

audioxpress

ADVANCING THE EVOLUTION OF AUDIO TECHNOLOGY

R&D Stories

How Optical Technology Enables a Generational Shift in MEMS Microphone Performance

By Jakob Vennerød and Kieran Harney (sensiBel)

Development & Design

The Impact of Mic Specs on Active Noise Cancellation Performance

By Nikolay Skovorodnikov (Knowles)

How ADC Hybrid Gain Control Can Enable Smooth, Pop-Free Microphone Preamp Gain Changes

By Michael Chandler-Page (Cirrus Logic)

Digital Signal Processing for Over-the-Counter Hearing Aids

By Alexander Goldin (Alango Technologies)

Audio Electronics

Distortion Measurement of OPA1656 and OPA1612

By Vojtěch Janásek

R&D Stories

Motion Feedback (MFB) for Active Desktop Subwoofers

By Zami Schwartzman

Sound Control

Over a Century of Home Loudspeaker Development

By Richard Honeycutt

Automotive Audio

Easing Deployment of Automotive Audio Functions

By John Whitecar (DSP Concepts)



Don't dream it...



Live it.

So real it sounds
like the orchestra
is in the room
with you!

The Ultimate Performance
MUSES05
Audio Op Amp

Single channel J-FET input
Very low noise: 7.5nV/√Hz
Ultralow distortion: 0.00003%
High Slew Rate: 40V/μS
Wide GBP: 12MHz





MANUFACTURING / DESIGNING TRANSFORMERS & ENCLOSURES SINCE 1917





November 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

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

© KCK Media Corp. 2023
Printed in the US

Audio DSP Intelligence

For the past five years, the audio industry has witnessed a tectonic shift in the platforms for digital signal processing available to developers of audio products. Now, we are seeing another significant evolution as the industry migrates to extremely sophisticated use cases of edge AI and on-device resources. All the while, users’ expectations of audio features continue to increase.

Voice recognition and voice engines were one of the biggest incentives for manufacturers to boost the processing power of built-in DSP and audio systems. Voice applications required embedded signal processing to ensure that voice signals that originated in even the most unpredictable conditions and captured by microphones could be clearly recognized. But the original voice assistant implementations have almost always depended on connectivity to the cloud and super-powerful centralized computing infrastructures. And those have now been proven to be unsustainable as a business model, and for any audio manufacturer to offer without recurring costs.

Fortunately, those same efforts related to voice have contributed greatly to the sophisticated systems that we now see embedded in true wireless earbuds and consumer wearable designs, where there is simply no connectivity to offload any of the processing. Users expect increasingly more advanced features to simply work—and all in the context of ultra-low power requirements. The considerable DSP power existing in those platforms is also now benefiting new generations of consumer audio products able to accept simple voice commands processed on-device, and which at the same time are able to run sophisticated room correction, noise reduction, and previously unimaginable algorithms.

A remarkable example was recently displayed by Apple with new features for its existing AirPods line of products, simply “activated” through a firmware update. Benefiting from its own silicon and full control of the hardware and software integration, Apple introduced Adaptive Audio, a term that defines a new listening mode that dynamically blends Active Noise Cancellation and Transparency (pass-through) to tailor the noise control experience based on the conditions of the user environment—in real time.

Naturally, very soon users will expect other systems to be able to offer this processing power and similar features. But not all manufacturers have access to an integrated Apple H2 SoC that uses cutting-edge 7nm technology with lower power consumption (Apple’s previous chip, the H1 used 16nm process nodes—and competitors in the space mostly still use 28nm technology). As any audio manufacturer quickly finds, in order to implement similar features to AirPods in a true wireless stereo (TWS) design, the limitations in terms of memory and processing power determine what new products can and cannot do. And those limitations are inherent to the existing chips from the major alternative vendors of Bluetooth systems for TWS products, or the traditional DSP platforms that are available for audio designs. As always, the audio industry alone is not able to determine what chips can or cannot do.

Fortunately, the hyper-enthusiasm with Generative AI is coming to the rescue. Cadence Design Systems and every other major semiconductor company is racing to support Generative AI solutions and run large language models (LLM). This is creating the incentives for a new generation of chips that no audio manufacturer or even many large technology companies could develop on their own.

Cadence is working on new solutions optimized for artificial intelligence (AI)-based speech recognition processing, expanding its family of Tensilica HiFi DSP solutions for audio and voice. New generations of HiFi DSPs normally double the power for audio processing as they also implement neural network (NN) processing improvements. This is good news for automotive developers focused on voice and immersive audio experiences, and also means that many other systems will likely benefit from new chips with additional HiFi DSP cores, energy efficiency for always-on use cases, and many improvements designed to accelerate neural network inferencing in addition to traditional DSP workloads. These new-generation DSPs are optimized to run machine learning applications, and able to power wake word detection, audio scene detection, speech recognition, and noise reduction, among other audio features.

Ultimately, these solutions will evolve from infrastructure to specialized AI chips for connected applications, and finally, to edge-AI, being able of unprecedented power all on their own.

That is precisely the inspiration that has already lead to the foundation of several new chip companies, focused on delivering purpose-built silicon for edge AI deployment. But even those young pioneers are now facing the massive transition that is required to support Deep Learning AI models that run directly on-device, including speakers, headphones and hearables. Having platforms designed specifically for edge applications, running trained models optimized for audio processing from large numbers of sensors will likely not be viable unless those are based on large-scale solutions that are being created with other uses in mind—as always.

J. Martins
Editor-in-Chief



The Team

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

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

OUR NETWORK



SUPPORTING COMPANIES

Amadeus	49	Menlo Scientific, Ltd.	59
Axign	27	MISCO	39
CEVA DSP, Ltd.	55	Oxford Digital, Ltd.	53
COMSOL, Inc.	9	Parts Express International, Inc.	68
Danville Signal Processing, Inc.	13	Profusion plc	2
DSP Concepts	33	Purifi Audio	37
Ellisys Corp.	21	QuantAsylum USA LLC	45
Hammond Manufacturing, Ltd.	3	Redco Audio	49
Hypex Electronics BV	65	SB Audience	25
KAB Electro-Acoustics	27	Triad Magnetics	67
Klippel GmbH	31	Warkwyn	37

NOT A SUPPORTING COMPANY YET?

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

The Audio Voice

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



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






axCache

Every issue of audioXpress published since 2001, on USB.

Order at: **cc-webshop.com**





● Development Features

8 **Motion Feedback for Active Desktop Woofers**

By Zami Schwartzman

Explore passive and active bass enhancement methods and learn how Servobass arrived at an optimized design using motion feedback (MFB) techniques to achieve highly performing active subwoofers in desktop and never before seen compact sizes.

14 **How Optical Technology Enables a Generational Shift in MEMS Microphone Performance**

By Jakob Vennerød and Kieran Harney

This article summarizes the evolution of microphone technology, outlines the operating principle of optical MEMS microphones, and describes the benefits of the sensiBel SBM100 optical microphone series over conventional solutions.

22 **The Impact of Microphone Specifications on Active Noise Cancellation Performance**

By Nikolay Skovorodnikov

Active noise cancellation (ANC) is a popular feature for consumer electronics headsets. To achieve high-performance ANC, we need to understand how microphone selection impacts ANC performance and influences the final user experience.

28 **No Zipper, No Pop! How ADC Hybrid Gain Control Can Enable Smooth, Pop-Free Microphone Preamp Gain Changes**

By Michael Chandler-Page

Learn more about the hybrid gain control (HGC) feature of Cirrus Logic's new CS530xP family of Pro Audio ADCs.

34 **Digital Signal Processing for Over-the-Counter Hearing Aids**

By Alexander Goldin

Learn more about the principles of digital signal processing (DSP), algorithms, and software that enable hearing improvements for people with hearing loss. These can be applied to true wireless stereo (TWS) earphones that use the same hardware components as over-the-counter (OTC) hearing aids.

56 **Easing Deployment of Automotive Audio Functions**

By John Whitecar

In this article, we delve into the signal processing challenges that designers encounter when aiming to create elaborate automotive audio systems that transform driving into a harmonious experience.



● Features

60 Distortion Measurement of OPA1656 and OPA1612

By Vojtěch Janásek

Learn more about different methods to measure distortion of the extremely popular OPA1656 and OPA1612 Burr-Brown audio operational amplifiers (op-amps).

● Departments

4 From the Editor's Desk

5 Supporting Companies

● Columns

SOUND CONTROL

46 Over a Century of Home Loudspeaker Development

By Richard Honeycutt

The history of home speaker development is always fascinating to revisit if we want to fully understand how electroacoustic devices impact our listening experience and their interdependence with room acoustics. This article explores that evolution and how we quantify a loudspeaker's quality.

HOLLOW-STATE ELECTRONICS

62 Tube Filament Supplies

By Richard Honeycutt

Delve into the finer details of the reasons to use DC filament supplies, advantages over using an AC heater supply, and alternative methods to eliminate induced hum.

● 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](https://www.linkedin.com/company/audioxpress)



audioxpress

Motion Feedback for Active Desktop Subwoofers



Explore passive and active bass enhancement methods and learn how Servobass arrived at an optimized design using motion feedback (MFB) techniques to achieve highly performing active subwoofers in desktop and never before seen compact sizes.

By
Zami Schwartzman
 (Thermoguide-Servobass)

Photo 1: Servobass ZRS-1 desktop subwoofers are shown underneath Genelec 8010A active monitors.

The dynamic loudspeaker as we still know it today was invented in 1925 by Edward W. Kellogg and Chester W. Rice who were issued US Patent 1,707,570 on April 2, 1929. In the patent, they described a cylindrical coil being immersed in a magnetic field, attached to a suspended cone, forced to vibrate by alternating current flowing through the coil thus creating sound waves.

This basic construction did not change much during the last century, although materials and design techniques have improved their performance a lot since it was widely used in radio receivers. However, from the start of loudspeaker research, it was clear that something had to be done with the sound wave coming from the back of the cone to prevent interference with the front wave.

Passive Bass Enhancement Methods

In 1954, Edgar Villchur developed the acoustic suspension principle. His sealed-box concept eliminated the back wave interference. The air trapped in the box was acting as a spring and together with the cone mass and acoustic resistance performed as a classic second-order system and when correctly calculated produced acceptable bass out of small, sealed-box speakers.

The “Bass reflex” or “ported box” principle was followed by enhanced bass in even smaller enclosures using a “Helmholtz resonator” port. Air mass trapped in a short tube open at both ends resonated with the compliance of the air volume in the speaker box and 180° phase shift of the speaker’s back wave could be obtained at a narrow range. When the resonator is tuned to a low frequency below the speaker’s resonance, the emitted sound from the port is summed with the speaker’s front wave.



In 1961, A. N. Thiele published an electro-mechanical loudspeaker model analysis. Following him and other researchers, in June 1972, Richard H. Small published a series of articles on loudspeaker design and analysis theory including closed-box, vented-box, and passive-radiator loudspeaker systems. That basic theory was successfully implemented with further adjustments and is still used as a woofer/enclosure optimization tool in various forms of computer-aided design. A quarter-wave length transmission line for phase reversal of the speaker’s back wave was also widely used for bass enhancement without significant benefits in terms of speaker size or sound quality.

All these methods are based on a narrow band method for back wave phase reversal. Although current design optimization tools can achieve significant bass extension, the bass notes still come out from a narrowband source with the drawback of sustained, slow transient response leading to perceived boomy, smeared, or boxy sound.

Active Bass Enhancement Technologies

The loudspeaker industry achieved excellent results in the midrange and high-frequency ranges with ongoing improvements of materials and design concepts, however, the challenge of high-quality bass reproduction with small-size speakers remains.

Modern filter theory and digital signal processing (DSP) enabled smart equalization of the low-frequency range. Flat frequency response can be achieved with a driver in a small sealed or ported box. However, this concept does not take care of the nonlinearity that is inherent to air compression, driver’s magnetics, and



Take the Lead in Acoustics

with COMSOL Multiphysics®

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

» comsol.com/feature/acoustics-innovation

suspension. Severe distortion is apparent at low-frequency high-excursion levels—therefore only a few loudspeaker makers use bass equalization with less than acceptable distortion and transient response.

DSP-powered feed forward combined with load current sensing by smart amplifiers have been intensively investigated by Dr. Wolfgang Klippel and




Figure 1: The block diagram for the Servobass ZRS-1 MFB plate amplifier



Photo 2: A plate amplifier designed for the ZRS-1 MFB desktop subwoofers



Photo 3: A demonstration of the ZRS-1 MFB desktop subwoofers with the Genelec monitors at the IPS High End show in Munich 2023.

other research groups with encouraging results. These methods rely on the driver's linear and nonlinear parameter characterization. Manufacturing consistency and material aging issues are critical for effective commercial implementation, therefore wide acceptance by manufacturers will take time.

Spatial sound, smart speakers, soundbars, and other consumer products have become massively popular in recent years with size and sound quality being a significant differentiator between brands. High-quality bass reproduction is part of it and passive bass enhancement methods cannot fulfill these requirements.

No doubt, DSP-powered technologies will eventually generate excellent results but engineering and implementation will take time. Many patents are being issued on the subject and speaker manufacturers will have to license the technology and probably need to buy DSP modules based on custom-made chips from a few vendors that will own these patents and dominate the supply.

Motion Feedback Bass Enhancement History

Another well-known but almost forgotten bass enhancement method is motion feedback (MFB) also known as servo bass control. This method attracted lots of interest when Philips published its MFB technology article in 1968. In the early 1970s, Philips introduced its MFB-enabled RH544 bookshelf speakers with 200mm woofer in a ~10-liter air volume sealed box that was specified down to 30Hz. This was followed by the RH545 Studio monitor with a 12" woofer in ~50-liter air volume sealed box specified down to 20Hz.

The MFB concept was invented much earlier. However, these RH544 and RH545 speakers were the first commercial implementation with revolutionary deep tightly controlled bass notes reproduced from a relatively small, sealed box. Their MFB implementation was based on a custom-made piezo sensor attached to the woofer's voice coil, feeding back real-time cone acceleration to an error amplifier, thus controlling the woofer's dynamics.

Although the sound was much appreciated by consumers and sound engineers, Philips stopped making these loudspeakers because customers wanted passive speakers with their choice of amplifiers. Also, the manual assembly of discrete component circuits of the 1970s turned out to be too expensive to generate a profit.

Since that time, MFB has received lots of research attention from private inventors and leading audio companies. Capacitive, magnetic, and optic position sensors, second-winding velocity sensors, MEMS,

and piezo accelerometers were proposed as solutions for sensing the woofer’s dynamics for error correction. Many MFB-related patents were issued, but all of them have been abandoned or expired after 20 years.

