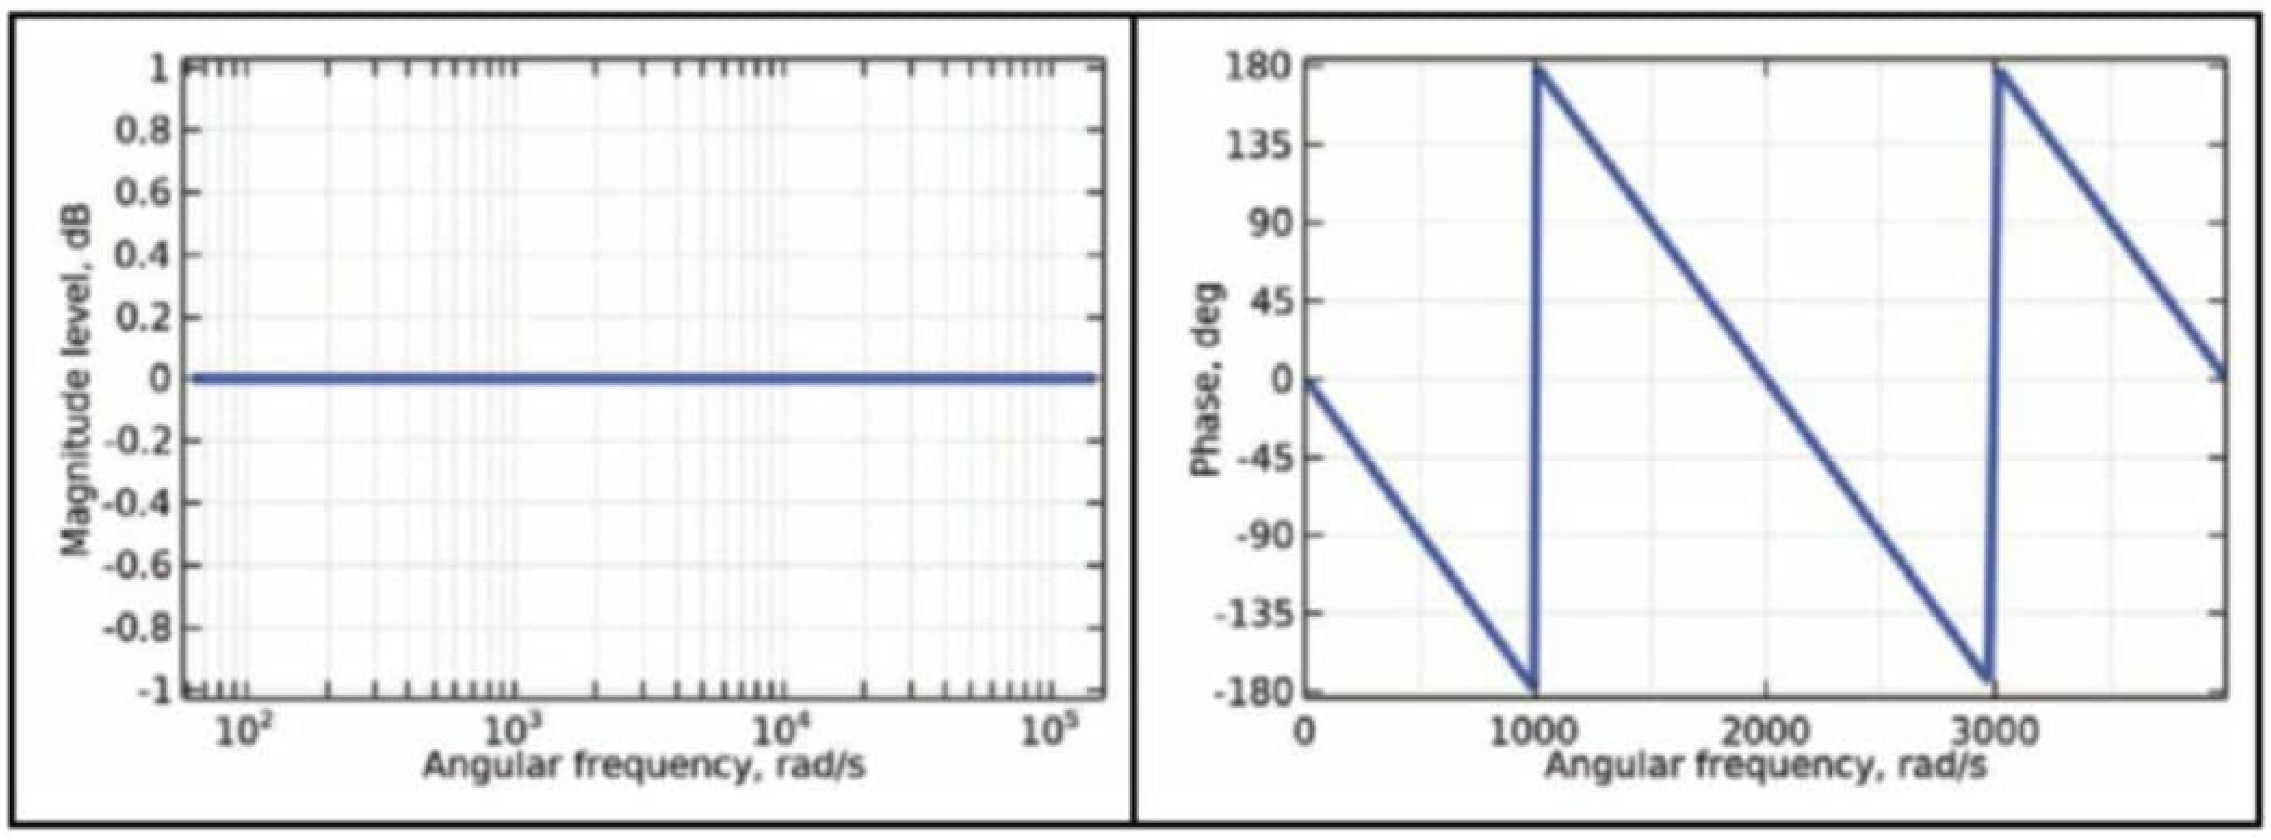

Figure 5: The magnitude level and phase response for the time delay

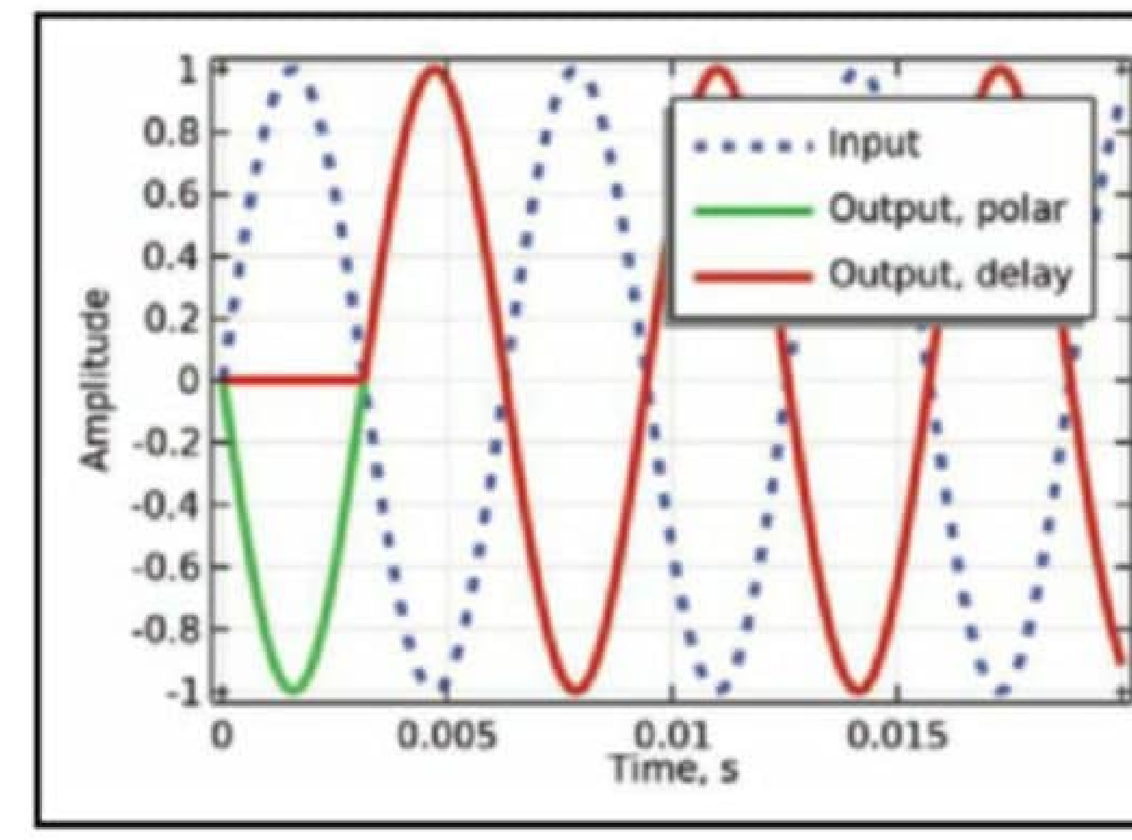

Figure 6: The input sinusoidal and the corresponding outputs from the polarity flip and the time delay systems, respectively

Where a polarity flip has a phase of -180 degrees at all frequencies, we see that the time delay will only have a phase shift of -180° at some discrete (angular) frequencies, namely at:

$$\omega_N = N\frac{\pi}{t_d},\ N = 1,3,5...$$

At these frequencies the time delay system will have a delay that fits with half a period, T/2, so at steady state the phase delay will be the same as for the polarity flip. If you look at the steady-state outputs from the two systems at one of those frequencies, there is no way to distinguish the polarity flip from the time delay. But we could make infinitely many systems that at these frequencies will have the same steady-state output as these two systems. What we need to remember is that there also is a transient response to consider, and here we will immediately see the difference. In **Figure 6**, three signals are shown—an input signal, the output from a polarity flip, and the output from a time delay at a frequency where the two systems have the same phase of 180°. It is clearly seen that the two outputs are not the same in the beginning but then look indistinguishable after half a period.

Not only do the two systems have different phase responses, but even for the few frequencies where they do happen to have the same phase, and hence phase delay, and share the same steady-state response, they can never have the same transient response. A major issue can be seen now, and that is that if you start from a drawing of an acausal/steady-state sinusoidal input and ditto output, you entirely miss which type of system caused the output you are looking at and what the associated transient behavior is. So, you cannot start from the output signal and work your way back to the system—you need to work from the system to the output signal.

One thing to note is that the polarity flip will have an associated phase delay that increases to infinity at low frequencies, as the apparent delay needed time for the steady-state response to look like that of the polarity flip will increase toward low frequencies as the slope drawn down to the phase value will increase. At low frequencies the period is long, and so a substantial time delay would be needed to achieve a 180° phase shift via a time delay of half a period.

But this is just an abstraction, and not a true delay. The polarity flip and other phase ambiguities can make it tricky to even plot the phase delay, compared to the group delay. The group delay of the polarity flip is zero, whereas both the phase delay and the group delay for the time delay system are equal to the time delay itself.

In conclusion, a polarity flipped signal has no inherent time delay, and the output will only look like a time-delayed version of the input when viewed steady-state. But that goes for any other output signal with a phase different than that of input signal, hence our calling the phase delay, and group delay for that matter, "apparent" delays [4]. The allpass filter discussed earlier is seen to also have the same steady-state response at the chosen frequency, and so shares the phase and phase delay, with the other two systems at this frequency, but with a different group delay again and different transient behavior.

Conclusion

The article has gone through some essential topics in signal processing and loudspeaker analysis. In an upcoming article, I will go through phase and delay aspects related to loudspeakers in more detail.

About the Author

René Christensen (BSEE, MSc, PhD) has been working with simulations in the loudspeaker and hearing aid industry for several years, and in 2021 he started his own company, Acculution ApS, with a focus on mathematical modelling of many different products, such as transducers. The overall aim is to use academic research results for engineering better solutions faster, training engineers in this type of virtual prototyping, and working with students across the world to progress the field.

References

[1] A. B. Carlson, *Communication Systems: An Introduction to Signals and Noise in Electrical Communication*, Singapore: McGraw-Hill, 1986.

[2] L. E. Kinsler, A. R. Frey, A. B. Coppens, and J. V. Sanders, *Fundamentals of Acoustics*, Danvers, Massachusetts: Wiley, 2000.

[3] R. Christensen, "#039: Allpass Systems: Phase and Time Delay," Acculution, April 23, 2022, www.acculution.com/single-post/039-allpass-phase-and-time-delay

[4] M. W. J. Leach, "The Differential Time-Delay Distortion and Differential Phase-Shift Distortion as Measures of Phase Linearity," *Journal of the Audio Engineering Society* (JAES), Volume 37, No. 9, pp. 709-715, 1989.

A (Z)OTL with 300Bs

Part 2 — The Final Details

In Part 1 of this article, Thomas Perazella presented the famous David Berning ZOTL amplifier and discussed its unique circuitry. In Part 2, he describes his adventure to dismantle and rebuild the amp, and how he fit it on a new chassis for its next owner.

By
Thomas Perazella

This is the completed 300Bs project, shown with tube cage in place.