A few manufacturers implemented MFB in commercial speakers, especially in active floor-standing subwoofers. Just to name a few, Velodyne Acoustics, Infinity, and Grimm Audio used acceleration sensors. Rythmik Audio used a second coil as velocity sensor. Others tried to use MFB and failed to get the loop stability, distortion, and noise performance. Somehow, MFB technology has always been considered complicated and expensive to implement. But is it really?

Increasingly, MFB is now getting attention from speaker vendors because the technology is all analog, mature, patent infringement-free, easy to understand, and easy to implement with current electronics and driver technology.

With MEMS accelerometers, Class-D power amps, switched-mode power supply (SMPS), and robust driver materials, MFB is simple, low cost, and easy to implement for deep bass reproduction by small drivers in small enclosures, thus saving cost and space. Small-size MFB subwoofers are especially attractive for production studios, home studios, bass guitar and electronic drum players, spatial sound installations, soundbars, and smart speakers.

The ZRS-1 Desktop Subwoofers Project

Small studio monitor speakers are widely used in home and professional production studios. Most of them are based on a two-way design with a 3.5” to 5” mid-bass driver in a ported enclosure and a tweeter.

The monitor’s mid-bass driver parameters are selected to follow Thiele-Small (T-S) or similar design formulas. In most of the small monitors the mid-bass driver’s maximum excursion is compromised by the need to fulfil other required parameters.

When maximum excursion level is exceeded, the driver’s nonlinearity can cause midrange modulation by the bass notes, compromising stereo image and clarity. Therefore, many small near field monitors roll off sharply and filter out the frequency range below 50Hz or 60Hz. Even with careful design, a ported monitor speaker is resonant, yielding sustained, over-hung low-frequency notes (boomy or boxy bass).

Adding a floor-standing subwoofer to complement studio monitors with the low-end frequency range causes room mode excitation and eventually although carefully placed, even high-priced studio grade subwoofers often do not blend well with studio monitors.




Figure 2: This graph shows a 6.5” driver in sealed box with 4 liters of air volume –with MFB (blue) and without MFB (red).




Figure 3: Crossover options include 100Hz, 200Hz, or wideband, which are measured with a microphone 2cm from dust cap.




Figure 4: The ZRS-1 Desktop Subwoofer’s THD, with a 30Hz tone with and without MFB

Resources

Thermoguide, Ltd. | Servobass
www.servobass.com

Thermoguide took the challenge and invested 3 years of research to create a small subwoofer that would complement small studio monitors. The goal was to create a small desktop subwoofer as a desktop add-on product.

It was clear that traditional bass enhancement methods could not cope with the acceptable size, therefore MFB was selected as a potential technology for this project.




Figure 5: The ZRS-1's third-harmonic attenuation, and a 30Hz fundamental tone with and without MFB.




Figure 6: The ZRS-1 Desktop Subwoofer's THD, with a 40Hz tone with and without MFB

About the Author

Zami Schwartzman is the founder and CEO of Thermoguide-Servobass. An electrical engineer who graduated in 1974 from the Technion in Israel, he worked for Motorola communications for 26 years, doing engineering and management of radio communication products development. He left Motorola 17 years ago and started his own engineering business—Thermoguide. Most of his business focuses on analog electronics, thermal design, and professional audio engineering. During these years he was deeply involved as a subcontractor with Waves Audio, a major software tools provider for the music industry. The Servobass desktop subwoofers described in this article were developed and engineered during COVID-19 as a complementary product that can be added beneath any nearfield studio monitor or be integrated into a three-way full-range speaker. Zami can be reached through his website: www.servobass.com.



Eventually, Thermoguide engineered and built six prototypes of the ZRS-1 MFB Desktop Subwoofer (**Photo 1**). The subwoofers measure 20cm×20cm×20cm and were intended to be placed underneath any small studio monitor to add the most desired 30Hz to 60Hz range, while blending in with the monitor's response, by having a common acoustic wave front.

An accelerometer-based MFB was chosen to demonstrate the feasibility of the ZRS-1 MFB design as a small near-field desktop subwoofer. The MDF enclosure contains the MFB control circuit, fourth-order Linkwitz-Riley crossover filters, a switching power supply, and a powerful Class-D amplifier, all designed and integrated for these subwoofers. The driver, a simple, off-the-shelf, 6" long-throw woofer was modified with an accelerometer bonded on the driver's voice coil former. Based on driver measurements and control theory, MFB loop parameters were set for flat frequency response of 20Hz to 400Hz.

A real-time voice coil temperature protection circuit was implemented to keep the voice coil temperature below 125°C during overstressed operation. And an overload warning circuit was designed to sense the amplifier's output voltage and turn a red light on when reaching close to the power supply rails.

The block diagram in **Figure 1** shows the basic features of the prototype design. The subwoofer was designed to be used in full stereo configuration, but it can also be configured as a single subwoofer feeding two monitors. In this configuration, Left and Right low frequency signals are summed for the subwoofer.

The custom-designed plate amplifier includes a 1kW peak ClassD power amplifier module, a 500W SMPS, MFB servo control loop, 24dB/octave Linkwitz-Riley crossover, real-time voice coil thermal protection, and studio-ready XLR in/out connectors. All these electronic circuits were integrated on a 20cm×20cm plate amplifier that can fit into a 20cm×20cm×20cm cube (**Photo 2**).

These plate amplifiers can be matched to any driver size in any sealed box size just by selecting a few control parameters. For that purpose, we also designed custom-made acceleration sensors that can fit any voice coil diameters from 12mm to 50mm, and to match with the plate amplifier.

With the subwoofer's frequency range from 30Hz to 300Hz, the crossover from the subwoofer to the monitor can be set to 100Hz or 200Hz, thus freeing the monitor to produce only the mid- and high-frequency range and minimize bass-midrange intermodulation. The combination practically forms a three-way desktop system with a powerful bass

extension. **Photo 3** shows a demonstration of the ZRS-1 MFB desktop subwoofers with the Genelec monitors at the IPS High End show in Munich 2023.

The ZRS-1 subwoofer (with 4 liters of internal air volume) responds flat from 30Hz to 300Hz with a powerful fast transient response. Total harmonic distortion (THD) and third harmonics are reduced by more than 10dB in the high excursion range of 30Hz to 60Hz thus the subwoofer can be driven to a higher SPL than what would be possible with an open-loop equalized speaker. By filtering out the bass range, any small monitor can be driven to a much higher SPL than specified by the manufacturer. This setup make the Desktop Subwoofers ideal for nearfield listening in home studios and even professional production studios. Gamers, bass guitar players, and other musicians, all can benefit from the powerful fast transient, effortless bass, and small size.

Another advantage of the desktop format is that direct nearfield listening to subwoofers significantly mitigates room modes, because the reflected sound is attenuated compared to the direct sound. The result is that bass notes sound the same in different environments, less sensitive to room size, acoustic treatment, or desk location.





Figure 7: The ZRS-1’s third-harmonic attenuation, with a 40Hz tone with and without MFB

THD and Third Harmonics

Other than flattening the frequency response down to the bottom audible range, being a negative feedback method, MFB significantly reduces nonlinear distortion. The ZRS-1 Desktop Subwoofer measurement plots (**Figures 2-7**) show frequency response, THD, and third-harmonic attenuation vs. SPL. The results show that with a driver in a small, sealed box, when driven by high excursion low frequencies, the MFB-enabled solution can generate up to 10dB more SPL for a given distortion level compared to the same driver in the same box driven to the same fundamental tone SPL without MFB. 

Turn your good loudspeaker into a great one.



State-of-the-art “Velvet Sound” data converters from AKM, SHARC floating-point DSP’s from Analog Devices, and integrated Audio Weaver graphical design from DSP Concepts at your fingertips.

Danville Signal has been making audio sound better for decades. We’re available to make your next system your absolute best. Let’s start a conversation.



 **Danville Signal** +1-507-263-5854 or dsp@danvillesignal.com



How Optical Technology Enables a Generational Shift in MEMS Microphone Performance

This article summarizes the evolution of microphone technology, outlines the operating principle of optical MEMS microphones, and describes the benefits of the sensiBel SBM100 optical microphone series over conventional solutions.

By
Jakob Vennerød and Kieran Harney
 (sensiBel)

The signal-to-noise ratio (SNR) and dynamic range of conventional capacitive microelectromechanical systems (MEMS) microphones have been improving incrementally over the last 20 years and are now approaching the limits of the technology. State-of-the art digital microphones are reaching an SNR of 73dBA and a dynamic range of 101dB, demonstrating the significant challenges of simultaneously achieving a low noise floor and a high clipping point in a digital microphone. Piezoelectric MEMS microphone technology promised to surpass the performance of capacitive microphones, but to date has only delivered an SNR of 64dBA with a dynamic range of 90dB.

Until now the only solution available for audio capture in premium systems requiring exceptional

electroacoustic performance has been the legacy electret condenser microphone (ECM), which is limited by the challenges of analog interface, large size, inferior part-to-part matching, poor stability, and the need for manual soldering.

The SBM100 series optical MEMS microphone has smashed the limits of SNR and dynamic range of miniature microphones. With 80dBA SNR (14dBA noise floor), 132dB dynamic range, 24-bit digital output and with low power consumption, the microphone is the long-awaited solution to electroacoustic challenges in high-performance audio and sound acquisition applications. These applications include enterprise video conferencing systems, 3D spatial audio capture, active noise cancellation (ANC) headphones, high-end smartphones, and many more.

So, how has sensiBel, a dynamic Norwegian high-tech company, achieved this significant breakthrough in MEMS microphone electroacoustic performance? Using MEMS technology, advanced complementary metal-oxide-semiconductor (CMOS) circuit design, a laser optical subsystem and miniaturization in electro-optical packaging, sensiBel has created the

Signal-to-Noise ratio (SNR)	20Hz to 20kHz	80 dBA
Equivalent input noise (EIN)	20Hz to 20kHz	14 dBA SPL
Acoustic overload point (AOP)	1kHz, 10% THD	146 dB SPL
Dynamic range		132 dB
Digital interfaces	PDM, I ² S and eight-channel TDM	

Table 1: Key product features of the SBM100 optical MEMS microphone

SBM100 product series of optical MEMS microphones that address the challenges in the market. This first-generation optical MEMS microphone breaks the performance barriers of conventional microphones and enables true studio-quality performance in a small package.

Optical MEMS Microphones

The SBM100 optical microphone is comprised of a laser optical subsystem and a digital ASIC for signal processing, packaged as a micro-assembly in a small 4.8mm×3.8mm×2.8mm top port surface mount package. The laser optical subsystem is comprised of a VCSEL, a photodetector, and a MEMS device including a diffractive optical element (DOE).

This technology is a breakthrough in microphone performance that brings great flexibility to audio system designers of modern acoustics systems in a wide range of applications. A digital microphone with an 80dBA SNR, wide dynamic range and low part-to-part variation significantly simplifies the design and cost of downstream signal processing for audio and sound acquisition, especially for multi-microphone beamforming applications. **Table 1** summarizes the SBM100’s key product features.

The SBM100 microphone has a flat frequency response up to 20kHz and a low frequency -3dB point at 10Hz. This means that the microphone has excellent phase matching at low frequencies, making it ideal for modern active noise cancelling applications that require excellent phase response and very low noise. The output is highly linear up to the acoustic overload point (AOP) of 146dB SPL. See the frequency response and linearity shown in **Figure 1**.

The microphone has a very low self-noise with an equivalent input noise (EIN) of 14dBA SPL. This equates to an SNR of 80dBA. See the SBM100’s input noise spectral density plot shown in **Figure 2**.

Miniature Microphone Technology Advancement

The ECM invented in the early 1960s represented a breakthrough in the ability of acoustic systems to acquire audio input information. The ECM enabled modern consumer devices to record audio cost effectively with a built-in microphone that was not prone to system vibration.

ECM microphones, however, are not reflow solder compatible and do not have good part-to-part sensitivity matching. With the explosion of consumer devices such as mobile phones and PCs in the 1990s, there was a need for a microphone that could be supported on high-volume surface mount technology manufacturing flows.

In the early 2000s, MEMS microphones were introduced that could withstand reflow solder temperatures and had good part-to-part sensitivity matching. MEMS microphones were available in both analog and digital output format.

Initially, digital MEMS microphones had an SNR in the low 60dBA range. Developments in both MEMS and ASIC technology have driven improvements in SNR and the current best-in-class digital MEMS




Figure 1: The SBM100’s frequency response and linearity




Figure 2: The SBM100’s equivalent input noise spectral density

microphone has an SNR of 73dBA. To achieve this level of performance, the MEMS capacitive element has evolved significantly, and the state-of-the-art microphone is comprised of dual membranes with a single backplate combined in a localized vacuum. It is increasingly difficult to further reduce the noise floor of the MEMS element and ASIC as we approach the fundamental capability of capacitive MEMS microphone technology.




Figure 3: Evolution of dynamic range in digital microphones




Figure 4: An illustration of SNR versus AOP for market leading digital microphones




Figure 5: Optical MEMS microphone development timeline

In parallel with the advancement of SNR in digital MEMS microphones, many applications demand the capability to perform at very high sound pressure levels (SPLs), for example, in loud environments or when the microphone is mounted close to a loudspeaker. The dynamic range of a microphone represents the range between the noise floor and the point where the output of the microphone is highly distorted, defined as the AOP where the total harmonic distortion (THD) is at 10%. Over the last 10 years, little progress has been made on advancing the dynamic range of digital MEMS microphones primarily due to inherent trade-offs between optimized noise floor and clipping point of the microphone. Specifically, the capacitance-to-voltage readout technology puts a heavy burden on the design of the analog-to-digital converter (ADC) and often requires a dual front-end solution to cover the entire voltage range of the output.

Figure 3 shows how dynamic range has improved slowly in digital microphones over the last two decades and how the SBM100 represents a significant breakthrough in performance.

The capacitive MEMS element sensitivity is a key contributor to the microphone noise floor. However, high MEMS sensitivity will generate high voltage output levels at higher SPLs and every 6dB increase of sound pressure will yield a further doubling of this voltage. So, at some point the ADC will be unable to manage these high voltage levels at the input. Similarly, a doubling of SPL will result in a doubling of the displacement of the diaphragm. With the gap between diaphragm and backplate of a capacitive MEMS microphone in the order of a few microns, the diaphragm and backplate will come into contact at high sound levels. The high voltage levels and mechanical constraints lead to a trade-off between SNR and AOP.

Figure 4 shows the progress in driving dynamic range by plotting SNR against AOP for various digital microphones. The plot represents the best-in-class digital microphones on the market today. The sensiBel SBM100 clearly represents a major leap forward in dynamic range by pushing the limits of SNR and AOP simultaneously.

History of Optical MEMS Microphone Technology

Work on optical MEMS microphones has been ongoing for more than 20 years as shown in **Figure 5**. In 2003, Sagberg, et al. [1] described work at SINTEF, a Norwegian independent research organization, presenting research on an optical microphone based on a modulated diffractive lens. In 2004 Lee, et al. [2] presented progress on an

optoelectronics-based micro-machined acoustic sensor summarizing that the technology should be able to achieve a noise floor in the range of 20dBA (74dBA SNR).

By 2010 various research labs reported results on proof-of concept designs to demonstrate viability of the technology. In 2014, Kim et al. [3] presented an optical micro-machined microphone that achieved 24dBA self-noise and predicted that the technology can achieve a noise floor of 15dBA. Around this time, researchers at SINTEF demonstrated a proof-of-concept prototype optical MEMS microphone solution with 14dBA noise floor (80dBA SNR), and based on this work, sensiBel was founded in 2016 to commercialize the technology. sensiBel has since significantly optimized and miniaturized the prototype and transferred production to high-volume capable fabrications. In 2022, sensiBel started delivering miniaturized SBM100 samples to customers. **Photo 1** shows the fully packaged SBM100 optical MEMS microphone including an ASIC with digital output.

Microphone Construction

The MEMS structure is a simply supported diaphragm with a reflective coating and a Diffractive Optical Element (DOE). **Figure 6** shows the basic construction and operating principle of the optical microphone system. There is no conventional backplate in the optical microphone construction and the DOE is spaced tens of microns from the diaphragm. The VCSEL produces a coherent optical beam that impinges on the DOE. A portion of light is reflected at the DOE and the remainder of the beam passes through the DOE and reflects off the bottom surface of the diaphragm. Consequently, two independent beams, A and B, are formed. The signal at the photodetector represents the optical interference of these two beams. The fundamental principle of operation is interferometric phase modulation. sensiBel has developed a photodetector-based optical signal detection system that acquires this phase-modulated interference pattern and converts it into an output that represents the displacement of the diaphragm.

This optical subsystem has been miniaturized and packaged with an advanced mixed-signal ASIC that manages the complex processing of the photodetector output. **Figure 7** shows the overall system’s block diagram.

The SBM100 series supports standard pulse density modulation (PDM), I²S and eight-channel time division multiplexed (TDM) output formats with 24-bit resolution. An I²C interface can optionally be used for configuration of advanced microphone

features. The system is packaged in a standard top port surface mount package that can be reflow soldered into the end application.

Interferometer Principles in Optical MEMS Microphones

The fundamental building block of the patented sensiBel optical MEMS microphone is the interferometer. An interferometer is basically comprised of a single coherent laser light source that is split into two different paths and then recombined, and interference of the combined light can be interpreted as a difference in path length between the two resulting beams. The interferometer is used in many diverse applications to make dimensional measurements at femtometer precision.




Photo 1: The SBM100 series technical sample. The first R&D proof-of-concept prototype has been optimized and miniaturized to 4.8mm×3.8mm×2.8mm—containing a fully packaged SBM100 optical MEMS microphone including an ASIC with digital output.




Figure 6: Basic optical microphone system construction and operating principle




Figure 7: A block diagram showing the overall system

Figure 8 shows the basic construction of a classic Michelson Interferometer. Beam A and beam B are split from a single source. When they recombine, the phase difference between the beams defines the intensity pattern that will be seen at the detector. If the path length of beam A and beam B are the same, there will be no phase difference and thus constructive interference. If the path length differs by a half wavelength, destructive interference




Figure 8: Basic Michelson interferometer

will result. Any intermediate phase difference will result in an intensity pattern that is a function of the phase difference.

The sensiBel optical microphone works on this same fundamental principle of interferometric phase modulation. A DOE is placed between the laser source and a reflective MEMS diaphragm. A proportion of the laser source is reflected by the DOE and the remaining light continues and is reflected at the MEMS diaphragm surface and returns to the detector. The phase difference between the two beams is a function of the distance between the DOE and the diaphragm. **Figure 9** shows a simplified representation of the construction of the optical microphone. A variation of SPL will cause diaphragm movement, which results in a path length difference between beam A and beam B. The light intensity pattern is detected by photodetectors and is interpreted to determine the diaphragm's position.

The light intensity at the photodetectors will change depending on whether the two light beams create constructive or destructive interference. The DOE will direct light into different directions depending on the relative phase between wave A and B, thus creating a phase-modulated interference pattern in the photodetector plane. The light patterns at the photodetectors are interpreted by a DSP in the ASIC (Figure 7) to generate the audio output of the microphone.

The photodetector current is represented by the common equation for two-beam interference from optics and wave theory:

$$I_D(x) = \alpha_i + \beta_i \cos\left(\frac{4\pi}{\lambda}x\right)$$

Where:

$I_D(x)$ = Intensity of photodetector signals

α_i = Midpoint intensity of the interferometer signal

β_i = Amplitude of intensity of the interferometer signal

λ = Optical wavelength of the source laser

x = Variable gap between the DOE and the MEMS

The intensity of signal at the photodetector is a cosine function of the gap x between the MEMS diaphragm reflective surface and the diffractive optical element. The plot in **Figure 10** shows how the output current will vary when the MEMS gap changes by one half wavelength.

The microphone output sensitivity is a function of the rate of change of the photodetector signal intensity. From the minimum point on the cosine wave to the maximum point the rate of change or output sensitivity is maximized. At the turning points the sensitivity is effectively zero.

Figure 11 represents the absolute value of the derivative of the cosine wave. Using this basic system construction, the dynamic range of the microphone is defined by the output with contiguous useful system sensitivity as shown. The digital processing of the photodetector output converts this output signal into a linear acoustic signal. Building on these fundamental concepts, sensiBel has created a sophisticated proprietary optical subsystem coupled with a digital ASIC in a miniaturized package that delivers 80dBA SNR with 132dB dynamic range.




Figure 9: The SBM100's basic operating principle

Advantages of the SBM100's Mechanical Construction

There are several aspects of the novel SBM100 optical MEMS microphone's construction that are very advantageous to microphone performance:

- There is no noise contribution from air movement through small backplate holes
- The gap between the diaphragm and the DOE is in the order of tens of microns, an order of magnitude larger than capacitive microphones.

Consequently:

- There is no significant squeeze film damping in the system. Squeeze film resistance between two surfaces is a significant source of noise and is proportional to $1/x^3$ where x represents the gap between the two surfaces.
- The diaphragm can move tens of microns, so there is no practical mechanical limit to the AOP.
- There is no risk of mechanical stiction between the diaphragm and a backplate due to material adhesion or the presence of contamination. Stiction (static friction) is a well-known failure mechanism in MEMS devices, where two surfaces come into contact and adhere.
- As the MEMS element is electrically passive, there is no voltage bias on the MEMS. In capacitive MEMS devices the diaphragm and backplate can adhere when the electrostatic attraction forces between the diaphragm and backplate exceed the restoring forces of the system. This is not a concern in the SBM100.

In contrast, the listed characteristics limit performance of a capacitive MEMS microphone. The main noise contributors for a capacitive microphone are the small gap between the diaphragm and backplate and the resistance of thousands of small holes in the backplate.

These noise contributors are shown in a basic lumped element model (Suvanto [4]) in **Figure 12a** and include back volume, backplate perforations, membrane to backplate gap, package vent and the membrane compliance. The basic construction of a capacitive MEMS microphone is shown in **Figure 12b**.

Squeeze film damping is a significant contributor of noise in a capacitive MEMS microphone because the gap between the diaphragm and perforated backplate is in the order of microns. Squeeze film resistance between two surfaces is a function of $1/x^3$ where x represents the gap between the two surfaces. In a capacitive MEMS microphone the




Figure 10: Photodetector signal output intensity vs. diaphragm position




Figure 11: Rate of change of microphone output sensitivity

gap is in the order of a couple of microns while the sensiBel microphone gap is an order of magnitude larger, reducing the backplate gap resistance significantly.

The movement of air through the perforated backplate is also a significant source of noise in capacitive MEMS microphones. In a capacitive MEMS microphone design, the backplate is perforated with thousands of small holes to reduce damping resistance between diaphragm and backplate. The total capacitance between backplate and diaphragm is a function of the plate area, so the density of hole




Figure 12: Noise sources in capacitive microphones (Suvanto [4]) (a); and basic capacitive MEMS microphone construction (b).

perforation will affect the capacitance. Consequently, there is a trade-off between maximizing the backplate hole and maximizing the system capacitance. Air movement through these holes is also a source of noise. Some new technologies apply a low-pressure/vacuum-sealed structure that reduces the damping resistance. However, such technologies have other design constraints, such as mechanical collapse due to the vacuum, and still suffer from the fundamental challenges in the ASIC-MEMS capacitive interface.

The AOP of a capacitive MEMS microphone is constrained by the small gap between diaphragm and backplate and the very large voltage output generated by high acoustic signals at the input.

The mechanical spacing between the diaphragm and backplate constrains the compliance of the diaphragm. If the compliance is too stiff the sensitivity and noise performance of the microphone will be compromised. If it is too compliant, the diaphragm will strike the backplate at relatively low SPLs. So, there is a trade-off between system SNR and AOP in a capacitive MEMS microphone which

restricts the dynamic range, typically in the order of 100dB to 110dB as shown in Figure 4.

The bias voltage of a capacitive MEMS microphone is a factor in the sensitivity of the system. The bias voltage is limited by the phenomenon of “electrostatic capture.” If the electrostatic attraction between the diaphragm and the backplate is larger than the mechanical restoring forces of the system, the diaphragm will stick to the backplate. The lower the bias voltage, the lower the system sensitivity and so the effective dynamic range of the system is constrained. In the case of the sensiBel optical microphone, the sensing element is electrically passive.

Very high signal levels create challenges for the ADC of a capacitive microphone. A very high SPL event will result in a very high voltage output from the MEMS capacitor. This can result in saturation of the ASIC analog front end, which typically is limited by the ASIC, defined by power supplies in the range of 1.2V to 3.6V depending on the application.

This results in the need for sophisticated design solutions to manage these high voltages. The capacitive MEMS diaphragm and backplate can also adhere together due to high voltage bias levels or contamination between the diaphragm and backplate.

The sensiBel optical system design is not constrained by these factors because of the design of the diffractive optical element and its distance from the diaphragm.

About the Authors

Jakob Vennerød is the co-founder and Head of Product Development at the Norwegian MEMS startup, sensiBel. Previously he worked on several research projects in audio and acoustics at SINTEF. He has extensive experience in acoustics, electronics, and signal processing, particularly in audio technology. He has a master’s degree in Electronic Engineering and Acoustics at the Norwegian University of Science and Technology, and has co-authored several patents related to microphone technology.



Kieran Harney is responsible for business development in the Americas for sensiBel. He has been in the MEMS industry for 28 years and he recently served as Managing Director of Audio Products at TDK where he played a leadership role in product development, marketing, and business development of MEMS microphones. Earlier in his career, he started a successful MEMS microphone business at Analog Devices (ADI) in 2005, that was later sold to InvenSense in 2013 and is now part of TDK.



Optical MEMS Microphones—The Future

The SBM100 product series represents the market entry of optical MEMS microphone technology and provides a generational shift in performance. It offers audio designers opportunities to improve system design and create applications with new performance levels that are otherwise difficult to implement today. Achieving 80dBA SNR (14dBA noise floor) and 146dB SPL AOP represents a dynamic range of 132dB, which is a breakthrough for digital microphones. The various digital interfaces (PDM, I²S and eight-channel TDM) offer flexibility and opportunities for simplification of the full audio signal chain.

But it doesn’t stop there. Work is already under way for the next-generation optical MEMS microphone based on the same fundamental breakthrough technology. The next generation will include further performance enhancements, reduced package size, added digital programmability, bottom port packaging, and more, to drive the optical MEMS microphones technology to very high-volume markets. 

References

- [1] Sagberg, Hakon, et al., “Optical microphone based on a modulated diffractive lens,” *IEEE Photonics technology letters*, 15.10 (2003), pp.1431-1433.
- [2] Wook, Lee, et al., “Fabrication and characterization of micromachined acoustic sensor with integrated optical readout.” *IEEE Journal of Selected Topics in Quantum Electronics*, 10.3 (2004), pp. 643-651.
- [3] Kim, Donghwan, and Neal A. Hall. “Towards a sub 15-dBA optical micromachined microphone,” *The Journal of the Acoustical Society of America*, 135.5 (2014), pp. 2664-2673
- [4] M. Suvanto, *The MEMS Microphone Book*, Mosomic, 2021.



Audio Over Bluetooth® LE Is Here. Ready?

Ellisys offers an array of protocol analysis solutions that support Bluetooth® Low Energy Audio, including Auracast™.



Bluetooth Vanguard™

Advanced Wireless
Protocol Analysis System

BT + BLE + Wi-Fi 3x + WPAN

www.ellisys.com/bv1



Bluetooth Explorer™

Dual-Mode Bluetooth®
Protocol Analysis System

BT + BLE + Wi-Fi 2x

www.ellisys.com/bex400



Bluetooth Tracker™

Ultra-Portable Bluetooth® Low
Energy Protocol Analyzer

BLE + Wi-Fi 1x

www.ellisys.com/btr1

The Impact of Microphone Specifications on Active Noise Cancellation Performance

By
Nikolay Skovorodnikov
(Knowles)

Active noise cancellation (ANC) is a popular feature for consumer electronics headsets. Designing effective ANC imposes particular specification requirements on audio chain components, including microphones. To achieve high-performance ANC, we need to understand how microphone selection impacts ANC performance and influences the final user experience.

The high-level idea of an active noise-cancelling (ANC) system is quite simple. The goal is to eliminate unwanted sounds such as airplane noise, highway noise, or people's chatter. Several techniques achieve this. First, it is essential to recognize passive cancellation. A common ANC device form factor that exemplifies this are the true wireless stereo (TWS) earbuds. The dominant design solution is a rubber ear tip that fits the ear canal tightly. When inserted into the ear, the device acts somewhat like an earplug. It provides low-pass acoustic filtering and significantly attenuates high-pitch sounds without any electronics involved.

The feed-forward (FF) ANC path is a second noise cancellation technique. In this technique, an external microphone picks up the unwanted environmental sound. The signal is inverted to create anti-noise and is played by the speaker. The ambition is to create perfect cancellation inside the ear. The FF signal is an exact inverted copy of the sound naturally coupling into the ear from the outside. FF ANC is usually effective in the low- to mid-band range.

The feedback (FB) ANC path is the third technique: The microphone inside the ear monitors the sound. The microphone's output is inverted and played back through the speaker. This way, the noise level inside the ear is reduced. FB ANC cancellation is effective in the low-frequency range. A typical example is shown in **Figure 1**.

Feed-Forward ANC

A concept block diagram of the FF system is given in **Figure 2**. The passive transfer function is a low-pass filter, subject to the mechanical design of the headset and the way it fits in the ear. Active transfer function includes a microphone, digital signal processing (DSP), and the transfer function of the speaker generating sound in the ear canal.