Instead of using an enclosed chassis, I wanted to use a more traditional open chassis to highlight the tubes and other components. It would be a somewhat less industrial look. To find an appropriate chassis, I went on the Internet and located five different models and presented the photos to the owner of the amp. He picked one model that had a sandblasted aluminum finish and rounded corners. **Photo 4** shows the chassis. It is interesting to note that the cost of the shipping from China was almost as much as the chassis itself.

Several things are noteworthy about the chassis. It is a modular design with top, bottom, front, rear, side, and corner pieces that are screwed together. The construction is quite substantial with the top being 5mm thick, the bottom 3mm thick and the rest of the pieces 8mm thick. This resulted in a sturdy chassis that also provided increased accessibility when mounting components. Two drawbacks were the difficulty in opening holes in the panels and accessing some of the screw locations when doing the assembly.

Construction

Working with a commercial product where all the small components are on PCBs with connecting harnesses is quite different from working with a prototype where components are added and moved during testing on perforated board stock. That, combined with the complexity of the circuit, led me to minimize any modifications to the boards. In addition, I decided to get new tube sockets and mount them on the new chassis so that for the point-to-point wired components, they could be moved from the old to new socket with the least chance of mistakes. Before any wires or components were removed, they were all marked either with tape numbers or markings directly on the old chassis.

Since some of the components connected to the tube sockets were cut to length, I decided to keep the layout of the tubes the same on the new chassis top plate. My first step was to produce a paper template of the locations of the top-mounted components and tape it to the plate for hole fabrication. This was more complicated than I thought because instead of cutting out a large rectangular hole as was done on the original chassis, reducing the strength of the top plate, I decided to drill four holes that matched the four round terminal pads on the bottom of the transformer. At first, I thought the four holes could be symmetrically located with reference to the bottom of the transformer, but I found that they were offset to one side. There were lots of erasures of pencil lines until the template represented the desired location of the holes. **Photo 5** shows the final form of the template used where you can see all the moves that were made.

With a 5mm thick aluminum top plate, making the various sized holes was a challenge. For smaller holes, I traditionally use hole punches. Component hole sizes ranged from ¾″ for the input signal tube sockets to 1-5/8″ for the 300B sockets. In total, there were five different sized component holes. Because of the thickness of the plate and the cost of hole punches, I decided against using them. Another approach is to use Rotabroches, but again in the larger sizes, the cost was prohibitive for a one-off project. For example, the 1-5/8″ size Rotabroach costs $132.

My final decision was to use traditional hole saws and Relton A9 aluminum cutting fluid. Using hole saws in these large sizes requires a drill press, a wood backing plate and clamps to hold both the plate and backer securely in place. Drilling these types of holes with a hand drill can result not only in poorly shaped holes but can also be quite dangerous if the hole saw binds in the hole during drilling.

Tube Cage

One of the drawbacks of the open chassis configuration with the 6JE6 beam pentode tubes used in the switching power supply is the exposed plate caps. Under certain conditions, voltages on those caps can reach 1,000V. This is definitely a situation

Photo 4: The new chassis had a sandblasted aluminum finish and rounded corners.

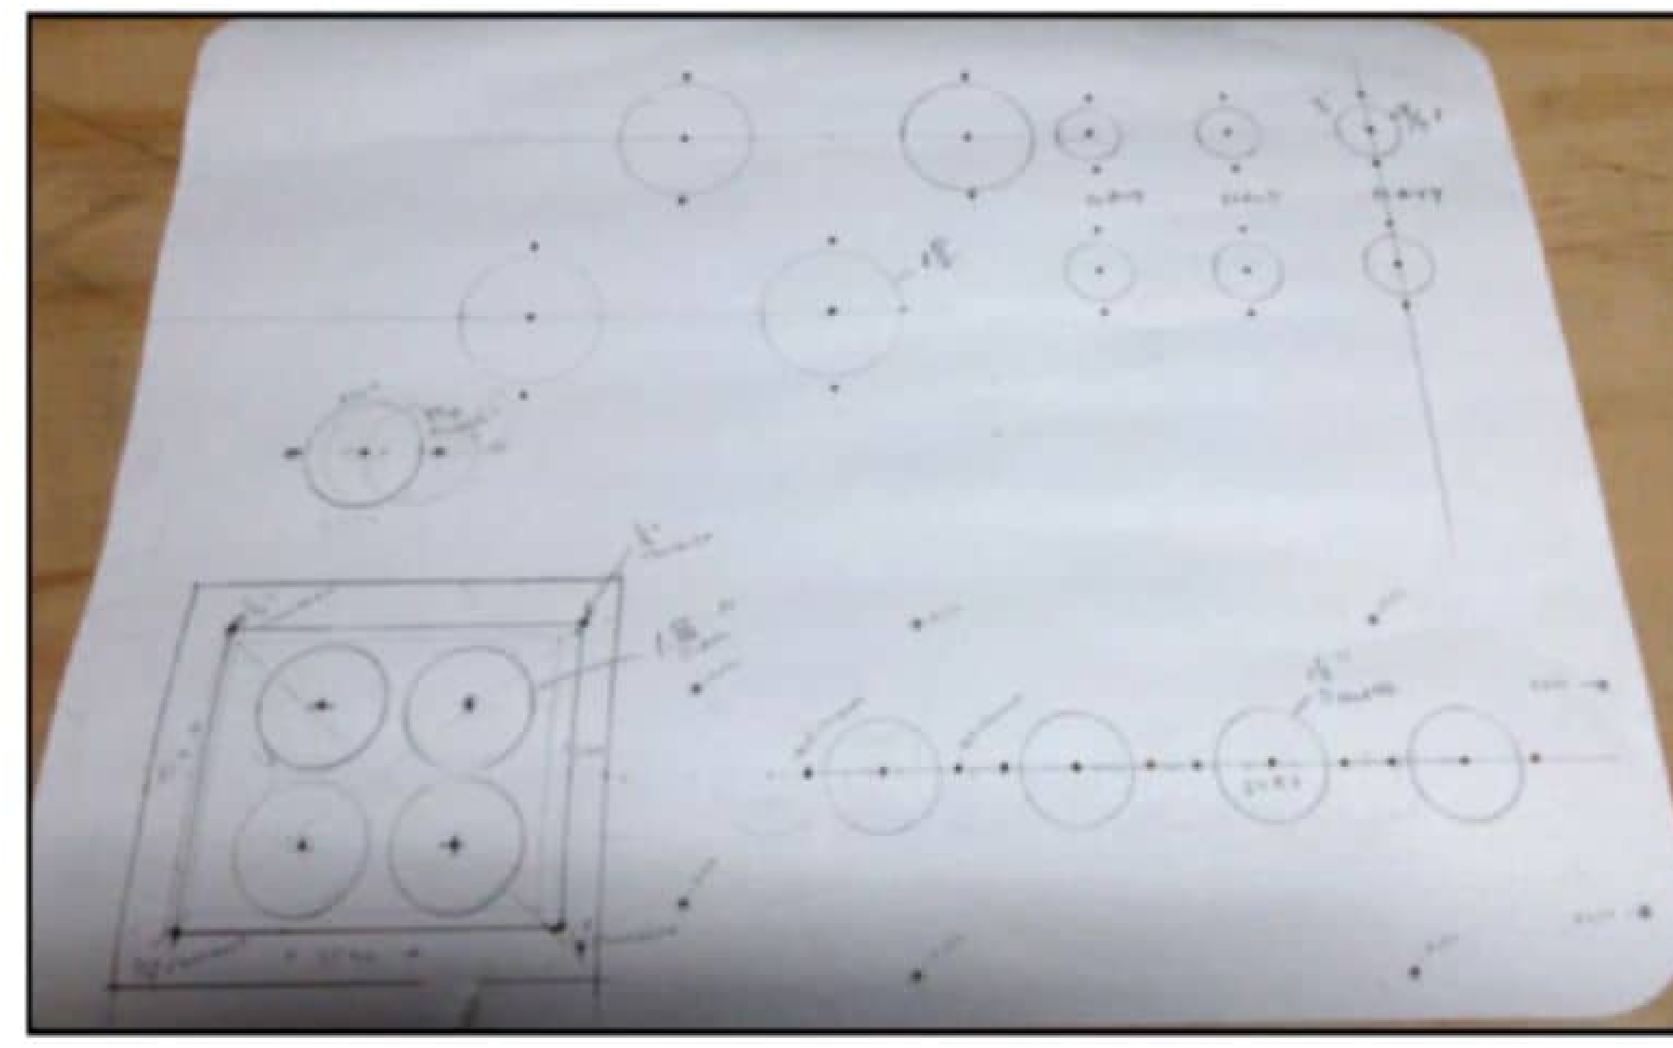

Photo 5: I drew the component layout template using pen and paper to make sure it was correct.

to be isolated from prying fingers. To maintain safety in this configuration I decided to make a tube cage to cover the four tubes. The original chassis had a perforated aluminum cover to enclose all the components. I used that material to make the new cage. After deciding on the dimensions, I laid out a pattern on the material and cut it to shape. Not having a sheet metal brake to bend the shapes, I cut a piece of scrap ¾" MDF into a block the same size as the top of the cage. Then I folded the aluminum up around the MDF form (**Photo 6**).

To form the base of the cage and provide reinforcement at the corners of the cage, I used strips of right-angle aluminum. Since I don't have a welder, I used my old standby, JB Weld epoxy. To make sure the base pieces remained aligned, I screwed them down to a scrap piece of wood with a layer of wax paper over the wood to prevent the frame from sticking to the wood (**Photo 7**).

After gluing the internal corner reinforcements, I glued the cage to the base. After all the epoxy was cured, I removed all the excess before painting. **Photo 8** shows the complete cage where you can see the smoothed epoxy at the junctions of the aluminum.

Paint

There were two major components that needed paint, the tube cage because it was raw aluminum and the power transformer housing because it was in bad cosmetic shape including some rust. After researching several different paints, I chose Krylon Fusion black textured paint. It has a fine textured surface to hide blemishes but does not collect dust like the old wrinkle finishes. **Photo 9** shows the cage after painting.

Tube Sockets

The original chassis had a conglomeration of whatever sockets David Berning had on hand. I replaced all of them except those used by the 6JE6 tubes. The reasons for keeping those sockets were twofold. First, the sockets were 9 pin Novar type, which are very difficult to find today. Some people state that the Magnoval sockets are appropriate replacements. However, the tubes designed for the Magnoval socket have larger pins and the smaller 6JE6 pins might not make good contact in the wider fingers. Second, the existing sockets were bottom mount with a lot of components already attached. The sockets could be removed from the original chassis with the components mounted and inserted into the new chassis. The rest of the sockets except for the 300Bs were replaced with Belton Micalex sockets. The 300B sockets were ceramic. All new sockets were top mount. **Photo 10** shows the top of the chassis with the new sockets mounted.

The early stages of mounting the tube sockets can be seen in **Photo 11**. You can see the identification labels I put for reference next to each tube socket and the tape markers for wires.

Photo 6: After deciding on the dimensions for the tube cage, I laid out a pattern on the material and cut it to shape.

Photo 7: I used strips of right-angle aluminum to assemble the tube cage base.