The shape of the active transfer function consists of these expected features: second-order high-pass effect in the low end from microphone and speaker low-frequency roll-off (LFRO). There is a pinna resonance from the speaker transfer function at higher frequencies.

Note that the phase response of the active path includes linear phase contribution from digital microphone ASIC latency and DSP processing latency. The idea is to pick the ANC filters to achieve complete cancellation so that active and passive signals are equal in magnitude and perfectly out of phase at every frequency. ANC filters are part of the DSP block, but it is convenient to separate those from the fixed components of the active transfer function.

The next step is to set the ANC filters in the active signal path. How accurately must the active and passive paths be matched to make ANC noticeable to the consumer? Both magnitude and phase differences will lead to less cancellation. **Figure 3** illustrates this.

Let's pick -10dB cancellation as a clearly audible target. Cancellation is perfect in the ideal case with a 0dB magnitude difference and 180° phase shift. However, a -10dB or better noise reduction target is achieved with a magnitude mismatch of up to 2dB and phase deviation from the ideal of 180° up to 20°. This exercise shows how accurately the ANC filters need to be set.

Having measured the passive and active transfer functions without ANC filters (unity filters), it is now possible to compute the target ANC filter curves for magnitude and phase. The curves are shown in blue in **Figure 4**. Unfortunately, this exact shape cannot be reproduced with available DSP filters in the ANC chip. Commonly there is a cascade of configurable second-order IIR bi-quad filters available. Also, there is a fundamental difference: The target phase response is rising in the frequency range of interest. A perfect fit in the entire frequency range is impossible without violating causality.

The best possible fit of DSP filters to the target filter curve is shown in red. Note that phase match is breaking at 40Hz to 50Hz, and both magnitude and phase deviate from the target in the frequency region of the pinna resonance (2kHz to 3kHz). These exact effects limit the frequency range of effective FF ANC.

It is also important to note that FF ANC may not only reduce the sound in the ear but also increase it. Too high of ANC filter gain at higher frequencies leads to the so-called waterbed effect, and the final FF ANC effect is above 0dB. This is not desired, and the solution is to reduce the ANC filter gain at those frequencies. However, this will also make cancellation worse at the frequencies where it is effective. ANC designers must take this trade-off into account. Luckily, passive attenuation is suitable in high frequencies and helps to mitigate the waterbed effect.

The previous paragraphs explained the FF ANC mechanism and the sources of its limitations. Now it is possible to analyze the effect of microphone characteristics on FF ANC.

An FF microphone is a part of the active transfer function. The shape of its magnitude and phase response adds to other components, like the speaker. The simulation in **Figure 5** shows a comparison of two FF ANC systems. The baseline solution has 30Hz LFRO microphone and 24μs microphone ASIC digital latency. The modified microphone has 10Hz LFRO and 3μs latency and assumes 30μs lower ANC processing latency. One can observe the change in the shape of the active transfer function on the chart below. The target filter curve changes and the fitted FF filter transfer function is updated accordingly.

As expected, the microphone's broader and flatter magnitude and phase response flattens the active transfer function, making target filter curves flatter and easier to fit with DSP filters in a wider frequency range. A better fit of ANC filters leads to better ANC cancellation in a wider frequency range and pushes the waterbed region to higher frequencies. Specifically, lower microphone LFRO




Figure 1: A typical example of the feedback (FB) ANC path




Figure 2: Feed-forward (FF) active noise cancellation path




Figure 3: Both magnitude and phase differences will lead to less cancellation.




Figure 4: It is now possible to compute the target ANC filter curves for magnitude and phase. The curves are shown in blue.




Figure 5: The simulation shows a comparison of two FF ANC systems.




Figure 6: This is a concept diagram for the FB ANC system.




Figure 7: A FB ANC system's open loop transfer function is shown here.

improves the low-frequency limit of FF ANC, while lower system latency improves the high-frequency limit of FF ANC.

Should it be concluded that FF mic LFRO for TWS ANC headsets should be as low as possible? It will improve ANC, but other system trade-offs are essential to consider (e.g., wind noise mitigation). If the microphone LFRO is too low, the system will be too sensitive to wind noise, which is rich in low frequencies. High-frequency FF cancellation improves with lower system latency. The total latency value includes microphone and DSP processing. All components of the system need to be improved for the best results.

Feedback ANC

A concept diagram for the FB ANC system is shown in **Figure 6**. Unlike FF ANC, the entire system functions inside the user's ear. Input signal into the FB loop is the residual noise after passive and FF cancellation. The FB microphone captures sound, and the DSP path applies FB ANC filters. Then, the resulting signal is played through the speaker. The open loop transfer function includes all system components, assuming no coupling from the output back to the input. Usually, it is measured in the analog domain by feeding a constant signal to the speaker and measuring the output of the DAC of the ANC chip.

To maximize noise cancellation in an FB system, the open loop transfer function must apply as much gain as possible while maintaining the system's stability. There are mathematical techniques to describe this trade-off. For example, phase and gain margin analysis can be used. **Figure 7** shows a FB ANC system's open loop transfer function. With filters set to unity, one can observe the starting point of algorithm tuning. Next, the filters are picked to maximize FB cancellation while the waterbed effect and stability are still acceptable. In this example, filters are constructed from 20dB gain, 180° phase inversion, and second-order bandpass filter from 1Hz to 425Hz.

The result of the tuned open loop transfer function is given in **Figure 8** as the FB ANC curve. Maximum achievable attenuation is more than -15dB at around 200Hz. The open loop transfer function at that frequency has maximum gain and phase inversion close to 180° in the low-frequency range, and at 1kHz to 3kHz, the FB system amplifies the signal. This is an unwanted waterbed effect. Phase and gain margins are reduced at those frequencies, but the system is still stable. The above 5kHz open loop transfer function is sufficiently attenuated, and the FB loop does not impact the signal.

The next step is to evaluate the impact of microphone specification on the operation of the FB ANC loop. As in the FF process, the microphone transfer function is one of the components of the open loop transfer function. The lower microphone LFRO makes it flatter at low frequencies and thus improves the gain margin. Lower system latency leads to less steep phase response and improves phase margins in the high-frequency range. **Figure 9** shows the impact of LFRO and latency improvement on FB ANC. Effective bandwidth is enhanced on both sides, reducing the waterbed effect.

In summary, the microphone requirements for the best FB ANC have the flattest possible magnitude and phase frequency response. Lower LFRO improves FB ANC in the low-frequency range, and lower microphone latency improves FB ANC in the high-frequency range. It is important to note that lowering the microphone LFRO too much can have negative consequences. Head or jaw motion and other forms of human physical activity can lead to high SPL, low-frequency signals generated inside the ear canal while wearing a TWS headset. Having LFRO sufficiently high protects the system from unwanted overload.




Figure 8: The result of the tuned open loop transfer function is given as the FB ANC curve.

Figure 9: This graph shows the impact of LFRO and latency improvement on FB ANC.

ROSSO-34CDN-PK
PEEK diaphragm compression driver

Scan for more information:

SCREW ON / BOLT ON

Find us on:

- SB Audience
- sbaudience
- sbaudience_
- SB Audience
- sbaudience.com




Figure 10: The block diagram for the transparency mode is similar to that of the FF system.




Figure 11: Here is the target open fit transfer function and active and passive signal path transfer function discussed in the FF section.





Figure 12: Like FF ANC, target curves are fitted with available filters in the DSP chip.

Figure 13: The microphone input-referred third-octave noise spectrum is compared to the same characteristic of the transparency mode system.



Transparency Mode

Earlier we stated that a TWS headset provides passive attenuation in the ear. Simply put, it acts like an ear plug attenuating high-frequency signals. Hearing the outside world without such passive attenuation is important in some use cases. This is achieved in what is called transparency mode. The block diagram for this (**Figure 10**) is similar to that of the FF system. The difference is the target signal inside the ear. Unlike FF ANC, where the goal was to cancel out the sound that couples into the ear, the purpose of transparency is to match the transfer function of the open ear.

Figure 11 shows the target open fit transfer function and active and passive signal path transfer function discussed in the FF section. As expected, the curve is generally flat with high-frequency resonances in the ear. The active signal path multiplied by the transparency mode filters combined with the passive signal path must match the open target curve. At this point, it is already clear that transparency mode will require lots of gain applied to the microphone signal.

Target transparency mode filter coefficients can be derived by solving a simple equation described as:

$$\text{TargetFilter} = \frac{\text{OpenFit-PassivePath}}{\text{ActivePath}}$$

Like FF ANC, target curves are fitted with available filters in the DSP chip. An example is given in **Figure 12**. Note that the filter magnitude is above 0dB, and lots of gain is applied to the microphone. Matching perceived loudness will be the most noticeable effect in transparency mode. Of course, the shape of the phase response curve can be audible due to temporal distortions. Given the flexibility of DSP filters available, the match of magnitude filter curve is considered more important in this case for the audible experience of the user in transparency mode.

As a result of this algorithm, the user of the TWS headset will hear sounds that closely resemble the experience with an open ear. However, it is critical to recognize that the microphone itself is a noise source. Microphone noise will be amplified by the DSP coefficients and played through the speaker. This means the noise floor in the ear will increase. If the environment around the user is a sufficiently loud microphone, the noise will be negligible, and the system will operate as intended. However, the microphone will dominate the in-ear noise floor in a quiet environment. Sometimes this effect is described as transparency mode hiss.

Figure 13 compares the microphone input-referred third-octave noise spectrum and the same

Resources

Knowles V2S Demo, <https://vimeo.com/811043420>

characteristic of the transparency mode system. The Integrated A-weighted noise difference is 13dB. In this example, the user might start hearing the transparency mode hiss if the environment is below 39dB SPL.

The level of transparency mode hiss drives the requirement to use microphones with as high as possible SNR. Higher SNR and, thus, lower microphone noise expands the range of use cases for transparency mode where the user will not experience the hiss.

Conclusion

We described the operation of FF ANC, FB ANC, and Transparency mode, along with the critical microphone specifications that lead to improved user experience. This exercise shows that microphone latency needs to be reduced as much as possible, and that a proper algorithm tuning allows better FF and FB ANC in the high-frequency range. It also reduces unwanted high-frequency amplification, known as the waterbed effect.

We have also seen how microphone LFRO needs to be set optimally. With proper algorithm tuning, lower LFRO allows for better cancellation in low frequencies for both FF and FB ANC. However, LFRO

cannot be set too low to avoid wind noise issues with the FF microphone and infrasonic overload issues with the FB microphone.

Finally, we have shown that microphone SNR must be as high as possible to reduce audible hiss and expand use cases for transparency mode.

Knowles has invested significant time in understanding how microphone selection influences the final user experience. From these findings, Knowles has developed a comprehensive portfolio of MEMS microphone models that can help manufacturers achieve the best noise cancellation possible. With these specifications in mind, these MEMS microphone models can ensure optimal performance in TWS ANC applications. 

About the Author

Nikolay Skovorodnikov is a Senior Manager Applications Engineering, MEMS Microphones at Knowles, based in Itasca, IL. Previously, Nikolay was a Graduate Researcher at University of California, Davis, Department of Chemical Engineering and Materials Science, and held research positions at Nano-Optic Devices. Nikolay received a Bachelor of Science degree from Lomonosov Moscow State University (MSU).



We are Axign. A semiconductor company on a mission to conquer the consumer audio amplifier market with our unique post filter feedback digital class-D audio amplifier solutions

THE RESULTS

- ✕ Power consumption reduction at all output power levels
- ✕ Excellent audio quality due to post filter feedback and digital control loop
- ✕ Cost competitive audio amplifier systems




axign
THE SOUND OF INNOVATION

FIND US ON **AXIGN.NL**

SpiralAirtm
Interconnects

*Coax
Is Cool
Again!*



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

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

KAB
Preserving The Sounds
Of A Lifetime

www.kabusa.com

No Zipper, No Pop!

How ADC Hybrid Gain Control Can Enable Smooth, Pop-Free Microphone Preamp Gain Changes

By
Michael Chandler-Page
(Cirrus Logic)



Learn more about the hybrid gain control (HGC) feature of Cirrus Logic's new CS530xP family of Pro Audio ADCs.

Pop and zipper noise when changing gain on digitally controlled mic preamps is now simple to eliminate, using the hybrid gain control (HGC) feature of Cirrus Logic's new CS530xP family of Pro Audio ADCs. This pop and zipper noise is a long-standing challenge for designers of pro-audio gear. Existing solutions for fine-resolution mic preamp gain control, with the feel of a traditional potentiometer but without audible pops or zipper noise, are often expensive and complex. HGC solves this problem by integrating control of the mic preamp gain into the ADC, together with zero-cross detection and precisely synchronized digital gain control. HGC makes it easy for pro-audio product designers to assure smooth pop-free gain control, while retaining their existing mic preamp design and simplifying both hardware and software.

Digital Gain Control of Mic Preamps

Historically, the mic preamp gain on analog pro-audio products such as mixing consoles was controlled by a potentiometer. The advent of digital audio mixing consoles created a need for mic preamp gain to be digitally controlled. Digital control enabled mic preamps to be located remotely from the control surface of the mixer. In turn, this enabled valuable new product architectures that replaced the heavy analog multi-core cables strung around studios and performance

venues with thin, light network cables carrying digital audio. Digital gain control also allowed mic preamp gain settings to be stored and recalled under software control. It is now a ubiquitous feature of mixing consoles and is increasingly featured in other pro audio products such as recording interfaces for laptops.

This digital gain control is most often implemented using analog multiplexers to switch different gain-setting resistors into the mic preamp circuit. The user interface on the control surface is typically a rotary encoder. Sound engineers still expect this "gain knob" to respond with the continuous gain response of a potentiometer. To achieve this, control software manages a digital "interpolation gain" applied after the ADC to interpolate between the analog gain steps. As the sound engineer turns the gain knob clockwise, the switched mic preamp gain increases in steps, while the digital gain interpolates smoothly between those steps (**Figure 1**).

When the analog gain increases by a step, the digital gain must simultaneously decrease by the same amount to maintain a smooth gain increase. The analog and digital gains sum to produce the system gain, which should be the same as the gain commanded by the knob.

To minimize pop noise on changing gain, the analog gain changes may also be aligned with the analog signal crossing the signal zero-level point. This ensures that the gain change occurs when the signal amplitude is zero, to

avoid the gain change causing audible transients. A typical system implementing this scheme, without HGC, is illustrated in **Figure 2**.

Implementing Digital Control of Mic Preamp Gain Is Tricky

Pro audio product designers over the years have encountered various problems and complexities with implementing the kind of scheme for digital control of preamp gain depicted in Figure 2.

Perhaps the biggest difficulty is that the digital interpolation gain is usually implemented in the audio system's DSP. This is often in a different part of the system to the mic preamp, and system constraints can make it difficult to achieve precise synchronization between analog gain changes and digital interpolation gain steps. Timing misalignment between these operations may cause audible transients, as illustrated below by misconfiguring Cirrus Logic's HGC demonstration hardware. These figures show a gain change comprising an analog +3dB increase and a -3dB digital gain step on a 1kHz sine wave. The analog gain change is simultaneous with the zero-crossing point, but the matching digital gain step is 4 samples (about 83μs) late. This causes a glitch visible in the waveform (**Figure 3** and **Figure 4**) and is clearly audible.

In isolation, these transients sound like "pop" noises, caused by the user turning the mic preamp gain control knob. If the knob is turned rapidly, the rapid stream of transients can sound like a zipper. In both cases, these noises are so unlike the behavior of a potentiometer, which the system is supposed to emulate, that it is distracting and unpleasant for sound engineers.

Another difficulty is that the system illustrated in Figure 2 is complex. It includes discrete circuitry for analog zero-cross detection to align analog gain changes with zero-crossing to minimize transients. If omitted, the pop and zipper noise is much worse, such that the audio channel may need to be temporarily muted when changing gain. The system also needs multiple control paths from the host. If the DSP and the analog input path are under the control of two separate communication processors, distributed host software must manage the state of the system gain split into two components, each located in different computing nodes. This can introduce substantial software complexity.

One potential solution is to use a Multiplying DAC (MDAC) instead of the switched-resistor gain selection. This gives high-resolution, pop-free gain changes, but MDAC devices may be too expensive for many pro audio products. Alternatively, more multiplexer paths with more resistor combinations




Figure 1: Interaction of analog gain steps and digital interpolation gain




Figure 2: A typical legacy pro audio input signal path with digitally controlled mic preamp gain





Figure 3: Real waveform view of glitch due to digital gain step misalignment

Figure 4: Zoomed-in waveform view of a digital gain step misalignment glitch, illustrating a 4-sample timing error



can give smaller step sizes to approximate continuous gain control, but that also increases cost and complexity, and may degrade audio performance.

In short, implementations of digitally controlled mic preamp gain have tended to be either susceptible to pop/zipper noise on gain changes or have incurred high cost and complexity to mitigate the pop/zipper noise.

Cirrus Logic CS530xP Makes Pop-Free Gain Control Easy

Cirrus Logic's new CS530xP family of Pro Audio ADCs introduces HGC, which greatly simplifies pop-free digital control of mic preamp gain. HGC achieves this by integrating into the ADC chip several elements of the legacy system shown in Figure 2.

A configurable daisy-chain SPI controller configures the gain of up to eight mic preamps. Each

ADC channel features an analog zero-cross detector to synchronize pop-free analog gain changes. Each ADC signal path includes a digital gain control to implement the interpolation gain emulating the continuous control feel of a potentiometer. HGC logic coordinates operation of all the above functions to execute precisely synchronized, glitch-free gain updates. The resulting simplified system architecture with HGC is illustrated in **Figure 5**.

When the user turns the gain control knob on the control surface, the host software writes registers in the CS530xP to configure and request a mic preamp gain change. This causes HGC to control the mic preamp, via SPI, to set the desired gain change. This is precisely aligned with an analog zero-crossing. Since HGC also has control over the digital audio signal path, it precisely aligns the digital gain change with the analog gain change. This precise alignment of analog and digital gain changes with a zero-crossing ensures a pop-free gain update.

The states of the analog gain and digital interpolation gain are both located in the registers of the CS530xP HGC block, under the control of the one host processor controlling the ADC.

Compared to the legacy system illustrated in Figure 2, the HGC system has simpler hardware for a smaller bill-of-materials, simpler software and creates less pop/zipper noise on gain changes.

How Hybrid Gain Control Works

The sequence of events for a gain change with HGC is illustrated in **Figure 6**. A gain change is initiated by the host software writing to the HGC registers to configure and trigger the gain update. This causes the HGC SPI interface to send the gain bit patterns to the daisy-chain of mic preamps, to configure one of them for the newly commanded gain value (point 1 in Figure 6). The gain bit pattern depends on the design of the mic preamp.

After the HGC SPI interface has shifted the gain bit patterns into the mic preamps, the SPI CSb signal remains asserted low until a zero-cross is detected on the channel that is changing gain (point 2 in Figure 6). This causes CSb to be de-asserted, which causes the new gain code to take effect in the mic pre-amp, which in turn causes an instantaneous zero-cross-aligned gain change in the analog signal path (point 3 in Figure 6).

After the gain change, HGC waits a short configurable duration (typically 4-8 sample periods, point 4 in Figure 6) to compensate for the signal path latency from the mic preamp, through the analog anti-aliasing filter, ADC, and decimation filter.

At this point, HGC executes the precisely timed digital gain step (point 5 in Figure 6), to counteract




Figure 5: Simplified mic preamp digital gain control with HGC

the analog gain change. It then commences a gain ramp to the final fine-gain value (point 6 in Figure 6). Ramping the gain over tens of milliseconds avoids introducing audible transients.

Achieve Pop-Free Operation with Existing Mic Preamp Designs

For some pro-audio product designers, the design and sonic characteristics of their mic preamp is a crucial brand asset and a differentiating feature. HGC allows product designers to achieve pop-free operation with their existing mic preamp designs. Any mic preamp design may be used if it complies with just three simple requirements:

- Has stepped gain control. Gain steps do not need to be equal sized, and it may have any number of gain steps of any values.
- Gain must be controllable by daisy-chain SPI, such that the gain can be changed by writing one SPI bit pattern into its daisy-chain SPI shift register, of any length between 1 and 32 bits per mic preamp. The encoding of SPI bit pattern to gain value is completely configurable in HGC.

- Gain change must take effect immediately on de-asserting SPI CSb at the end of the daisy-chain SPI operation.

A simple and compact implementation is to use daisy-chainable SPI-controlled analog multiplexers




Figure 6: HGC analog gain change sequence

TBM

Tone Burst Measurement

Maximum Short-term SPL of Subwoofers according to IEC and ANSI/CEA Standards

- ▶ Stepped sine bursts (amplitude, frequency)
- ▶ Harmonic distortion threshold
- ▶ Reveals mechanical speaker nonlinearities
- ▶ Compatible with KLIPPEL module
- ▶ In-situ Room Compensation (ISC)

BURSTING WITH BASS!

www.klippel.de | info@klippel.de

to switch different value resistors into the mic preamp circuit. Alternatively, a standard logic shift register such as the 8-bit 74HCS595 may be used with daisy-chain SPI to give static logic outputs in a manner similar to a port expander, which in turn may control analog multiplexers or directly control discrete transistors.

HGC supports mic preamp gain control for each channel of the ADC (up to eight channels, in the case of the Cirrus Logic CS5308P). It also includes control of up to 64 bits of additional SPI daisy chain devices, through the same SPI interface the mic preamp gains. These can be used for control of auxiliary hardware functions associated with the input paths, without requiring extra microcontroller I/O or serial interfaces. Applications include switching +48V phantom power, enabling an analog high-pass filter, controlling LEDs by the input connectors, or switching between balanced and unbalanced inputs.

Easy to Control from Existing Software Architectures

Control of gain changes with HGC is similar to existing schemes, in which both analog gain and digital interpolation gain are computed and set by the host software. This means that the host software modifications required to use HGC are typically minor and of limited scope. To command a HGC gain change, the host software writes four pieces of information to the HGC registers of the CS530xP ADC chip:

- **Gain selection bit pattern:** The pattern of SPI bits to set the mic preamp to the selected analog gain setting, according to the design of the preamp.
- **Analog gain:** The absolute analog gain value corresponding to the new bit pattern for the selected mic preamp, to the nearest 0.125dB.
- **Digital gain:** the value of the digital interpolation gain required to achieve the system gain set by the user, to the nearest 0.125dB.

- **Gain update signal:** Indicates that the three items above have all been written, such that HGC may now execute the synchronized gain change.

As with prior digital gain control implementations, the algorithm to choose which analog gain step to select for a particular system gain request is implemented in host software, according to the system designer's preference for managing dynamic range.

The analog gain register fields are used to determine the difference in mic preamp analog gain between the previous setting and the new setting. This difference is used to determine the size and direction of the digital gain step, to counter the analog gain step before the digital gain is ramped to the new value specified in the digital gain field.

The host may command a gain change at any time, and it is guaranteed to be executed within a time bound of approximately the number of active channels multiplied by the zero-cross timeout configuration setting (20ms maximum). In practice, with real-world audio signals it is likely to execute much faster than that due to audio zero-cross detection.

HGC Solves a Long-Standing Design Challenge

HGC on Cirrus Logic's CS530xP family of ADCs, and other forthcoming pro audio products, makes it easy and inexpensive to implement digitally controlled mic preamps with smooth, quiet gain changes without pop or zipper noise. This resolves a long-standing design challenge with implementing digital control of mic preamp gain, which until now has presented pro audio product designers with a difficult trade-off between frustrating pop and zipper noise on gain changes, or high complexity and larger bill-of-materials (BOMs).

HGC makes it easy to achieve these benefits with existing switched-gain mic preamp designs and is easy to control from existing host software architectures. It simplifies hardware design for mic preamp gain control, including integrated control for auxiliary functions such as phantom power, signal path relays, or LEDs. This new capability is conveniently integrated with state-of-the-art ADC performance of 123dB dynamic range and at least -110dB THD+N, at a power consumption of just 25mW/channel (Cirrus Logic CS5308P).

HGC enables the premium user experience of potentiometer-like digital mic preamp gain controls, with no pop or zipper noise, to be delivered across all pro audio product tiers at an affordable cost. 

About the Author

Michael Chandler-Page is a systems engineer at Cirrus Logic where he leads engineering product definition for the company's latest Pro Audio family of converters. His career in the semiconductor industry includes a combined 12 years at Cirrus Logic and Wolfson Microelectronics (which was acquired by Cirrus Logic in 2014). Prior to joining Wolfson, Michael worked at Soundcraft, Sony Pro Audio Lab and Peavey Electronics, where he developed professional and commercial audio products, as well as conceiving and developing the AES50 pro audio networking technology. Michael earned a Master of Engineering degree in electrical and electronic engineering at Loughborough University, UK.





Get your product to market faster with **Audio Weaver**, the audio development platform that makes audio innovation easy.



- Easy-to-use graphical design tools
- Tools for prototyping, debugging, and tuning
- Open IP for playback, conferencing, and far-field voice
- Expert design, tuning, and system integration services
- Pre-integrated with multi-core SoCs for easy development

Pre-integrated SDK

Qualcomm

SAMSUNG

AIROHA



Advanced 3rd party IP offered through **Audio Weaver**

Fraunhofer



intelliGo
Technology Inc.



Pre-integrated Reference Products

TYMPHANY




AmTRAN

SINTEC

New Offerings

- Integrated Machine Learning support in Audio Weaver
- Comprehensive use-case support for all Bluetooth products
- Conferencing Toolkit enables easy assembly of complete solutions leveraging customer internal IP, DSP Concepts IP, and 3rd party IP
- Multi-instance features allows processing of audio across multiple independent cores in an SOC via a single layout
- Multi-threading support in Audio Weaver allows you to specify exactly which core or thread each processing block should run on



Digital Signal Processing for Over-the-Counter Hearing Aids

This article looks at the principles of digital signal processing (DSP), algorithms, and software that enable hearing improvements for people with hearing loss. These can be applied to true wireless stereo (TWS) earphones that use the same hardware components as over-the-counter (OTC) hearing aids. What are these algorithms and how can they improve hearing?

By
Alexander Goldin
(Alango Technologies)

A recent World Report on Hearing by the World Health Organization (WHO) shows that 1.5 billion people worldwide have some degree of hearing loss, and 430 million people need rehabilitation services. Shockingly, less than 3% (on average worldwide) use hearing aids. Can anything be done to change this? For example, can true wireless stereo (TWS) earphones that use the same hardware components serve as over-the-counter (OTC) hearing aids? The answer is yes, if their software is upgraded with algorithms specifically tailored to the needs of the hearing impaired.

Hearing Loss, Hearing Thresholds, and Hearing Tests

In audiology, hearing loss is quantified by measuring hearing thresholds for each ear at different frequencies. During a hearing test, the person being tested typically sits in a soundproof booth wearing headphones. The audiologist then plays a series of tones of varying frequencies and loudness levels, one ear at a time, and the person indicates when they can hear each tone. The test continues at different frequencies until there is a complete picture of the person's hearing thresholds for each ear.

These thresholds are expressed relative to those of an average individual with healthy hearing, usually around 0 to 10dB of sound pressure level (SPL) and then plotted on an audiogram, which is a graph that shows the hearing thresholds relative to those of an average individual with healthy hearing. Audiograms, which illustrate these relative thresholds, are used to summarize the collective data. A sample audiogram is provided in **Figure 1**.

The audiogram indicates that the tested person has moderate hearing loss in the frequency region around 3000Hz, with a minimal loudness level 40dB higher in the left ear and 50dB higher in the right ear compared to a person with healthy hearing. However, simply amplifying sounds around 3000Hz by 40dB and 50dB, respectively, will not restore the person's hearing to normal. Hearing loss is not




Figure 1: Hearing test and audiogram example

a linear reduction in ear sensitivity that can be compensated for by a simple or even frequency-dependent linear sound amplification. Instead, hearing loss is both frequency and sound level dependent, meaning that the degree and type of hearing loss can vary significantly depending on the specific frequencies and sound levels involved. **Figure 2** illustrates this concept, although the numbers are given for illustration purposes only.

One of the main characteristics of hearing loss is “loudness recruitment,” where the perceived loudness of sounds grows much faster for people with hearing loss than for those with normal hearing. As a result, the difference in sound level between “barely heard” and “normal” becomes very short for people with hearing loss, which can explain why they often have difficulty hearing others speaking to them in a normal voice. This phenomenon is illustrated in **Figure 3**, where the loudness for people with normal hearing grows steadily, while it remains at zero until the personal hearing threshold is reached for people with hearing loss starting to grow abnormally fast afterward till the point where it becomes equal to people with normal hearing.

This phenomenon has led to the development of multichannel wide dynamic range compression (WDRC) for hearing loss compensation in hearing aids. However, hearing loss is a complex, nonlinear, frequency-dependent, and cognition-related phenomenon that cannot be universally compensated by amplifying sounds. Nonetheless, proper hearing amplification can still benefit most people with hearing loss to some extent. Our goal is to make such amplification as efficient as possible. It’s important to note that all thresholds and corresponding plots shown in Figure 3 are individual and frequency-dependent.

Hearing Aid Categories

Hearing aids are the most commonly used hearing enhancement devices. Air conduction hearing aids amplify ambient sounds based on the user’s specific hearing loss and/or preferences, and functionally consist of a microphone, an amplifier with controlled amplification, and a receiver (speaker) as shown in **Figure 4**.