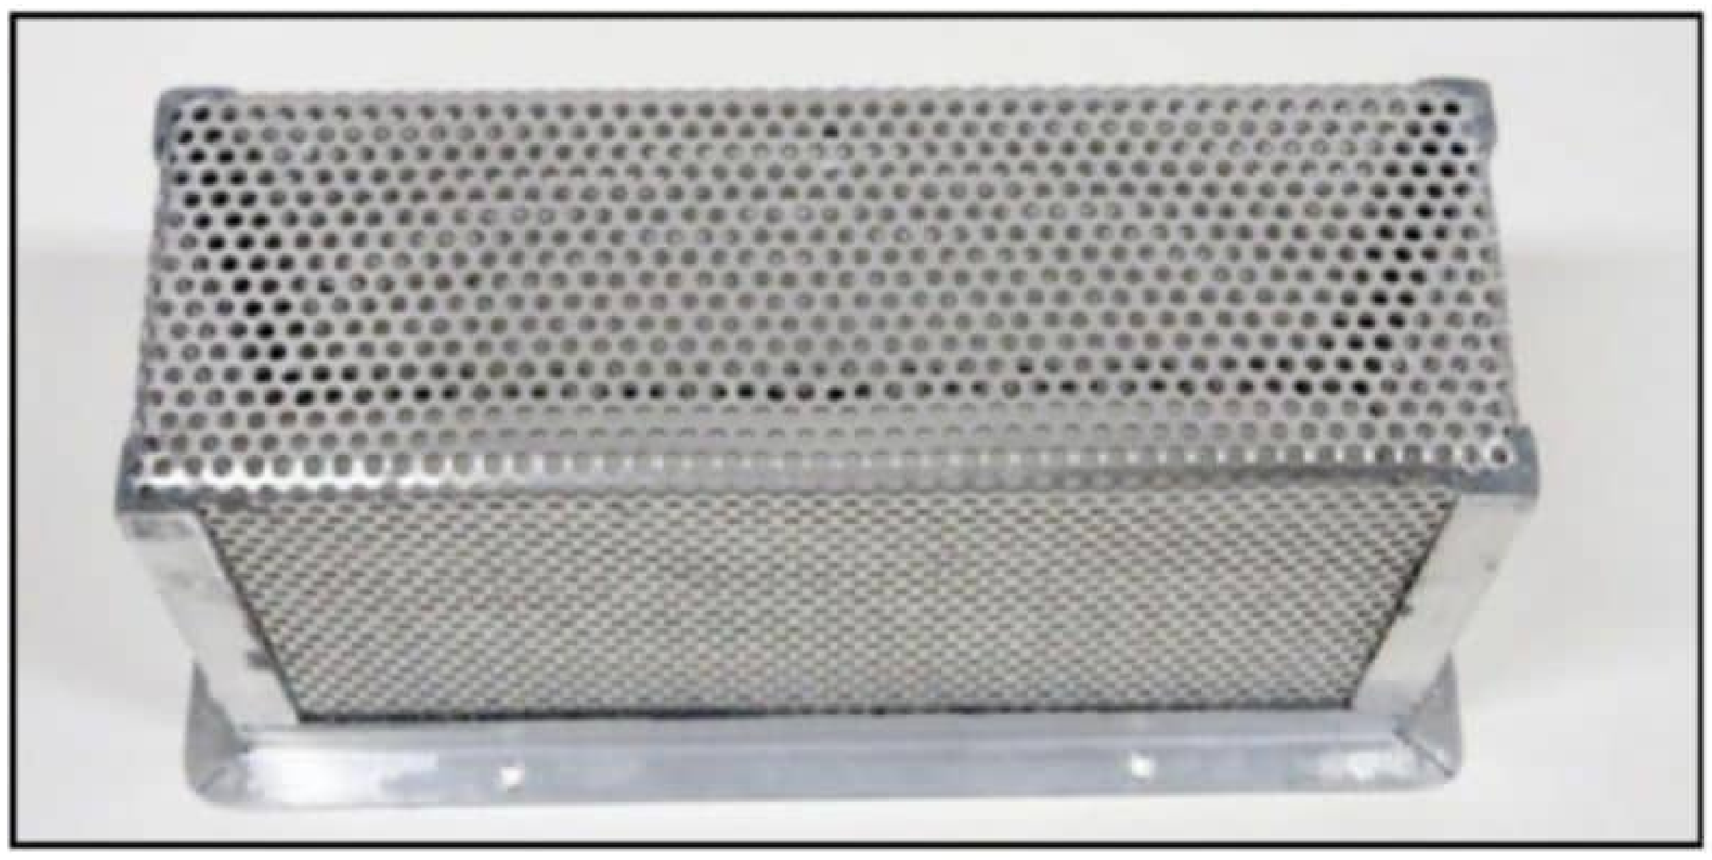

Photo 8: This is the completed tube cage.

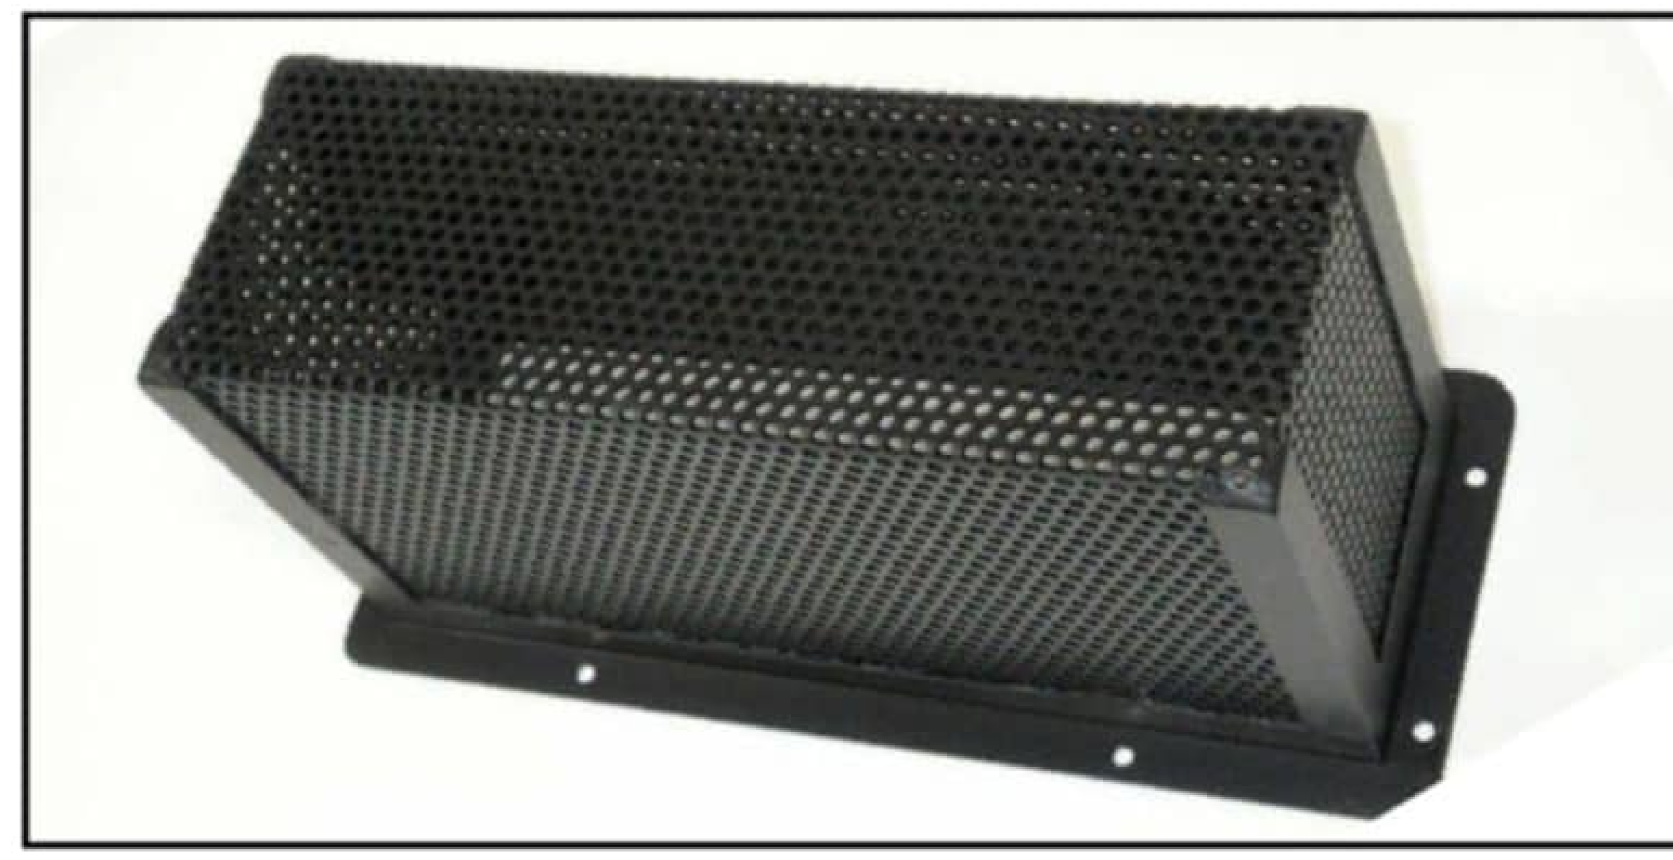

Photo 9: I chose Krylon Fusion black textured paint to complete the look of the tube cage.

You can also see some thin circles of heat shrink tubing and silicon insulating material on the bottom of the 300B sockets as additional protection from arc over between the pins and chassis.

Circuit Board Modifications

As previously mentioned, I tried to keep the circuit boards as original as possible. There were a few modifications that were necessary because of the smaller depth of the new chassis. The transformers in the impedance converters have their outputs rectified by strings of BAV21 small signal diodes that are on a PCB mounted in the center of the toroidal transformer. The same board is used for many different amplifier versions. As such, it has positions for the maximum number of diodes used in the largest amplifiers. In this amp, not all positions on the board were used, allowing me to shorten the boards to better fit the space available. The whole impedance converter board was originally mounted in a vertical orientation in the old chassis. In the new chassis, there was insufficient height for that orientation. Fortunately, there were a few rows of spaces in the board that were unused that only had wires passing over them. I was able to cut the board along one of those rows and fold it 90°. I had to make new brackets to mount the modified shape, but it worked well.

Another modification involved the power supply board. It was also mounted vertically in the original chassis, but to fit in the new chassis, it had to be mounted horizontally. A few power resistors that were wired between the 6JE6 socket pins to the board were also moved directly to the board.

About the Author

Thomas Perazella is a retired Director of IT. He received a Bachelor of Science degree from the University of California, Berkeley campus. Audio has been his passion for more than 50 years and he is a member of the Audio Engineering Society, the Boston Audio Society, the Philadelphia Area Audio Group, the DC HiFi Group, and the DC Audio DIY Group. He has written for *Speaker Builder* and *audioXpress* magazines, authored several articles in professional audio journals, and taught commercial lighting at the Winona School of Photography. Recently, he received a patent on a cost-effective high-efficiency LED lighting system for commercial and residential buildings. He is also a Past President and Treasurer of the Rockville Chapter of the Izaak Walton League of America, one of the oldest national conservation organizations in the US.

Mounting the Power Supply

Some of the larger power supply components are the power transformer, the large high voltage filter capacitors, and the choke between the two stages of the filters. Because of its weight and bulk, the power transformer was one of the last things to be mounted. The 6JE6 tube sockets with components already attached from the old chassis were mounted to the bottom of the chassis and connected to the power supply board.

The filter capacitors are axial leaded round cans, configured with a vinyl sleeve on the case which is connected to the negative polarity. For each section, there is a series/parallel arrangement with two upper caps in parallel connected to two lower caps in parallel. Two balancing resistors are across the first stage pairs with a center connection through a resistor to the center of the second stage to make sure all caps are balanced. Since the cases of the caps on the top half are at one half of the total voltage with respect to ground, David Berning used additional insulation between all the caps and the chassis before mounting them. Likewise, I cut some pieces of fish paper to fit between the caps and chassis. The paper was secured with Pliobond before mounting the caps to one of the chassis side pieces. The choke was mounted to another side piece and the leads were extended to reach the new arrangement of the power supply.

The Balance of Chassis Parts

The rest of the chassis parts included a newly purchased IEC input power connector containing a power socket, internal fuse holder with space for a spare fuse, and a power switch mounted on the right rear as viewed from the rear. On the left side of the rear were the input RCA jacks and speaker binding posts. There is a sub assembly that is shielded containing the jacks, binding posts, and output filter. I did not want to recreate that assembly so I drilled the appropriate holes and mounted it as is. As a result, the arrangement is somewhat unconventional with the RCA jacks vertical and the binding posts horizontal. **Photo 12** shows a bottom view of the completed chassis.

In the front, the volume control and an LED pilot light were mounted on opposite sides. The volume control is a 50K Goldpoint

V24 stepped attenuator. A problem occurred when trying to mount that control through the thick front plate. The threaded section of the mount was too short to allow the fastening nut to be attached. I had to remove some of the panel's thickness to allow the mounting. To do that, I used a hole saw large enough to allow clearance for the nut as well as a socket to cut a slot partly through the panel. Then I drilled a clearance hole in the center for the control. I used a Dremel tool with a flat-ended burr to remove the remaining material between the slot and hole. After that, the control was successfully mounted and a large matching aluminum knob added. Finally, the power transformer was mounted and the appropriate wires connected. Only after all the components were mounted were the tubes inserted and the cage mounted over the 6JE6s. To allow for servicing of the 6JE6 tubes, the holes in the top plate for the cage screws are tapped so they can be easily removed.

Photo 10: This is the top of the chassis with the new tube sockets.