While microphones and receivers are typically high-quality, linear components, the amplification is generally nonlinear, with properties and characteristics that define the features and quality of hearing aids. Hearing aids can be divided into two main categories based on the type of amplification control:

1. Prescription hearing aids are Class 1 medical devices distributed through professional hearing




Figure 2: Sound levels and their perceptions




Figure 3: Loudness and hearing recruitment





Figure 4: Hearing aid structure

healthcare channels and may be designed to cover all types of hearing loss from mild to profound. The amplification parameters of prescription hearing aids must be changed through professional fitting services and cannot be adjusted by users themselves.

Figure 5: Standard DSP tasks in hearing aids



2. Over-the-Counter (OTC) hearing aids are a new category recently defined by the US Food and Drug Administration (FDA) to provide a solution for people with mild to moderate hearing loss. OTC hearing aids offer similar functionality to prescription hearing aids for people with mild to moderate hearing loss but can be sold over the counter and administered by users without the help of hearing care professionals. OTC hearing aids are designed to enable users to directly control the amplification characteristics or perform a hearing test and set application parameters accordingly, to provide an optimal user experience.

There are no constraints on OTC hearing aid form factors, creating opportunities for innovation, increased competition, and reduced prices. There is a big hope today that OTC hearing aids will lead to much wider use of hearing amplification devices by people who, for whatever reason, do not benefit from prescription hearing aids.

Basic DSP in Modern Hearing Aids

All modern hearing aids are digital, i.e., analog electrical microphone signals are first digitized, then processed in the digital domain, and finally converted to an analog signal and reproduced through a receiver (loudspeaker). The typical DSP blocks of a high-end hearing aid are shown in **Figure 5**.

Most of the DSP in hearing aids is done in frequency sub-bands. The full-band input signals are divided into frequency sub-band signals, as shown in **Figure 6**, with each sub-band conveying the information for the corresponding narrow frequency range.

DSP is sequential, block by block. After all the processing is done, the full-band signal is resynthesized (combined) from the processed sub-band signals.

In some cases, the processing blocks can receive sub-band signals from multiple sources and combine them into a set of sub-band signals. Acoustic beamforming is one type of this processing, shown in **Figure 7**.

Splitting into and combining frequency sub-bands can be accomplished using various and well-known methods such as the short-time Fourier transform (STFT) or filter banks. Finite impulse response (FIR) or infinite impulse response (IIR) filter banks can be used, each having its own advantages and disadvantages. For example, IIR filter banks provide the shortest delays but are less suitable for implementing other algorithms such as adaptive beamforming, noise reduction,




Figure 6: Sub-band processing scheme




Figure 7: Acoustic beamforming sub-band merging

and others. When STFT is used for decomposition, the sub-bands contain complex numbers and are said to be processed in the “frequency domain.”

The number of sub-bands, their frequency distribution, and overlap depend on the method of sub-band decomposition as well as a trade-off between frequency resolution, processing delay, computational complexity, available data memory on a given chip, and other affected parameters.

Personalized Hearing Amplification

Modern digital hearing aids use multichannel WDRC to achieve personalized amplification. This involves digitally amplifying and compressing the wide dynamic range of sounds heard by a healthy ear into a narrower range. As ambient sound increases, amplification is naturally reduced compensating for the loudness recruitment. For example, the original range of sounds from distant rustling to a pop concert is compressed and reproduced as a range from a vacuum cleaner to a pop concert, as shown in **Figure 8**.

“Multichannel WDRC” means that the input signals are processed in “frequency channels,” as shown in **Figure 9** for a four-channel system.




Figure 8: Wide band dynamic compression principles




Figure 9: Multichannel WDRC scheme



www.purifi-audio.com



⦿ Ultra High Performance Amplifier Modules

⦿ Ultra Low Distortion Driver Family

⦿ Unrivalled Specifications...

sales@purifi-audio.com



WARKWYN
Consultation • Analysis • Sales

Warkwyn utilizes and represents the most sophisticated analysis and measurement equipment available.



#1 Electro-Acoustic Hardware and Software Measurement Systems in the World.



Incredibly Unique Particle Velocity Measurement Systems for Noise and Vibration Localization. Hear it **and** See it.



Precision Made Type 1 & 2 SPL Measurement and Vibration Analysis Systems, Committed to a Quieter World Since 1969

For more information on Warkwyn's services, or for more information on our fantastic partners, please contact info@warkwyn.com or 612.759.4042 www.warkwyn.com

The channels are amplified/compressed differently according to a specific hearing loss. Each channel usually contains signal frequencies that correspond to a certain predefined frequency range and partially overlap. The channels can correspond to the frequency sub-bands one-to-one or combine several and even an unequal number of sub-bands, depending on the required frequency resolution. In practical implementations for hearing aids, the number of frequency channels can vary from 4 to 64.

In general, gain/compression parameters are set according to individual hearing thresholds measured directly or indirectly during a hearing test. **Figure 10** illustrates the input/output characteristics of a four-channel WDRC corresponding to a high-frequency loss.

No hearing loss at 500Hz corresponds to the magenta “unity gain” line where the input and output SPLs are equal. Hearing loss progresses with

Figure 10: Example compensation for a high-tone frequency loss




Figure 11: Acoustic feedback





Figure 12: Acoustic feedback cancellation

frequency and amplification is increased accordingly. For all frequency channels, the amplification is reduced with the signal level. It becomes unity at a specific level, typically near to the end of the loudness recruitment region for the frequency corresponding to the channel.

Acoustic Feedback Reduction

Howling, whistling, and squelching are common problems reported by hearing aid users that are due to acoustic feedback. As shown in **Figure 11**, acoustic feedback occurs because the receiver and the microphone are not perfectly separated, so the microphone signal reproduced by the receiver is picked up again by the microphone.

When no amplification is applied, the feedback signal at the microphone is naturally attenuated by the air conduction. However, when the microphone signal is amplified and reproduced by the receiver beyond a certain level, the signal enters an infinite microphone → speaker → microphone loop, with the sound amplified more and more in each circle, making the system unstable and uncontrollable. This devil’s loop does not occur simultaneously throughout the entire operating frequency range. Feedback problems begin earlier at frequencies with higher overall system gain. For hearing aids, the frequency range between 3kHz and 5kHz is the most problematic. The Feedback Reduction block is introduced to reduce the effects of acoustic feedback and allow higher amplification.

There are two major approaches to reducing acoustic feedback: feedback cancellation and feedback suppression. Feedback cancellation, shown in **Figure 12**, attempts to simulate the acoustic path from the receiver to the microphone, estimate the feedback signal and subtract it from the microphone signal.

This is generally done using an adaptive filter controlled by the result of subtraction. While this problem looks similar to acoustic echo cancellation (AEC), there is a fundamental difference that makes it more challenging. In AEC, the signal to be cancelled (the acoustic echo) and the signal to be preserved (the speech in the near field) are not correlated. Therefore, minimizing the cancellation error (the result of the subtraction) is the valid objective for the adaptive filter coefficients.

In acoustic feedback cancellation, both the signal of interest and the feedback are very similar signals with a high degree of correlation, especially at low frequencies. If an AEC approach is used directly, it will cause the adaptive filter coefficients to converge to incorrect values, resulting in cancellation of the signal of interest. To avoid the erroneous

convergence, a decorrelation, in the mathematical sense, is required between the output of the hearing aid and its input. Decorrelation can be accomplished, for example, by a small frequency shift or scaling of the output that is not yet perceived by the user as signal distortion or pitch shift. In general, frequency shifts of up to 3Hz can be used, which makes the task of decorrelation easier, but does not completely solve it. Ultimately, a compromise must be found between the convergence speed of the filter, its stability, and the cancellation of the signal of interest.

Feedback suppression does not attempt to simulate and predict feedback but works like an adaptive frequency equalizer that depends on the spectrum of the microphone signal, as shown schematically in **Figure 13**.

Acoustic feedback often manifests as a sudden and unrealistic increase in spectral amplitude within narrow frequency ranges. When detected, corresponding frequencies are attenuated to mitigate the unpleasant artifacts. However, acoustic feedback occurs rapidly in hearing aids, providing limited time for feedback suppression algorithms to analyze the microphone signal spectrum and

respond before the feedback becomes perceptible. Additionally, a challenge in feedback suppression is determining the appropriate time to release the suppressed frequency range and provide the required amplification.

To achieve the best performance, the feedback suppressor and canceller can be combined in one algorithm. In addition, binaural decisions can be made by comparing the signal analysis in the two hearing aids to detect feedback. In-ear microphones can be used to analyze the differences in arrival time, and many other “inventions” can be made.




Figure 13: Acoustic feedback suppression



MISCO is a leader in innovative, high-performance single and multi-channel audio amplifiers for OEM applications.

Digital Signal Processing (DSP) is available on many models to provide numerous user-controlled features including crossovers, filters, frequency equalization and dynamic range control. MISCO amps are available in a variety of stock configurations, with power supply and upgrade options, or we can design a custom solution that exactly fits your audio needs.

The field of acoustic feedback cancellation offers much scope for specific heuristic approaches and know-how for differentiation.

Acoustic Beamforming

Most modern digital hearing aids are equipped with two microphones, positioned apart from each other (front and back) for acoustic beamforming. This technique generates a directional “sensitivity beam” toward the desired sound, preserving sounds within the beam while attenuating others. Typically,




Figure 14: Acoustic beamforming principles




Figure 15: Typical microphone polar patterns

sounds from the front are preserved, while sounds from the sides and rear are considered unwanted and suppressed. However, different beamforming strategies can be applied in specific environments. Acoustic beamforming relies on a slight time delay, typically in tens of microseconds, for sounds to reach the two microphones. This time difference is dependent on the direction of the sound and the distance between the microphones, as illustrated in **Figure 14**.

Front sounds reach the front microphone first, side sounds arrive simultaneously, and rear sounds reach the rear microphone first. Acoustic beamforming utilizes this time difference to create a variable sensitivity called a “polar pattern” or “polar sensitivity pattern” that is directionally dependent. This is achieved through DSP techniques. **Figure 15** shows the four well-known theoretical polar patterns that can be achieved in a free field with dual microphone beamforming.

While sensitivity to front sounds stays consistent, sensitivity to sounds from other directions decreases with one angle having zero sensitivity. However, these theoretical polar patterns only apply in a free field and are not realistic in practical situations due to plain acoustic wave distortion caused by the user’s head and torso.

Acoustic beamforming can be either fixed or adaptive. Fixed beamforming has a static polar pattern that is independent of the microphone signals and acoustic environment, making it easy to implement and computationally efficient. Adaptive beamforming, on the other hand, changes the polar pattern in real time to steer the null of the sensitivity polar pattern toward the direction of the strongest noise, optimizing array performance.

Adaptive acoustic beamforming can be implemented in either the full frequency band or in frequency sub-bands. The full band scheme optimizes an average performance with a uniform polar pattern across all frequencies in the hearing aid frequency range, which may not be optimal for individual frequencies. In the sub-band scheme, polar patterns are tailored to different frequency regions to optimize performance across frequencies. This can be more efficient for practical scenarios where noise comes from different directions in different frequency regions.

Noise Reduction

The primary concern of individuals with hearing loss is difficulty communicating in noisy environments. While acoustic beamforming can effectively reduce noise from specific directions, it has a limited ability to mitigate diffuse noise—noise that emanates from

multiple sources without a distinct direction. Diffuse noise arises from multiple noise reflections off hard surfaces in a reverberant space, which further complicates noise directionality. For example, noise inside a car is often diffused, as uncorrelated noises from various sources such as the engine, tires, floor, and roof are diffused by reflections off the windows.

Compensating for hearing loss with WDRC amplifies sounds that are below the user’s hearing threshold, but it can also amplify noise and reduce the contrast between noise and the desired sound. Noise reduction technologies aim to mitigate these issues by suppressing noise prior to amplification by WDRC, thereby improving the output signal-to-noise ratio (SNR).

Today, most noise reduction technologies used in hearing aids are based on spectral subtraction techniques. The input signal is divided into frequency sub-bands or channels, and the noise is assumed to have relatively low amplitude and be stationary across sub-bands, while speech exhibits dynamic spectral changes with rapid fluctuations in sub-band amplitude. If the average noise amplitude spectrum is known or estimated, a simple “spectral subtraction” strategy can be employed to suppress noise, as shown in **Figure 16**.

The upper plot illustrates the sub-band signal amplitudes of a speech in noise signal across different frequency sub-bands. The relatively stationary noise amplitudes are averaged by the red line, while the speech signal amplitudes are significantly higher. In the lower plot, the sub-band noise amplitude shown by the red line in the upper plot has been subtracted from the sub-band input amplitudes using the spectral subtraction technique, which is applied to all sub-bands. Finally, the output full-band signal is reconstructed using the same method used for splitting, with the noise effectively reduced.

Digital Signal Processing

Alango Technologies has developed a full range of state-of-the-art DSP algorithms, software, and turnkey reference designs that can be scaled for various types of hearing enhancement devices, including OTC hearing aids, as well as TWS earbuds and headsets with amplified transparency and conversation boost.

Figure 17 displays the DSP technologies that are integrated into the Hearing Enhancement Package (HEP) developed by Alango.

While some of the technologies share names with those shown in Figure 4, they are based on Alango’s proprietary algorithms, which have been refined based on our experience and feedback received from customers.




Figure 16: Spectral subtraction principles




Figure 17: Alango DSP block for hearing aids




Figure 18: Five frequency channels system example

Alango's DSP software blocks use a uniform frequency sub-band approach. Sub-band splitting and combining is achieved through weighted overlap add (WOLA) methods based on STFT, where the STFT window and frame sizes are parameters. Optimal values for these parameters depend on the product, as they involve trade-offs between processing latency, frequency resolution, computational complexity, memory size, and power consumption. Two main versions exist, with STFT window lengths of 4ms and 8ms, providing frequency resolutions of 250Hz and 125Hz, respectively. The corresponding algorithmic/processing delays are 3ms and 6ms, respectively, for the 4ms and 8ms STFT windows.

Binaural Synchronization

The human brain processes information from both ears in a way that is not fully independent. This

allows us to identify the direction of sound arrival and filter out sounds that come from directions other than the direction of interest. However, if the sound processing in two hearing aids is not synchronized, it can reduce the brain's ability to locate sounds accurately, especially in noisy or reverberant environments. As a result, it can create an unnatural and confusing perception of sounds.

To address this challenge, Alango's DSP hearing enhancement technologies support two types of binaural synchronization: full and partial. Full synchronization requires that microphone signals from both ears be available in each device. This can be achieved through either wired or wireless (e.g., near field magnetic induction) connections. However, in practical situations where dedicated hardware is not available, wireless (e.g., Bluetooth) connections between the two sides of the head may be unreliable or associated with transmission delays. In such cases, Alango supports partial synchronization via a messaging mechanism, in which the processing blocks on the two sides exchange limited information about their states.

All information exchanged during the operation is encapsulated into a single data chunk that is formed on one side and accommodated continuously on the other side. The device software developer only needs to query the relevant APIs and transmit the data between the devices. This approach allows the two hearing aids to work together seamlessly, despite the lack of full synchronization. It ensures that the brain can effectively process sound information from both ears, improving the user's ability to locate and identify sounds in their environment.

Hearing Personalization

The basic principles of personalizing hearing with Alango technology are not different from those commonly used, as outlined previously. In Alango's processing approach, frequency bins resulting from the STFT are combined into "channels," as shown in **Figure 18** for an example with five channels.

Frequency channels are created based on the number and positions of their central frequencies, which are configured according to the product requirements. Each channel can have its own dynamic processing characteristic that corresponds to the user's frequency loss at the central frequency of the channel. With Alango technology, the dynamic characteristics of each channel are highly flexible and defined by six points. An example characteristic for one channel is illustrated in **Figure 19** by the magenta line. The plot displays the output SPL generated by the device in relation to the input SPL of sound in the frequency range that corresponds




Figure 19: Alango's six-point dynamics processing

to the channel's central frequency. The green line signifies the "unity line" of no amplification, where input and output levels are equal. The difference between the green and magenta lines represents the amplification.

To prevent the amplification of ultra-low-level sounds, no amplification is specified before point P6. Amplification begins at point P6 and increases (expands) to a maximum at point P5. In the P5-P4 region, the amplification is linear, preserving the original SNR. The P4-P3 region corresponds to "loudness recruitment," after which hearing is normal, and no amplification is required. Consequently, amplification gradually decreases (compresses) from the maximum at point P4 to no amplification at point P3. After P2, the output sound pressure is limited to prevent potential ear damage. It's important to note that this is just an example, and other characteristics can be defined.

Although specifying points P1-P6 directly is possible, the Alango auxiliary software simplifies the process by transforming the results of a hearing test into optimal values of P1-P6 for the user.

Optimal hearing amplification is nonlinear, and the difference between input and output levels (amplification) changes depending on the input level. As a result, natural sound perception and localization can only be achieved when amplification on both sides is coordinated (binaural amplification). Alango technology accomplishes this by exchanging the sub-band signal levels used to calculate amplification between two devices.

Feedback Cancellation

Alango's acoustic feedback prevention algorithms leverage both feedback cancellation and feedback suppression principles as previously outlined. To achieve feedback cancellation, proprietary adaptive filters based on LMS principles are utilized. To avoid unnecessary calculations, the range of frequencies where the adaptive filters are active is a parameter. Decorrelation is performed by frequency shifting through multiplication of each complex sub-band signal by a complex factor of unit amplitude rotating with the speed as required for the frequency shift. For uniform frequency shift across all frequencies, the same factors are used in all sub-bands, while different factors with phases proportional to the sub-band central frequency are used for frequency proportional shift. In practice, a combination of both is used to obtain the best results, which can also depend on the application profile, such as conversation, live music, nature sounds, and others.

Normal acoustic signals, such as speech, are generally wideband, while feedback typically occurs at a specific frequency where the amplitude quickly rises. To prevent feedback, Alango's algorithm analyzes microphone signal sub-bands and identifies those with amplitudes significantly exceeding adjacent bands, which often occurs at a specific frequency. If this situation persists beyond a predefined time (e.g., 50ms), the algorithm attenuates the signal in that sub-band.



Loudspeaker Design Cookbook 8th Edition

For the First Time Ever...

Available in
Full Color
and a
Hardcover Collector's Edition



Learn to Design Professional Sounding Loudspeakers for:

- **Residential**
- **Automotive**
- **Home Theater**

Vance Dickason's top-selling *Loudspeaker Design Cookbook* is an essential must-read for serious speaker designers, audio engineers, and anyone looking to master speaker-building technology. With more than 140,000 copies in print, author Vance Dickason's *Loudspeaker Design Cookbook* is not only regarded as the loudspeaker engineering "bible," but the #1 resource on the construction of high-quality loudspeaker systems for two-channel, home theater, pro sound PA speakers, and studio monitors, as well as car audio applications.

Get your copy at: CC-Webshop.com/LDC8

Alango's binaural feedback suppression algorithm takes advantage of the fact that feedback usually does not occur simultaneously on both sides of the head. It analyzes the spectra differences of microphone signals from both sides and attenuates sub-bands with significant and lasting differences. This binaural feedback suppression works for both full and partial synchronizations.

Acoustic Beamforming

Alango's acoustic beamforming technology supports multiple microphones, but the dual microphone configuration is more common for hearing aids and TWS earbuds due to space and interaural connection limitations. The algorithm operates in independent frequency sub-bands, allowing each sub-band to have its own polar sensitivity pattern. This outperforms full band beamformers in practical scenarios with multiple noise sources or quasi diffused noise, and near acoustic barriers like the human head. Alango supports both full and partial beamforming synchronization between ears, improving directivity in the front direction and sound source localization for comprehension.

Wind Noise Reduction

Wind noise can be a significant issue for hearing aid users, as turbulence affecting the microphones is amplified and worsened by hearing loss. To mitigate this problem, Alango has developed wind noise reduction (WNR) technology, which uses information from two or more device microphones to detect and reduce wind noise. The algorithm estimates the correlation between device microphone sub-signals in each frequency band and identifies and attenuates sub-bands affected by wind.

If the microphone signals from both sides are available in each device (full synchronization), signals from the less windy side can be used for both ears, potentially improving speech intelligibility during conversations in windy conditions. However, this approach may reduce the ability to locate sound sources. When full synchronization is not possible and the microphone signals from the opposite side are not available, partial wind noise reduction synchronization can still be effective in reducing wind artifacts.

Noise Suppression

"Alango's noise reduction technology operates in frequency sub-bands and uses a "spectral noise gating" approach to classify sub-band signals as "noise" or "signal" based on a proprietary voice activity detector detecting human speech. Noise reduction is then applied to the "noise" sub-bands, while the "signal" sub-bands are passed unmodified. This technology adapts quickly to changes in noise level and helps to reduce "musical noise," a common problem with spectral subtraction. For the best user experience, noise reduction on both ears should be synchronized, which Alango supports through both full and partial synchronization options.

Own Voice Attenuation

Hearing one's own voice amplified is one of the main complaints of hearing aid users. It is particularly acute when the ear canal is sealed to enhance the low-frequency range and reduce the likelihood of acoustic feedback. Ideally, the user's voice must be detected and removed from the microphone signal. However, distinguishing the user's voice from other voices in the microphone signal is very difficult and unreliable, to say the least, and requires full synchronization between the two ears.

Alango's approach uses an additional bone conduction sensor integrated into a hearing aid that is attached to the user's ear. The user's speech causes tiny vibrations of the skull bones, which are picked up by the sensor. The signal from the bone-conducted sensor sounds very similar to the user's speech picked up by the microphones but is limited to about 1000kHz or less. However, the bone-conducted signal is completely isolated from all ambient noise and can therefore be reliably used to detect and attenuate the user's voice.

While adding an additional sensor may seem costly and space consuming, it saves both compared to the NFMI chip and its antenna used in hearing aids for binaural synchronization. In addition, the same bone conduction signal greatly improves the quality of the outgoing speech signal for voice communication via a Bluetooth connection to a phone, which is becoming a "must have" for hearing aids. Small, inexpensive MEMS bone conduction sensors are now being manufactured by several companies for use in TWS earbuds.

Alango Dereverberation Technology

Speech intelligibility for people with hearing loss is drastically reduced by reverberation in the room. Although the problem of restoring the original signal from a reverberant signal is generally unsolvable, some improvements are possible in practice. The Alango HEP includes reverberation suppression technology that significantly improves perceptual quality and reduces listener fatigue from additional mental effort by people with hearing loss.

Voice Communication and Streaming Audio

Modern hearing aids are also communication devices similar to TWS earbuds. This functionality requires DSP tasks such as acoustic echo cancelation, microphone beamforming, noise reduction, automatic gain control. In addition, the incoming voice during a phone call needs to be modified in the hearing aids to improve

About the Author

Dr. Alexander Goldin is a technology visionary, entrepreneur, and business leader committed to advancing innovation and improving people's lives. Alexander earned his PhD in mathematics and physics from the Russian Academy of Science. In 1991, he moved to Israel where he served as a technology leader for several international companies including IBM Research Lab, Johnson & Johnson, and Altec Lansing. In 2002, he founded Alango Technologies to develop innovative digital audio processing technologies that could revolutionize the quality of communication between people. Under his leadership, Alango has become a global leader in voice and audio processing technologies, serving a wide range of industries, including telecommunications, automotive, hearing aid, and consumer electronics.



intelligibility and compensate for the user’s hearing loss. All of these task DSP technologies are offered by Alango in its basic Voice Communication Package (VCP). In addition, Alango offers advanced proprietary voice communication DSP technologies such as OnlyVoice and EasyListen.

OnlyVoice technology uses the bone conduction voice sensor, which is also used to reduce the user’s own voice. However, the opposite is used for voice communication. Instead of attenuation, the user’s voice is emphasized and isolated from ambient noise to be transmitted to the other side of the connection.

EasyListen significantly improves the intelligibility of incoming speech from the far end by dynamically slowing it down in real time. The accumulated delay is quickly recovered during pauses.

With improved speaker technology and Bluetooth connectivity, hearing aids are becoming suitable for listening to streaming audio. As with voice communication, incoming audio must be personalized in terms of frequency response and dynamic range, or the user will miss much of the content, which cannot be compensated for by changing the volume. The Alango Audio Enhancement Package (AEP) enables such customization by applying multichannel WDRC with user-specific parameters to the audio signal and reducing the spectral difference so that soft, typically high-frequency sounds are not masked by strong low frequencies.

When audio streaming is used to listen to podcasts or news, the incoming voices of participants are modified in a manner

similar to telephone conversations to improve their intelligibility for hearing-impaired users. For such use cases, Alango’s EasyListen technology can further improve speech intelligibility by slowing down the often-rapid speech of podcasters and newscasters.

The Hearing Enhancement Roadmap

The possibilities for DSP technologies in the realm of hearing are endless. In addition to the solutions we have already discussed, there are many other exciting opportunities to explore. For example, we could harness the power of neural networks to develop even more advanced noise reduction techniques or leverage low latency wireless connectivity to create smart remote microphones with broadcasting capabilities. Additionally, active occlusion effect management and dynamic vent technologies offer promising avenues for addressing common hearing problems. Beyond this, we can even consider innovative approaches like hearing sense substitution via tactile sense and real-time speech tempo manipulation for improved intelligibility. 

Resources

- Alango Technologies, Ltd. | www.alango.com
- BeHear – Alango Technologies, Ltd. | www.alango-behear.com



The QA403 Audio Analyzer...\$599

- 110 dB THD+N
- +18 dBV Output (4 ranges)
- +32 dBV Input (8 ranges)

- USB Powered
- Fast, automated testing
- REST Programmable



www.QuantAsylum.com
Designed and built in the State of Colorado USA

Over a Century of Home Loudspeaker Development



By
Richard Honeycutt

Photo 1: Some early phonograph horns were conical, with a flared mouth.

The history of home speaker development is always fascinating to revisit if we want to fully understand how electroacoustic devices impact our listening experience and their interdependence with room acoustics. This article explores that evolution and how we quantify a loudspeaker's quality.

Often, when we think of the historical development of home audio—either for listening rooms or for home studios, we focus primarily on the electronics. But certainly, the device providing the interface between the electrical signal and the acoustical end product is at least equally important. The development of the electroacoustical link between the electronics and our ears has undergone a sea change fully as great as that of the electronic part or a sound system. But unlike the electronics, which have generally converged into hollow-state and solid-state streams, electroacoustic devices have diverged into multiple streams, of which different ones are favored by different listeners, and tend to be used for different applications. Let's start with how we quantify a loudspeaker's quality.

Frequency Response

Probably the first question that comes to mind when we think of loudspeaker quality is the frequency response: Can I feel the tympani and bass drum, hear the low bass viola, and experience the clear high notes of the piccolo and the shimmer of the violins? As we become more experienced listeners, we learn to notice whether singers sound like real people and whether a trumpet sounds like a trumpet rather than an oboe.

The frequency range of human hearing is considered to be in the 20Hz to 20kHz range for average healthy young adults, although prior exposure to loud sounds will decrease the upper limit in some individuals. We have no reason to

think that this range has changed significantly during the last century, but the range of frequencies a loudspeaker can reproduce certainly has—both because of advances in speaker technology and as a result of expectations changing as the frequency range of source devices (from telephones to AM radios to shellac phonograph records to FM radios to vinyl records to CDs).

Maximum Sound Output

A loudspeaker designed for use as a telephone receiver doesn't need to put out much acoustical power. But as loudspeakers have been called upon to provide sound for lecture halls, theaters, cinemas, large public gatherings, and outdoor concerts, the requirements for maximum sound output have increased tremendously.

The maximum sound output a speaker can produce depends upon thermal and mechanical limitations. The mechanical limits are imposed by the need to produce a high-volume velocity—cubic meters of air pumped back and forth per second, which translates into mechanical structures (cones, diaphragms, etc.) that must combine a certain area with the necessary force and range of linear motion. The thermal limit results from the fact that no mechanical device is 100% efficient, meaning that some percentage of the electrical power fed to the speaker is converted to heat rather than to sound. When this waste heat becomes excessive, it can lead to damage or destruction of the speaker. The more efficient a speaker is, the less its maximum output is constrained by thermal limitations.

Distortion

The sound output does not stop abruptly when the mechanical or thermal limit is reached, unless a very powerful spike is involved. Normally, the instantaneous acoustical power output is gradually reduced as the electrical power approaches the limit. This causes the sound waveform to lose its resemblance to the electrical waveform. We call this process nonlinear distortion, or simply distortion. Different loudspeaker mechanisms produce different amounts and kinds of distortion.

Harmonic distortion (HD) involves a speaker's nonlinearity adding harmonics (sound at integer multiples of the input signal frequencies). These change the timbre of the sound—making a trumpet sound like an oboe, for example. Intermodulation distortion (IMD), (or simply modulation distortion) adds extraneous frequencies that are not harmonically related to the input signal. These frequencies effectively constitute wrong notes and therefore have a more objectionable sound than those produced by HD. Depending upon the signal frequency and the program material, the acceptable maximum HD can range from 1% or 2% to about 15%. Maximum acceptable IMD has been determined by different investigators to be in the range of 0.1% to 2%.

Directionality

The directional pattern of a loudspeaker's radiation does not seem to have been of much importance to users of early models. Before speakers' widespread use in cinemas in the 1930s, it was enough that they reproduced sound with some accuracy.