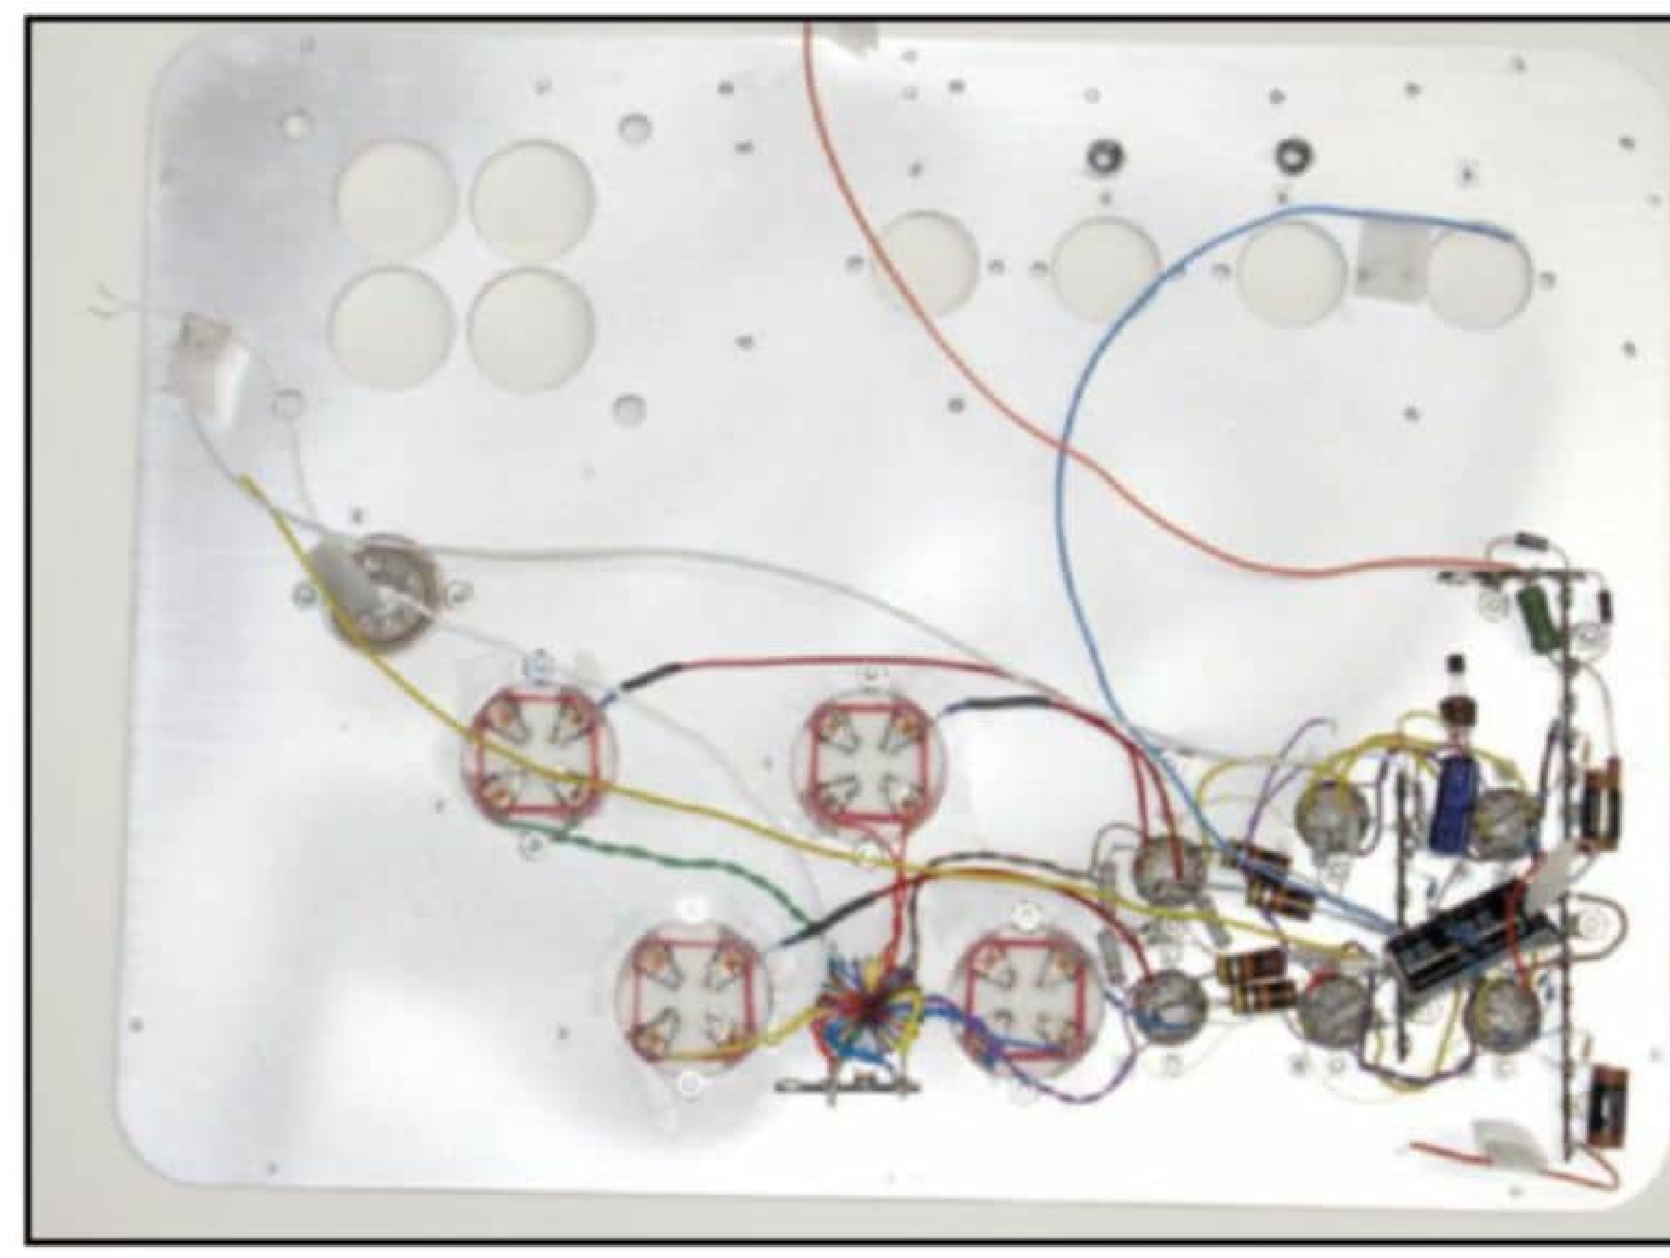

Photo 11: The bottom of chassis is shown with a partial wiring of sockets.

Initial Testing

As usual when first powering on a new component, I connected the amp and brought the voltage up slowly while measuring the input current. There were no excessive surges, which is always a good sign. However, there was also no high voltage on the tubes. As part of the design and a protection for the speakers, the amp has a servo circuit that looks for DC offset and shuts the power supply down and illuminates an LED if DC problems in excess of what the servo circuit can correct are found. It was illuminated.

After much work with an oscilloscope and several calls to David Berning, I found the problem. When working with high-frequency power circuits such as the switching power supply, radiated signals can cause problems. With the rearrangement of some of the components on the new chassis, I was concerned from the beginning that some of the new wire layouts could cause problems. In fact, that layout was the problem.

There are two operational amplifiers (op-amps) that act as integrators and inverters in the DC servo protection circuit. The first section was immune to the stray RFI because it necessarily had a very long time constant to properly function to only respond to very low frequencies. The second section simply functioned as an inverter and had no limitation on frequency response. It was picking up the stray RFI and triggering the protection circuit. The solution was to limit the bandwidth of the inverter section by adding a 0.1μf capacitor across the 100K resistor between the output terminal of the inverter and its negative input terminal. Problem solved. The amplifier then started normally.

Photo 12: Here is a bottom view of the completed chassis.

Before starting on this project, I had tested the amp on the bench and stored the waveforms. With the new amp, the waveforms were the same. I connected the amp to speakers and fed it from a music player. The sound was also the same as before so, from a sound standpoint, I considered the project a success.

The New Look and the Owner's Response

To get an idea of the tube locations on the new chassis, I photographed it from the front without the tube cage in place. **Photo 13** shows that view. And **Photo 14** shows a rear view detailing the input and output connections.

I did not send a photo of the complete amp to the owner so as not to spoil the surprise when he saw it. My wife went with me to the owner's office to deliver it. She was very impressed with the results and wanted to accompany me to see his reaction. It was very positive. He had comments like it "looks marvelous" and "cannot wait to get it going." It is always good to get positive reinforcement when completing a project.

DIY Excels Again

Having been involved in many DIY projects over the years, I am firmly convinced of the value in creating and implementing your own projects. However, as you can gather, a project such as this is not for the beginner. For starters, I was working on someone else's equipment, which was a unique piece. There is therefore great risk of irreparable damage if you fail. And now for the de rigueur section. There are lethal voltages in a device such as this, including both mains AC and high voltage DC. If you do not have a lot of experience working on this type of equipment, do not attempt it.

Inevitably, with any DIY project, you will also have some unexpected results. You will need not only a good understanding

Resources

"Analog Pulse Modulation," Tutorials Point, www.tutorialspoint.com/principles_of_communication/principles_of_communication_analog_pulse_modulation.htm

D. Berning, "Comparisons between the Berning ZOTL Impedance Converter and a high-quality audio output transformer," The David Berning Company, https://davidberning.com/technology/comparison

C. Hansen, "Power Transformer Parameters, Selection, and Testing," *audioXpress*, November 2022 (pp. 58–61); *audioXpress*, December 2022 (pp. 58–62); *audioXpress*, January 2023 (pp. 58 – 66), February 2023 (pp. 48-52); *audioXpress* March 2023 (pp. 56–61); and *audioXpress*, April 2023 (pp. 62–66).

"History of 300B tube," Audiophile Club of Athens (ACA), www.aca.gr/index/forums/fen/hiend2?row=2495

Sources

J-B Weld Twin Tube - 2oz
J-B Weld | www.jbweld.com/product/j-b-weld-twin-tube

Fusion All-in-One Textured Finish
Krylon |
www.krylon.com/en/products/all-purpose-spray-paint/fusion-all-in-one-textured

Photo 13: This is the front top view with no tube cage in place.

Photo 14: The rear view of the completed project shows the connectors.

of the circuit operation but also the necessary test equipment to troubleshoot and calibrate it. As an example of the importance of understanding of basic circuit operations that is not directly related to this project, but valuable none the less, I was recently repairing a high-powered amplifier for a friend using the service manual supplied by the manufacturer. In one section for setting a null value, the nomenclature in the manual did not match the part numbers of the devices in the schematic. The pin numbers of the differential amplifiers to measure certain voltages during the adjustment process were also wrong. As a result, the manual was essentially asking you to produce a 21V potential across an emitter to base junction of a PNP transistor, which is not possible if the device is functional. So caveat emptor.

On the bright side, DIY is one of the most productive and satisfying parts of the audio experience. So, if you are not currently a DIYer, start small, but do get involved. You'll love it.

Power Transformer Parameters, Selection, and Testing

Part 8 — Additional Transformer Core Types and Their Attributes

The initial group of articles detailed the history of transformer cores, various construction methods, their materials, a few testing methods, their losses, commercial power quality, and the design of current transformers. In Part 8 of this series, our author focuses on additional transformer core types and their advantages and disadvantages.

By
Chuck Hansen

Toroidal Cores

Toroidal cores are similar to C and E cores in that they are wound from a continuous strip of silicon steel or other magnetic alloys. They are circular rather than elongated ovals and don't have any core gaps. They are available in a number of different cases, as shown in the cross-sectional cutaway views shown in **Figure 23**.

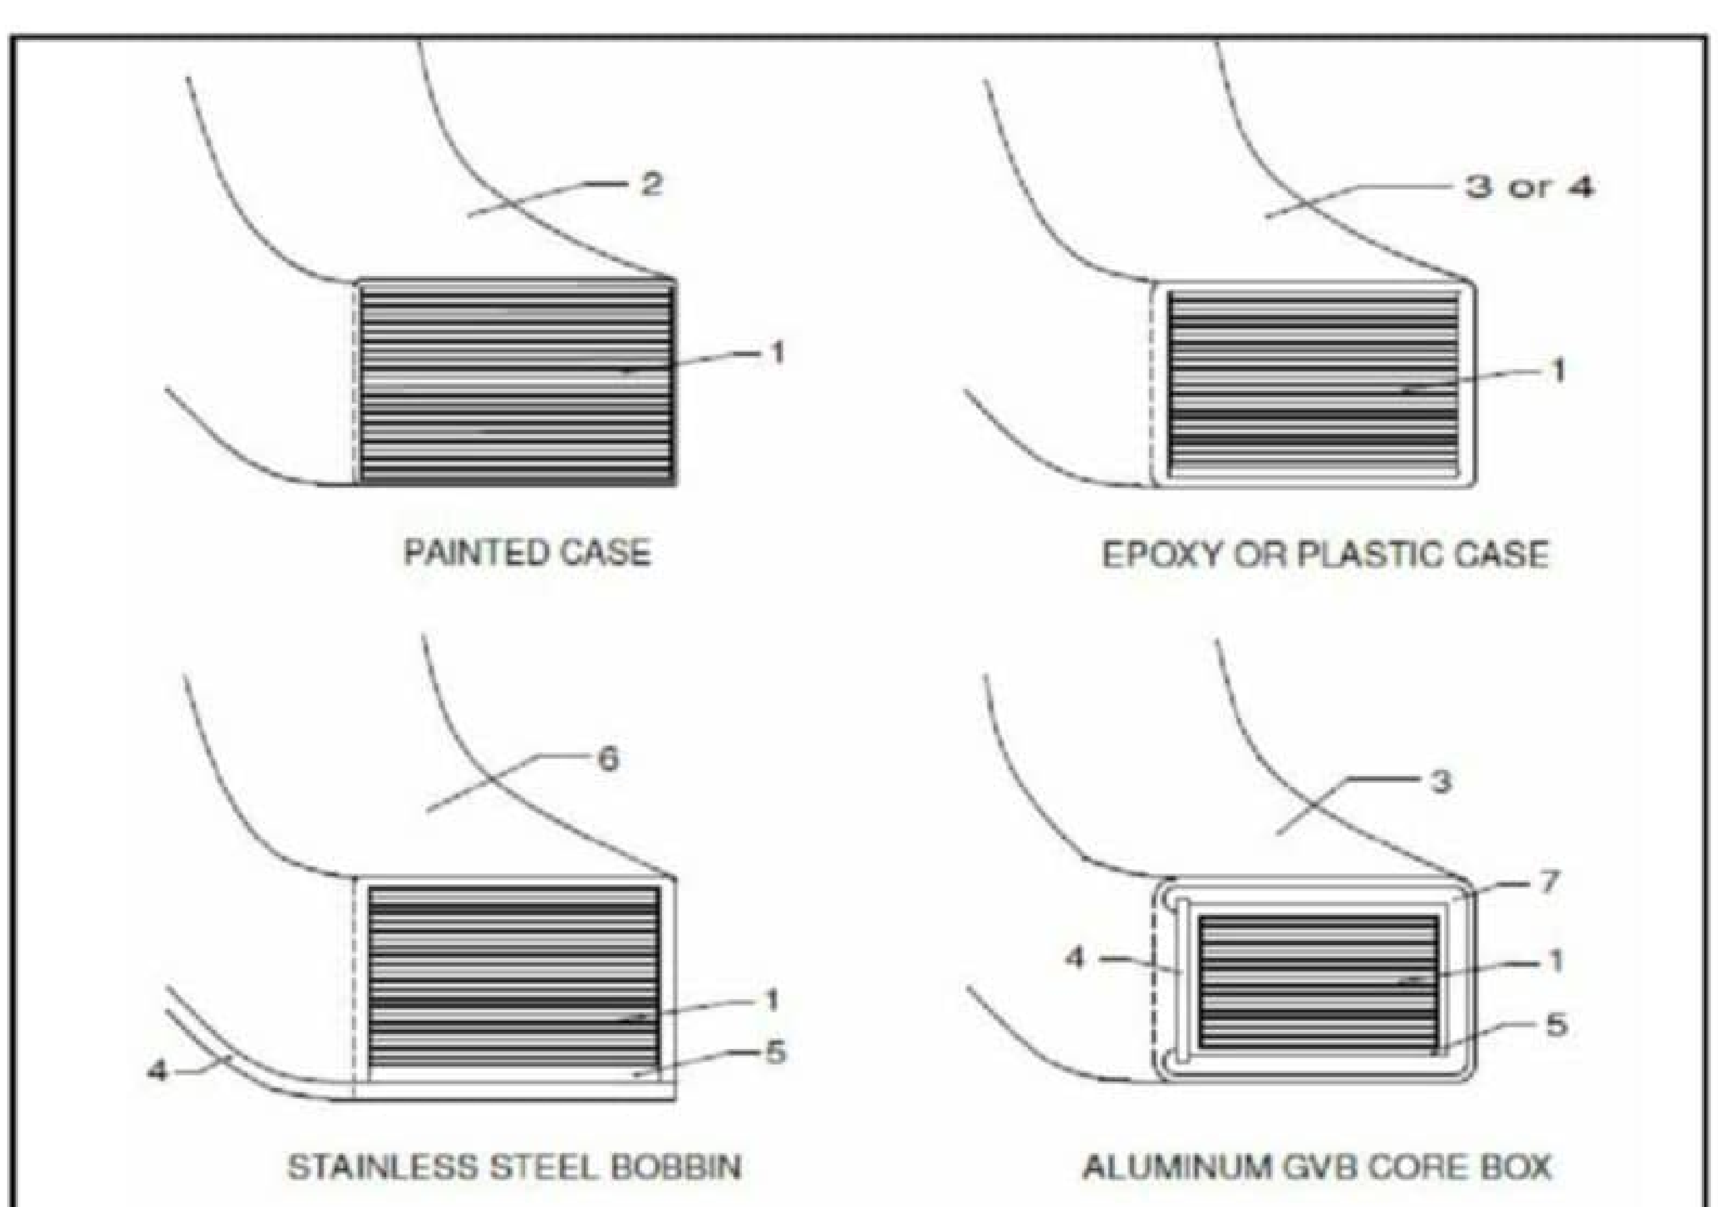

Figure 23: Cutaway views of toroidal transformer core/case construction: 1- Tape wound core, 2– High temperature paint, 3- Epoxy coating, 4- Plastic case (phenolic or glass reinforced nylon), 5- Silicone damping medium, 6- Non-magnetic 316 stainless steel (used for nickel alloys), 7- Anodized aluminum GVB (guaranteed voltage breakdown). (Image Source: Magnetics Inc., Design Manual TWC-300U, Core Box Selection; www.mag-inc.com/Products/Tape-Wound-Cores)

The stainless steel and aluminum core boxes that are used with nickel alloy cores do not wrap around the full cross-section of the core, since that would make a single-turn short-circuit. The metal part of the case is interrupted on one side with an insulated plastic spacer.

Specialized programmable machinery is required for automated winding of a toroidal core. The ring-shaped toroidal core is clamped horizontally between three elastomer drive wheels (**Photo 18**). The winding head has a vertical circular shuttle with a split in it that allows the shuttle to be slipped over the core.

In order to wind the wire around the toroid, the winding machine has to first load a known length of wire onto the shuttle by passing it through the center of the toroidal core. The winding head is equipped with its own set of drive wheels. Once the shuttle is loaded, it's direction of rotation is reversed and wire is "un-wound" onto the toroidal core as it is rotated at the rate that ensures uniform winding around the inner circumference of the core. This results in the wire on the outer circumference of the core being spaced farther apart than on the inside. There is a YouTube link (www.youtube.com/watch?v=82PpCzM2CUg) that shows this winding process.

As the core rotation completes each full turn, new inner turns are added uniformly on top of the layers below them. This ensures the finished transformer ID will be as circular as possible, since they have to be concentric with the transformer mounting gasket and support hub.

It does get more complex if the toroidal transformer is bi-filar wound, with two windings "hand-in-hand" as shown in Photo 18. We have some current transformers wound that way to provide identical signals to two circuits that don't share the same ground and power supply. You have to wind two identical lengths of wire onto the winding shuttle at the same time, then "un-wind" them side-by-side onto the core with the same tension on each wire, for up to 1,000 turns. Winding heads are available for both layer winding and random winding (**Photo 19**).

Advantages and Disadvantages of Toroidal Transformers

Advantages of Toroidal Cores

Toroidal transformers have high efficiency. The tape wound 3% grain-oriented silicon steel core has all the grains in the preferred magnetic flux direction, unlike the standard E-I transformer, with about 40% in the wrong direction. There is no air gap, resulting in a stacking factor of up to 97.5% of its weight, and a high power-to-weight ratio.

Toroidal transformers operate at lower temperatures than other transformer core types with similar specifications. The copper windings closely surround the silicon iron core, and, being an excellent thermal conductor, help dissipate the core losses.

Since all the windings are symmetrically spread over the entire round, gapless core, a higher maximum flux density is possible. Toroids can operate at flux densities between up to 1.8 tesla (18 kilogauss), whereas E-I transformers operate between 1.1 and 1.5 T. High flux density allows for more consistent electrical efficiencies.

Toroidal transformers are also acoustically quieter than E-I transformers. The absence of an air gap typically provides an 8:1 reduction in acoustic noise. The core is only in compression or tension, the magnetic forces do not exert a bending moment on the core, so the circular shape is more stable mechanically. In addition, the windings tightly envelope the entire core, which effectively reduces magnetostriction, the main source of the familiar mechanical hum at twice the line frequency from standard laminated E-I transformers.

Toroidal transformers radiate about 10% of the stray magnetic field of E-I transformers because of the inherent efficiency in its construction. The windings cover the core and act as a shield. The

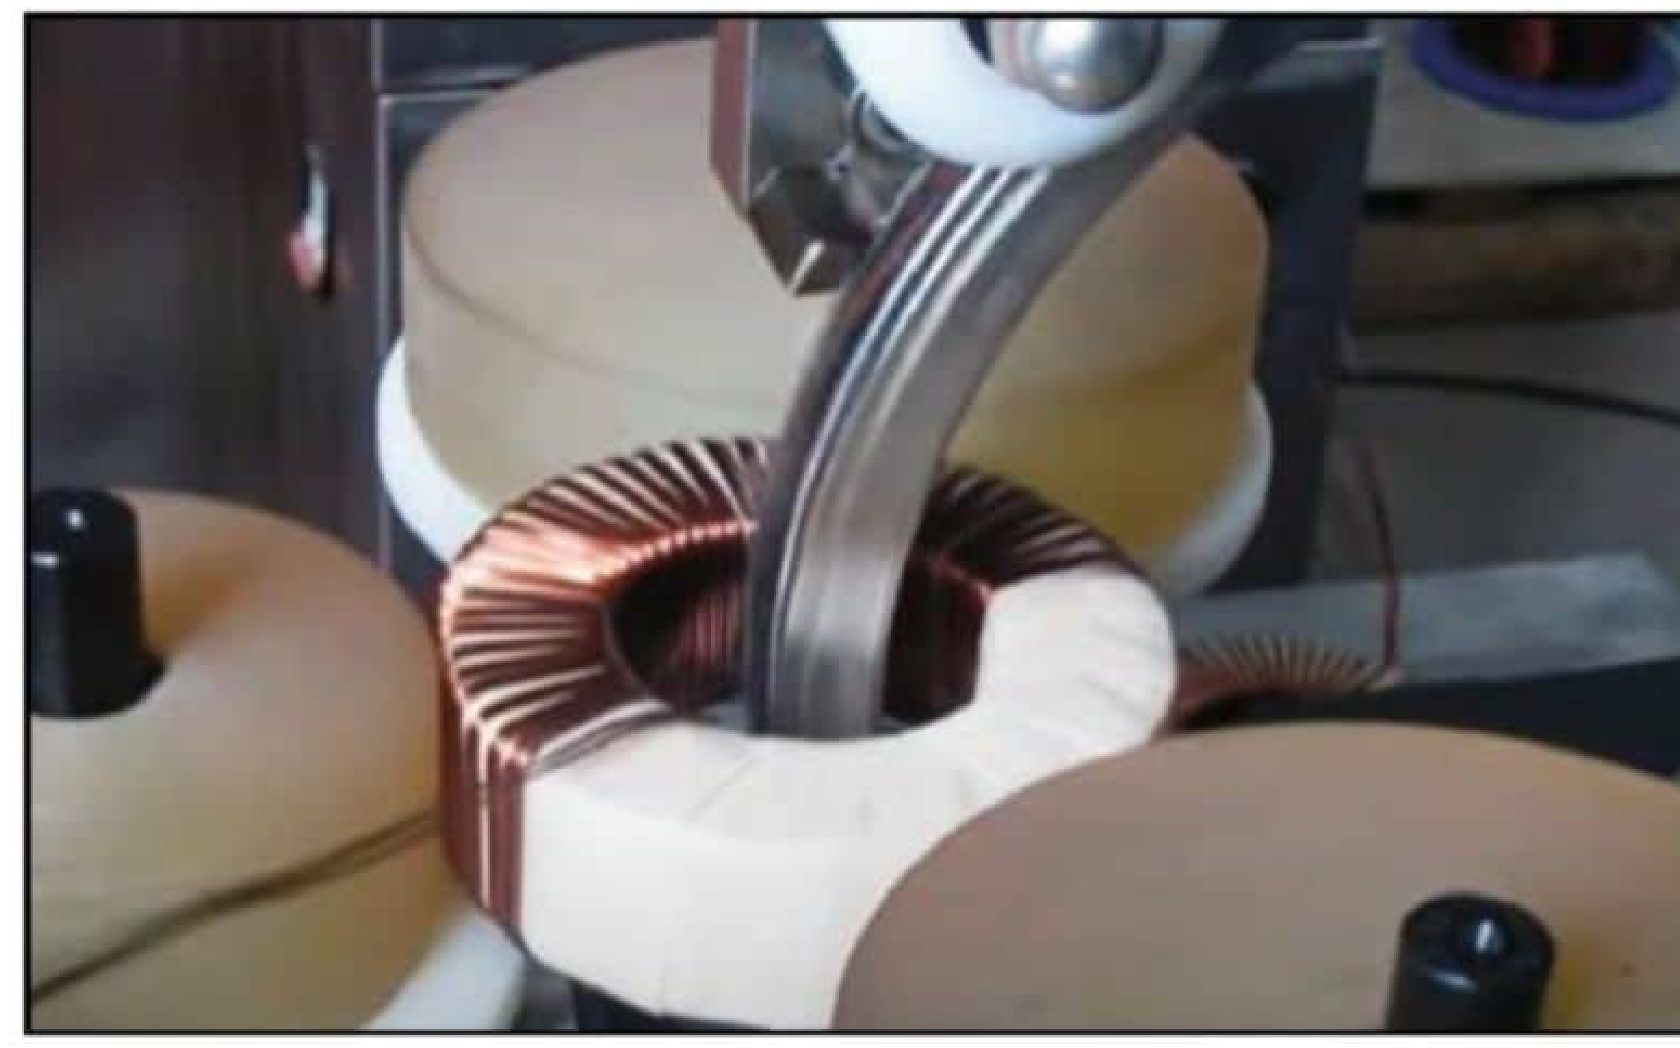

Photo 18: Toroidal Core Winding Machine (Image Source: VolMetAPG, Eda Wu)

magnetic interference field is contained to transferring energy from primary to secondary. This may eliminate the need for special shielding and makes toroidal transformers especially suitable for applications in sensitive electronic equipment, such as low-level preamplifiers and medical equipment.

Toroids require as little as 1/16 the steady-state no-load excitation current of conventional E-I transformers. This can translate into big savings in large transformer applications such as industrial controls.

Another mechanical benefit of toroidal transformers is their flexible dimensions. With the cross section held constant, the height and diameter can be varied to accommodate the application. For instance, for 400Hz applications there are up to 11 different combinations of 4-mil core inner diameter (ID), outer diameter (OD), and height available for the same net cross-sectional area. The largest standard core OD is 7" at 2in^2 core area.

For 60Hz applications there are fewer different combinations of 12-mil core ID, OD, and height available for the same net cross-section area, but the largest core OD is 6.210" at 2.25in^2 core area. Custom designs can be made with very large dimensions.

If you are purchasing a toroidal transformer rather than winding your own, there are VA ratings available from 0.5VA to 10kVA, manufactured by Acme, Avel Lindberg, Hammond, Plitron, Triad, and others.

The small size saves real estate. This efficiency yields the most easily identifiable feature of the toroid, its size benefit. Toroidal

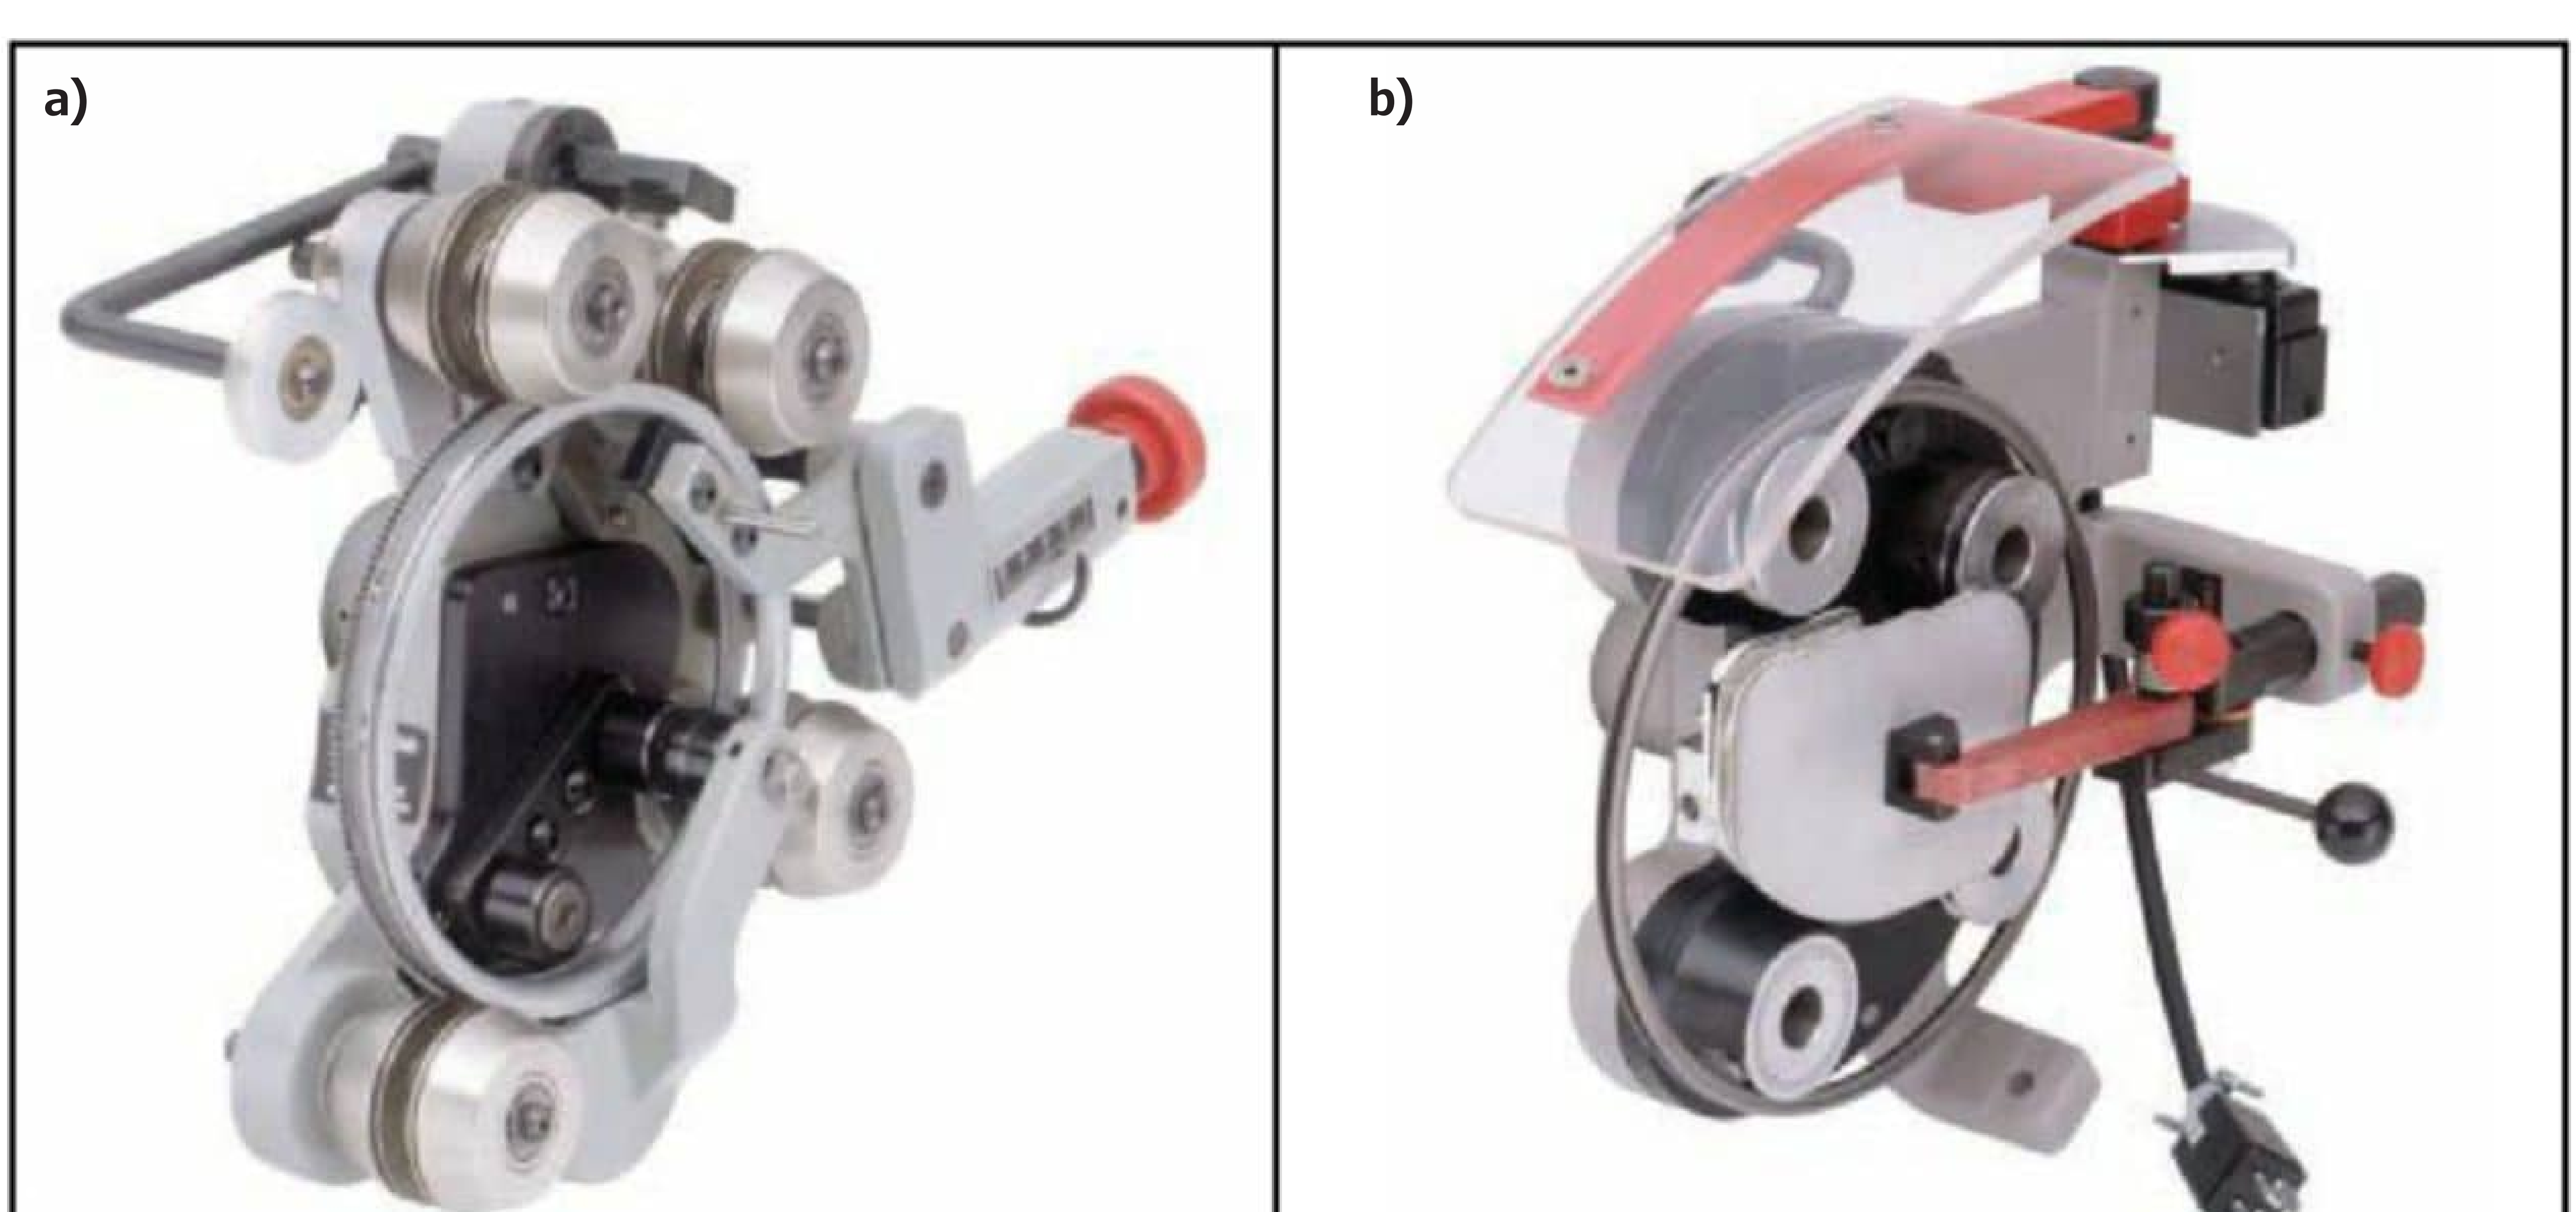

Photo 19: Jovil Universal 4U Toroidal standard layer-winding head (a); and a Jovil 4 toroidal random-winding head (b).

Photo 20: Sheet metal manufacturing with electrical discharge manufacturing (EDM) (Image Source: SK Electric)

transformers are about half the size and weight of standard E-I transformers of the same VA rating, so they are easily integrated into pre-existing designs. Only the R core and O core have less magnetic material for a given core area (more on those later).

Toroidal transformers are easy to mount. A single center bolt is all it takes to mount a toroidal assembly. Quicker, easier mounting means reduced assembly time by the equipment manufacturers who select toroidal transformers.

Disadvantages of Toroidal Cores

Toroidal transformers have higher inrush currents than E-I types. When a toroidal transformer is first energized, a transient current of up to 10 times the rated transformer current can flow, depending on the transformer efficiency and impedance of the power source. Worst-case inrush happens when (1) the primary winding is energized at the instant the primary voltage crosses zero (see Figure 14 in Part 5, *audioXpress* March 2023), and (2) the polarity of the remanence in the iron core, Br, has the same polarity as the voltage half-cycle if it was left at the highest point when the transformer was previously de-energized (see Figure 8 in Part 4, *audioXpress* February 2023). In this case the core at start-up will initially be saturated, the winding inductance is reduced to near the leakage inductance value, and only the primary winding and power source resistances will limit the inrush current. Harmonic-rich waveforms can be generated that may cause problems to other connected equipment until the transformer comes out of saturation. This can cause a momentary magnetostriction noise in the core.

Photo 21: Comparison of toroidal core and O core transformers (Image Source: Bridgeport Magnetics, bridgeportmagnetics.com)

It is inherent difficulty to wind wire through the center of a torus. Unlike a split core, like C cores or E cores, specialized machinery is required for automated winding of a toroidal core, as shown in Photo 18 and Photo 19.

Winding the coil evenly, especially with a large number of turns and less space available for turns inside a smaller inner diameter (ID) circumference, either (1) the winding ID will have jumbled turns stacked on each other, or (2) the OD turns will be spaced widely apart in order to have a smoothly layered ID.

Toroidal core cross-sections are not necessarily square, depending on the ratio of the core tape width to the wound outer diameter. The ID, OD, and height come in 0.125” increments.

O Cores and R Cores

O cores and R cores are also tape-wound cores. The core tape is cut from flat sheets of cold-rolled grain-oriented silicon-steel.

O Core Construction

A silicon-steel sheet is slit by means of a computer-assisted electrical discharge manufacturing (EDM) device (**Photo 20**). The width of the slit continuously changes from narrow at each end to widest near the center, so that when it is wound on a mandrel to form a core, it has a net circular cross-section with the required ID, OD, and cross-sectional area. More on this in the description of R core construction.

The rectangular cross-section of a toroidal transformer is taller than round cross-section of an O core transformer with the same VA rating (**Photo 21**).

Advantages and Disadvantages of O Cores

Advantages of the O core

The O core windings can be wound with the same equipment as a toroidal transformer. The circular cross-section of the O core (and its R core relative) provides the lowest circumference for a given area. The circumference of a circle is about 12% less than that of a square with same area, and 18% less than a rectangle with the same area and a diameter to height ratio of 2:1, like a typical toroidal core. This reduces the total wire length in the transformer, results in the lowest winding resistance and lowest copper loss.

O core transformers have the advantages of lower flux leakage, lower exciting current, lower overall losses, and higher efficiency than toroidal core transformers with the same VA rating.

Disadvantages of the O core

The biggest disadvantage of the O core is its price. It is much easier and faster to slit fixed widths of core tape from a roll of silicon steel, than to EDM machine flat sheets into varying widths and lengths of core tape that will result in the desired number of different O core sizes. It is up to the transformer user to decide if the higher cost justifies the performance. Like toroidal transformers, O cores have higher inrush currents than E-I types.

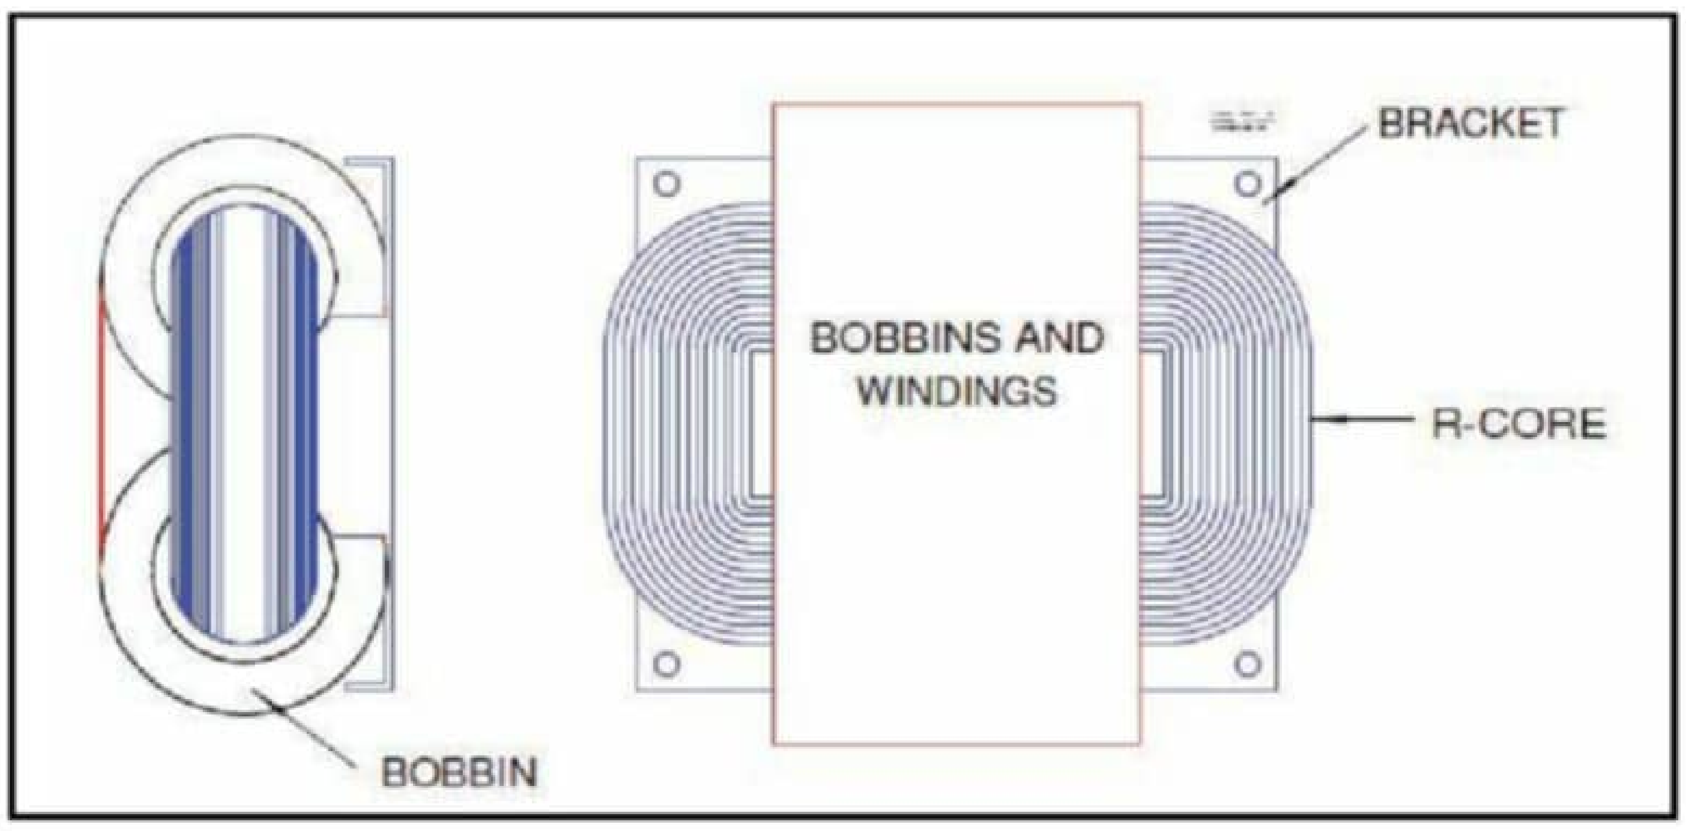

Figure 24: R core transformer construction (Image Source: Elliot Sound Productions, SK Electric)

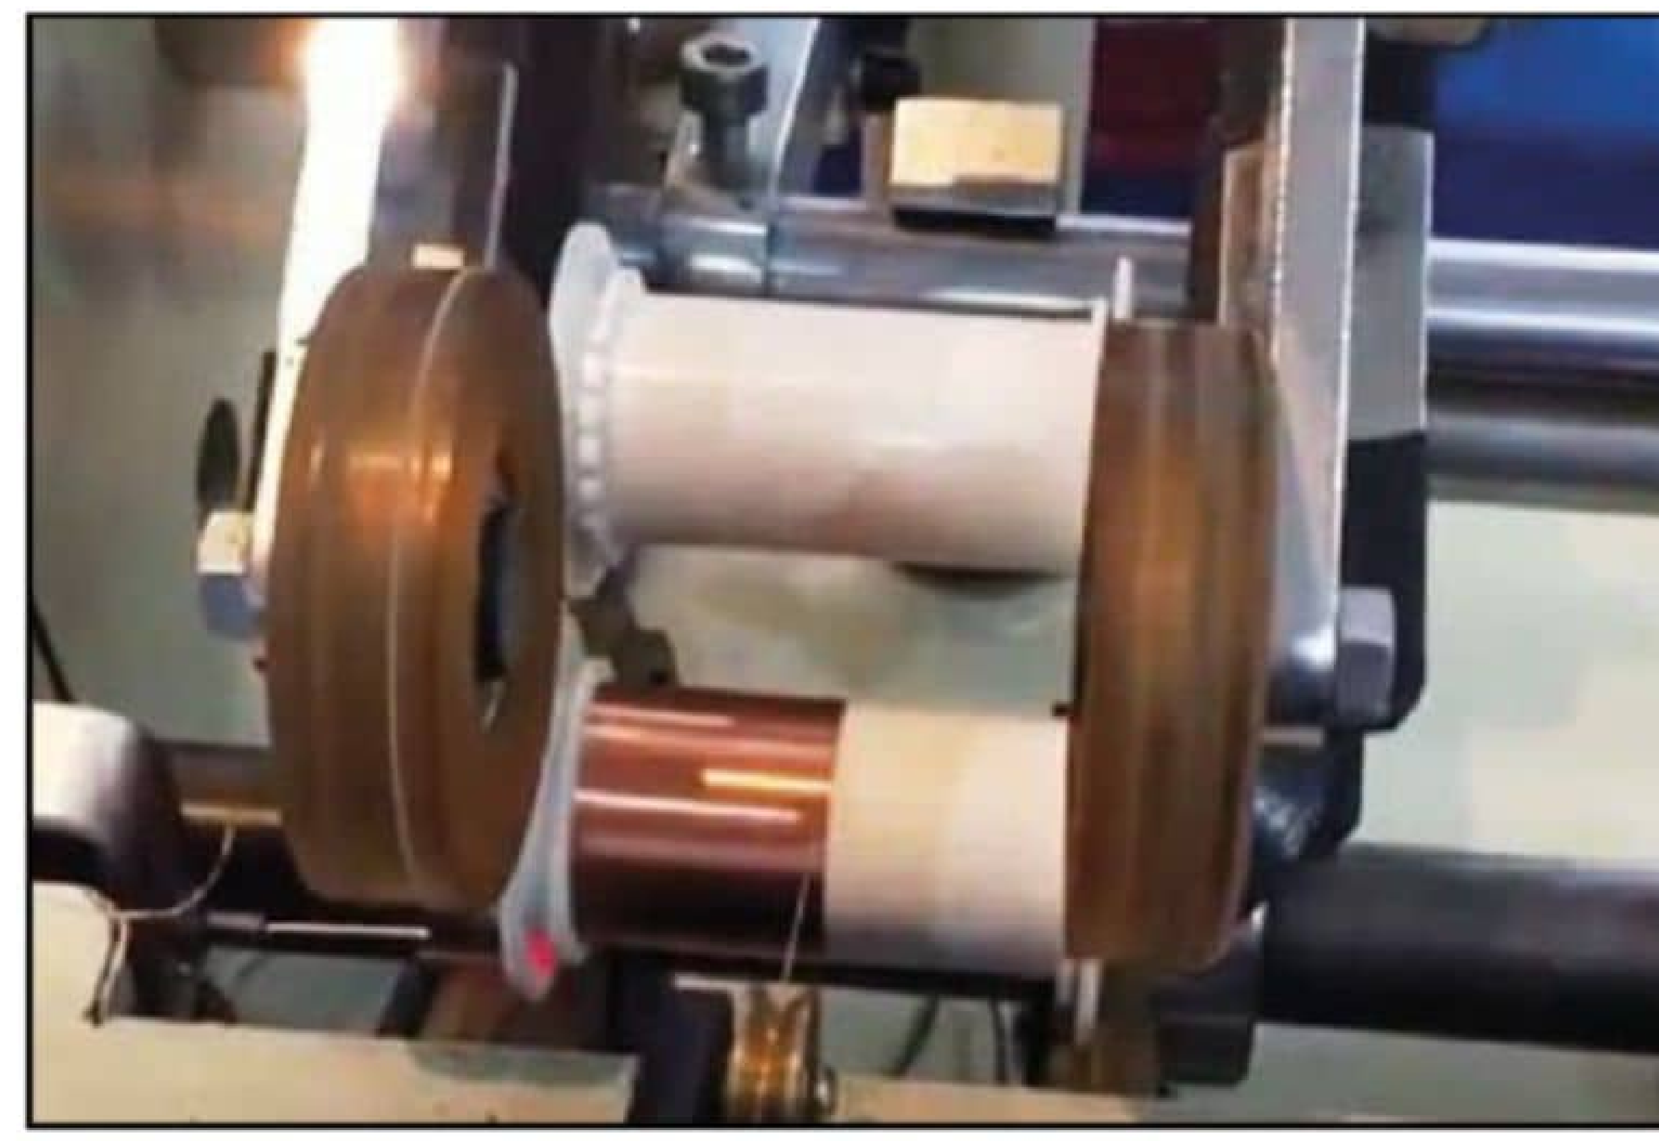

Photo 22: R core Winding Machine (Image Source: Zhongshan Shili Winding Equipment Co., Ltd., www.zsshili.com)

R Core Construction

Like the O core, R cores are wound from a continuous strip of variable-width silicon-steel, but formed into a shape with two straight sides and two curved ends, so the finished core winds up with a circular cross section (**Figure 24**).

Unlike a C core, the R core is not sliced in half to assemble the windings. Instead, the windings are wound on a pair of two-piece polybutylene terephthalate (PBT) plastic bobbins that are assembled over the straight sections of the core. They are then free to rotate while the core is clamped in a fixture.

Unlike a toroidal winder, where the core is rotated in a circle at a rate that ensures uniform winding around the inner circumference of the toroidal core, gear teeth on the R core bobbin end plates are rotated at both ends by large rubber drive wheels. At the same time the winding machine moves linearly back and forth across the straight section of each bobbin (**Photo 22**).

The end plates also prevent the wire at each end from tumbling off the stack-up of the winding layers. This process has to be repeated for each secondary and primary winding. Some R core transformers have half the primary and half the secondary turns wound on each bobbin. Then either the winding directions or the connections from one bobbin to the other need to follow the core flux path. Other R core transformers have the primaries on one bobbin and the secondaries on the other bobbin. Finally, the magnet wires are terminated with PVC or other stranded-insulated lead wires in the inner notched area of the end plates.

I have included a reference to a YouTube video (www.youtube.com/watch?v=WCZjsw_WDYs) that shows how an R core transformer is wound.

Once the windings are completed and wire leads added, a steel mounting bracket with rounded vertical legs is installed up through the curved ends of the core. The rounded legs are twisted and bent against the ends of the bobbin end plates to hold them in place (**Photo 23**). A potting compound is added to reinforce the whole transformer assembly.

Advantages and Disadvantages of R Cores

Advantages of R Cores

The circular cross-section of the R core (and O core) provides the lowest circumference for a given area. The circumference of a circle is about 12% less than that of a square with same area, and 18% less than a rectangle with the same area and a diameter to height ratio of 2:1, like a typical toroidal core. This reduces the total wire length in the transformer, results in the lowest winding resistance and provides the lowest copper loss.

R core transformers have the advantages of lower leakage flux, lower exciting current, lower overall losses, and higher efficiency (up to 97%) than E-I and C core transformers with the same VA rating. In comparison with E-I transformer of same rating, the R core may allow designers to place transformers closer to other components because of the lower operating temperature and lower leakage flux.

Because of its high efficiency, low leakage flux and low electrical noise, the R core transformer is becoming popular for use in high end audio equipment.

Disadvantages of R Cores

The biggest disadvantage of the R core is its price. As with the O core, it is much easier and faster to slit fixed widths of core tape from a roll of silicon steel. To accurately produce the variable-width strip for the R core (and O core) transformers, SK Electric developed and patented a unique automated cutting technique and set of equipment. Like toroidal and O core transformers, R cores have higher inrush currents than EI types.

The bobbins need clearance around the core to be able to rotate at a fairly high speed without skidding on the core while winding the bobbin. This added diameter increases the total length of the windings.

Some sources claim it would be more difficult to produce high power transformers by this method.

Photo 23: R core toothed gear-drive bobbin end plates and notched splicing area on a finished transformer. The curved steel mounting plate retainer lugs are visible above the curved ends of the core.

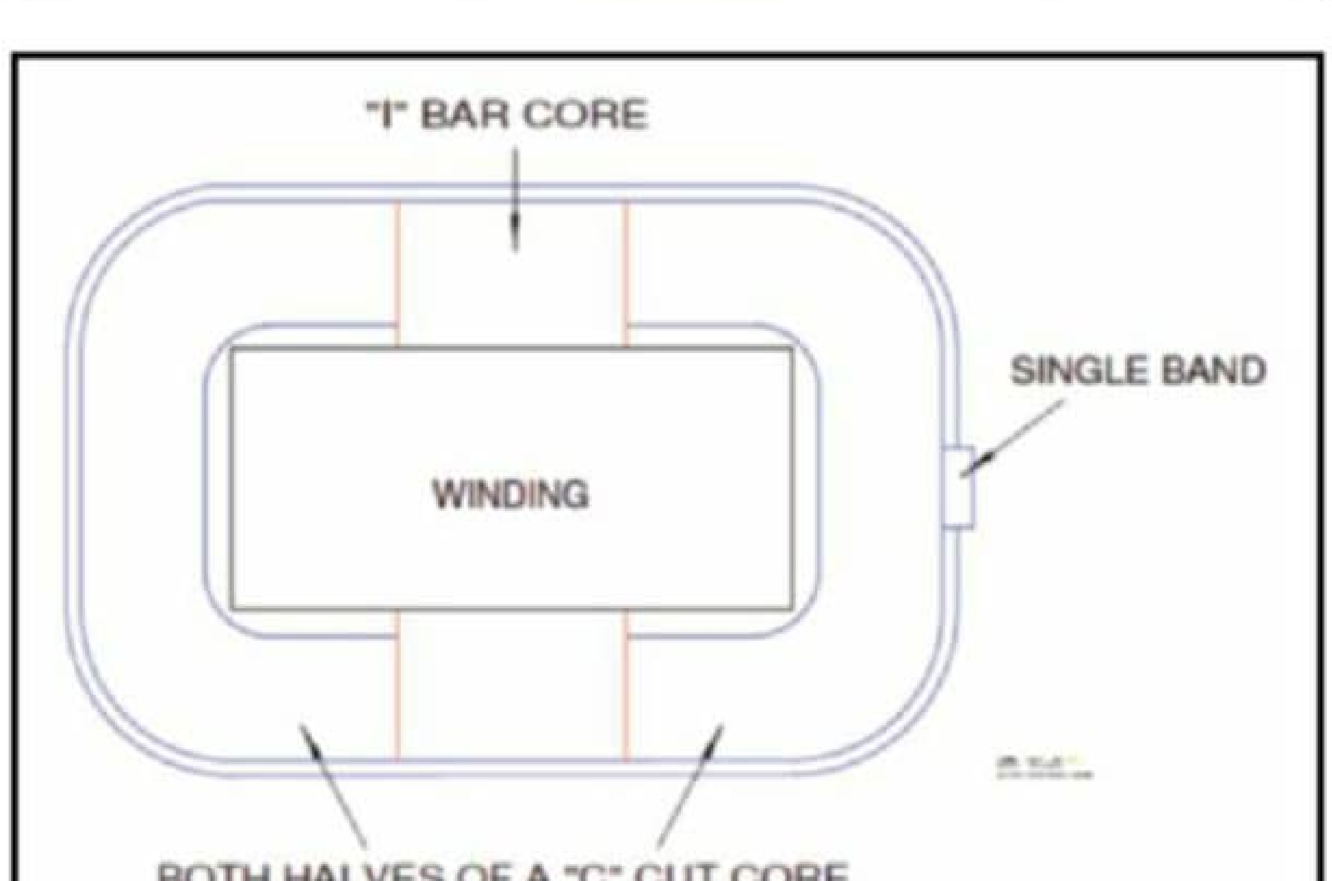

Figure 25: A C-I core is made up of two "C" sections of a tape-wound cut core and one laminated "I" bar put together (Image Source: Magnetic Metals)

Hybrid of Laminated and Tape-Wound Cores

C-I Cores

The C-I core configuration is made of two sections of one tape-wound cut core (rectangular or toroidal), and one laminated "I" bar on which the windings will be placed using a pre-wound bobbin, or even wound directly on the "I" section, as shown in **Figure 25**.

It combines the advantage of low-cost flat laminations ("I") and those of standard cut cores ("C") which are more expensive to manufacture but readily available in many standard sizes.

Advantages of C-I Cores

C-I cores use the same available core sizes as existing tape-wound C-cores, and inexpensive tooling for the laminated "I" section.

Windings can be wound on bobbins or core tubes. Since the "I" bar is removable, the four edges can be smoothed over (optional) and the copper wires wound directly on the insulated "I" bar. This could reduce the wire length and ultimately reduce the resistance of the winding, reduce the cost and simplify the production of prototypes significantly. It is easy to assemble using existing C-core bands and seal clips.

For transformers or inductors requiring air gaps, insulation can be placed onto the "I" sections to be used as air gaps. The tolerance stack-up only applies

Photo 24: Amorphous metal foil sheet (Image Source: Metglas)

to one cut C core instead of two as in the shell-type, and the "I" bar is made of stacked up lamination in which the tolerance can be more accurately controlled to fit the bobbin.

A shell-type transformer using double C cores requires clamping bands and seal clips on each individual core. The C-I core is held together by the same number of bands and seal clips as a single C core transformer.

Disadvantages of C-I Cores

It has wider mechanical gap tolerances than cut laminations or two C cores. While the I core grain direction is the same as the mating C core halves, the laminations meet them at a less efficient 90° angle.

Amorphous Cores

Metglas Amorphous Cores

In 1973 AlliedSignal scientists developed the Metglas amorphous alloy. AlliedSignal bought Honeywell in 1999 and took the Honeywell name. Hitachi Metals, Ltd. (Japan) purchased Metglas from Honeywell in 2003 and continues development and production to this day.

Amorphous cores are wound from melt-spun sheets of amorphous metal foil (**Photo 24**). Metglas Amorphous Metals have a unique non-crystalline nanocomposite structure and possess excellent physical and magnetic properties that combine strength and hardness with flexibility and toughness. They are typically made of iron, silicon, boron, copper, and niobium.

The melt-spinning process (**Figure 26**) forms thin ribbons of metal alloys with a particular atomic structure. The process involves casting molten metal by jetting it onto a rotating drum, which is cooled internally, usually by chilled water or liquid nitrogen. Cooling is sufficiently rapid from the molten state through the equilibrium melting point, that crystallization is avoided. The amorphous ribbon is then shaped by a series of rollers and dies to the desired thickness and width.

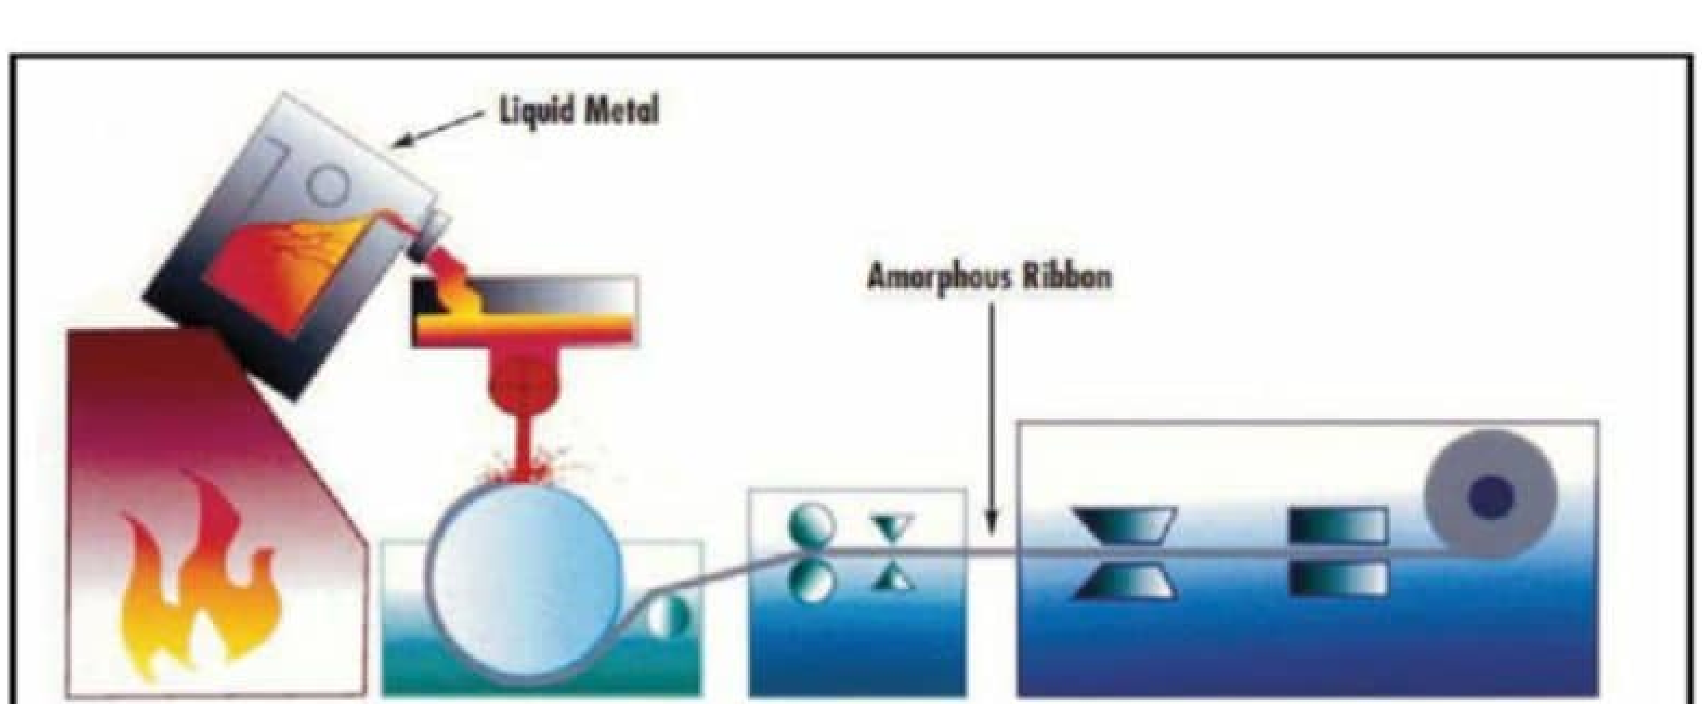

Figure 26: Amorphous metal melt spinning process (Image Source: Metglas)

Advantages of Amorphous Cores

Amorphous alloys can be annealed to obtain either round-loop or square-loop characteristics. The core losses are 20% to 30% lower than silicon steel, so they operate at a lower temperature.

With the increased overall efficiency, they are widely used in high-efficiency 50Hz and 60Hz utility pole-top, commercial, industrial and distribution transformers.

Disadvantages of Amorphous Cores

A larger core cross-section is required for same power rating compared to silicon steel laminations. This is due to the lower maximum flux density. This also causes added weight and size, but this is less of a concern with permanent transformer installations. With the larger core cross-sectional area, longer winding conductors are needed, which causes increased copper losses for the same current rating.

The low curie temperature, the temperature at which its maximum flux density starts to rapidly decrease, limits it's use to applications below 150°C (see Figure 2 in Part 1, *audioXpress* November 2022). The degree of sensitivity to temperature changes depends on the composition of the magnetic alloy, other added metals, the level of impurities, and on the alloy processing from ingot through finished sheet.

The NASA John Glenn Research Center is currently working on a novel nanocomposite soft magnetic material with an operating temperature limit of up to 400°C, with minimal increase in core loss. The new material uses small amounts of cobalt, and tantalum in place of the niobium, and are annealed at high-temperature in a magnetic field.

Next Month

Next month in Part 9 we will learn about National Electric Code Transformer Wire Lead Color-Code Standards, then discuss the effects on efficiency when operating a 60Hz transformer both normally (line voltage input to the lower secondary voltage output) and backward (rated secondary voltage stepped up to the primary winding voltage). We will also perform the same tests at 400Hz.

About the Author

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