Early speakers all tended toward beaming of high frequencies because of their size and shape. To my knowledge, the Bell Telephone Laboratories' Symposium on Auditory Perspective reported in the January 1934 issue of *Electrical Engineering* marks the first time real attention was paid to optimum speaker directionality. The conclusion of the participating Bell Labs engineers was that for a long shoebox-shaped hall with the speakers positioned at the front end, good audio perspective (but certainly not surround sound) required two- or preferably three-channel stereo. Each speaker was essentially omnidirectional at the lowest frequencies, narrowing to about a 90° horizontal × 40° vertical pattern at mid and high frequencies.

Predecessor of the Loudspeaker: The Acoustic Horn

Home music systems predated loudspeakers. The acoustic phonograph used an acoustic horn to adapt the high mechanical impedance of the phono stylus

to the low impedance of the air at the horn mouth, by attaching the stylus to a metal diaphragm that fed the throat of the horn (**Photo 1** and **Photo 2**). Little was known about how horns operated prior to the 1919 publication of Arthur Gordon Webster's theoretical investigation of the subject. Hence, a variety of horn shapes were introduced, often based upon ease of manufacture or visual aesthetics. The frequency responses of these two horns were limited at the low end by the flare profile (conical, exponential, etc.) and mouth diameter and at the high end by the mechanical characteristics (mass and stiffness) of the stylus-diaphragm assembly. Although users may not have noticed, the horn shown in Photo 2 had a more extended low-frequency response because of the quasi-exponential flare. The high-frequency response on axis would have benefitted from beaming resulting from the same cause. In any event, reproduction of these phonographs was pretty much limited to roughly 500Hz to 3kHz at best.

The maximum sound output of an acoustic phonograph was limited by the minuscule acoustic power fed from the record groove into the horn throat. With such low acoustic power, distortion mainly resulted from the mechanical recording/playback process. The horns transitioned from omnidirectional at low frequencies to very directional at high frequencies.

Loudspeakers as We Know Them Today: The Jensen Speaker

Although British physicist Sir Oliver Joseph Lodge invented an experimental moving-coil loudspeaker in 1898, it was actually Danish-American engineer Peter L. Jensen who invented



Photo 2: Other horns had a more decorative design.



Photo 3: This sculpture in Napa, CA, commemorates Peter Jensen's first commercial loudspeaker.



Photo 4: The WE CW-929 "loudspeaking telephone" used a design similar to Jensen's.

the first practical moving-coil speaker in 1915. Jensen's speaker differed from early telephone receivers in that, while both used magnets whose fields interacted with the field of a "voice coil" (an electromagnet fed by the audio-frequency electrical signal), the coil in the telephone receiver was stationary, while Jensen's voice coil was attached to the diaphragm, resulting in better linearity and reduced distortion (**Photo 3**). The section of Jensen's speaker that converted electrical energy to sound worked into what amounted to an air spring at the horn throat; thus, it was called a compression driver. The Jensen speakers came to public attention in the first public address system, on Christmas Eve in 1915 in San Francisco, CA, to enable a gathering of 75,000 people to distinctly hear all the speeches and announcements. The public address systems were marketed under the name of Magnavox and were used all over the United States.

By 1918, Western Electric was producing "loudspeaking telephones," such as the CW-929, that were used as speakers for radio receivers (**Photo 4**).

Cone Speakers

Due to the reliance on small horns, the Jensen, Western Electric, and similar speakers only excelled the performance of acoustic horns in one area: maximum sound output. The bass response only extended down to about 500Hz, although the human ear/brain system is able to produce the sensation of lower frequencies from the harmonics of those frequencies. Atwater Kent and perhaps other radio receiver manufacturers sought to improve the bass response of speakers by attaching a paper cone to the diaphragm of what was essentially a telephone receiver earphone. The improvement was slight.

It was not until General Electric's Edward Washburn Kellogg and Chester Williams Rice developed their moving-coil cone loudspeaker (patented in 1925) that truly improved bass performance became possible. Rice's patent drawing is shown in **Figure 1**. Note that the speaker only had a 4" cone, and the cone's outer edge was unsupported: the addition of a surround to prevent the back wave from cancelling out the front wave at low frequencies, and the move to larger cones were later improvements. But while the immediate effect of the Kellogg-Rice speaker on home audio quality was not very great, the introduction of the concept itself was revolutionary, and led to the design used for most speakers sold today.

Many of the first radio speakers using the Kellogg-Rice design were mounted in open-back cabinets, but later, engineers realized that if the speaker were mounted in a sealed cabinet—approximating an infinite baffle—better bass response could be realized.

Cone-Driven and Combination Horn Speakers

As has been true throughout much of home audio history, improvements in home audio devices grew from strides made in cinema audio and live sound. The Bell Telephone Laboratories and its manufacturing partner Western Electric shared a long competition with RCA for domination of the cinema sound market. Both produced full-range horn systems whose electroacoustic conversion efficiency exceeded that of direct-radiator speakers by a factor of 20 or more, increasing the maximum sound power for a movie house. Acoustical research by both companies led to increased understanding of the physics of horns, and thus to improved and extended frequency response. By decreasing the maximum cone/diaphragm excursion required of the cone or compression driver, the horn improved efficiency. Although the Altec A2 "Voice of the Theater" (VOTT) full-range horn system shown in **Photo 5** (introduced in 1945) was far too large for home use, it inspired engineers and inventors to create superior home speakers




Figure 1: This is C. W. Rice's patent drawing for the first moving-coil cone speaker.

based upon similar concepts. To help the bass efficiency more nearly match the mid-high-frequency efficiency, a tuned port was added to the woofer enclosure. Although such "bass reflex" systems became more popular in the late 1940s and 1950s, the VOTT systems were the best-known examples of their early use in professional audio. Altec advertised the frequency response as 50Hz to 20kHz, but no measurement tolerances were given.




Photo 5: This Altec Lansing A4 cinema speaker was based on a Western Electric design.



HOLOPHONIX 128 and 64.


The new immersive sound processors,
now available in two formats.

Perfectly optimised for real-time sound spatialization and
immersive experiences, the new HOLOPHONIX processors feature
the latest technological innovations designed by Amadeus,
in coordination with IRCAM research center.



They have chosen HOLOPHONIX :

Aix-en-Provence Festival, Avignon Festival, Beijing Stadium, Bob Wilson,
Chaillot National Theatre of Dance, Chaumet, Comédie Française,
ETH Zurich, French National Library (BnF), La Scala, Lyon National Opera,
Pantheon, Tsinghua University, and many more...

DO IT YOUR WAY


Panels/Boxes/Cables

WITH REDCO AUDIO'S

"Build Your Own" web pages

Full control for design, editing and more.

WE'LL BUILD IT JUST THE WAY YOU WANT!



REDCO AUDIO

www.redco.com 203-502-7600




Photo 6: The “Shearer Horn” used a folded bass horn to increase the efficiency at low frequencies.

Although the first prototype was designed and built before the introduction of the VOTT systems—in 1935 because of collaboration among MGM, Lansing Sound, Western Electric, and RCA—the “Shearer horn” (**Photo 6**) arguably had the greatest effect on home audio of any full-range professional horn system. (The horn was named for Douglas Shearer, head of MGM’s sound department.) Using RCA’s folded bass horn suggested by Harry Olson, a Stephens multicellular mid/high horn, and a Lansing compression driver, this system, also marketed as Western Electric Diaphonic and RCA Photophone horns, had the great advantage of providing full-range horn loading without requiring a large bass cabinet.

In 1923, young Paul Wilbur Klipsch was on a holiday from his engineering studies at New Mexico A&M, and he visited his mother. He noted that she had moved their Edison phonograph to the corner of the room, and that it sounded much better there. Sometime during his stint as a maintenance engineer for electric railroads in Chile, a friend gave him a folded horn speaker to use with his radio receiver. These two events led Klipsch to research why the phonograph sounded better in the corner and why the small, folded horn speaker was able to produce good (for the time) sound at a comfortably high listening level, when being driven by only a 0.7W amplifier.

He found that a sound source would produce a 6dB louder sound if mounted against a floor rather than up in the air; backing it up to a wall




Figure 2: The original Klipsch folded bass horn was designed to take advantage of corner loading.

Photo 7: The early production Klipschorn used a mid-high frequency horn mounted on top of the folded bass horn.



References

- [1] P. W. Klipsch, “A Low-Frequency Horn of Small Dimensions,” *Journal of the Acoustical Society of America*, 13, October 1941, pp. 137-144.
- [2] A. N. Thiele, “Loudspeakers in Vented Boxes,” *Proceedings of the Institution of Radio Engineers Australia*. 22 (8), 1961, pp. 487-508.
- [3] F. V. Hunt, *Electroacoustics*, Cambridge, Harvard University Press, 1954, p. 168.
- [4] Ibid, p. 169.




Photo 8: The Voigt corner horn used only a single driver. High frequencies were dispersed by a reflecting “distributing cone.”

added another 6dB, and adding a second wall so that the speaker was in the corner boosted the sound level by yet 6dB more. Thus, simply by positioning a speaker in the trihedral corner of a room, you could make the sound 18dB louder. But to avoid interference problems from sound reflections off the walls and floor, this effect was only useful if the speaker was within a quarter-wavelength of the actual corner.

He also found that, as Webster had described in his seminal paper, the efficiency of a sound radiator depended upon what we would now call the mechanical impedance match between the speaker diaphragm and the air into which the speaker radiated. Horns improved this match, potentially increasing the efficiency of a speaker by up to a factor of 25 or more.

Using this understanding plus his knowledge of physics and engineering, he designed a folded horn loudspeaker created to fit in the corner of a room. **Figure 2** shows the drawing of the folded bass section of the original Klipschorn, as published in a professional journal in 1941 [1]. Tests indicated a smooth response from 40Hz to 400Hz. Higher frequencies were reproduced by a high-frequency horn and driver mounted atop the bass horn. **Photo 7** shows an early production Klipschorn. With time, improvements were made in the low- and high-frequency horns and drivers and the crossover network. The current model AK6 has a frequency response flat within $\pm 4\text{dB}$ from 33Hz to 20kHz, and a power sensitivity of 105dB at 1W at 1m.

The Klipschorn was neither the only nor even the first corner horn produced. In 1934, Paul Voigt introduced a non-folded corner horn fed by a single speaker (**Photo 8**). Later, he partnered with O. P. Lowther to develop a range of Lowther corner horns (**Photo 9**), many of which were installed in cinemas. Lowther has since added some folded corner horns to its line.

In 1963, Klipsch and Associates introduced the La Scala “Industrial” (Paul’s term for “pro”) speaker, a three-way full-range folded horn designed for portable live sound applications. The prototype was clearly inspired by the Shearer horn, as you can see from **Photo 10**. In later production versions, the bass horn was



Photo 9: This is a production model of the Voigt-Lowther non-folded corner horn.



Photo 10: The early Klipsch La Scala speakers were three-way versions of the Shearer horn.

rotated 90°, resulting in the appearance shown in **Photo 11**. Many clones of this “W-box” design were manufactured by a variety of home stereo and pro-sound speaker manufacturers.

The Rebirth of the Bass Reflex Speaker

Through the years, many speaker designers were aware that the low-frequency efficiency of a speaker could be doubled if the “back wave” (radiation from the rear of the cone) could be utilized instead of being absorbed within the speaker enclosure. Albert L. Thuras of Bell




Photo 11: Two La Scalas are shown serving as L-R speakers in a home theater.




Figure 3: The dashed line shows the improvement in bass response afforded by a properly tuned port.

Telephone Laboratories patented a design using a tuned port leading from the rear of a loudspeaker diaphragm to the outside atmosphere, and included the graph shown in **Figure 3**, illustrating the effect of the improved bass efficiency. Although this concept was employed in the VOTT speakers, it was poorly understood and thus mainly of interest to experimenters and hobbyists until the 1961 publication of A. Neville Thiele’s theoretical analysis [2]. Variously called “ported,” “vented,” or “bass reflex” speakers, these provide both advantages and disadvantages, but many prominent home stereo speaker manufacturers now include such speaker systems in their product lines.

If the port is replaced by a long—usually folded—tube lined with acoustically absorptive material, it acts as an acoustical transmission line (TL). Until very recently, this design has not been as well understood as ported or sealed designs and required much tweaking and careful tuning in the design stage and even during final testing. Some highly respected audiophiles are great fans of transmission line speakers, and several companies have successfully marketed TL speakers. TL speakers are sometimes referred to as acoustic labyrinth speakers.

Notable Sealed Loudspeaker Systems

To produce really low bass using sealed enclosures, some manufacturers in years past used really large cabinets designed for corner locations. Electro Voice created the Patrician (**Photo 12**), employing a 30” woofer, 12” low-midrange, and a compression driver feeding a diffraction horn—claiming a frequency response of 15Hz to 23kHz. Such speakers often found their way into radio station lobbies, but seldom into home living or listening rooms. This created a market for smaller bass-capable speakers. Edgar Villchur created the “acoustic suspension” speaker to serve this market. Not just a small, sealed speaker system, this design incorporated a powerful speaker magnet with optimized enclosure volume and damping in the cabinet to replace the mechanical springiness otherwise supplied by a stiff spider and cone surround. The Acoustic Research AR-1 (**Photo 13**), the first commercial acoustic suspension speaker system, used this design to capture a significant share of the home hi-fi market.

Planar Speakers

In his 1954 book, the late Dr. Frederick V. Hunt, former head of the Acoustics Laboratory at Harvard University, stated that, [the internal

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.



acoustical impedance of an electrostatic transducer combines with the low acoustical impedance of the air into which it radiates, such that] “there should be an essential ‘rightness’ about the employment of electrostatic mechanisms for the transduction of airborne sounds” [3]. Hunt goes on to list more than a half-dozen references to papers discussing electrostatic speakers, published between 1926 and 1933 [4]. One might well ask why, with such a long history, have electrostatic loud speakers (ESLs)—rather than electromagnetic ones— not become the predominant design technology?

Let’s begin with a brief introduction to the operation of an electrostatic speaker. If a thin conducting sheet or diaphragm (a metal foil or metallized plastic film) is mounted parallel to a conductive backplate (“stator”) and a polarizing voltage source is connected between the diaphragm and the stator, the diaphragm will be electrostatically attracted to the stator. If an alternating voltage is superimposed upon the polarizing voltage, the diaphragm will move in synchrony with the alternating voltage. Because of the inverse square law, the positive (away from the stator) and negative

(toward the stator) excursions of the diaphragm will not be equal, and significant even-order harmonic distortion will result. This distortion can be eliminated if the diaphragm is sandwiched between two stators,



Photo 12: The Electro Voice Patrician boasted 15Hz to 23,000Hz frequency response.



Photo 13: The AR-1 was the first commercial acoustic suspension speaker.

HFFx Auto



The Ultimate Automatic High Frequency Audio Enhancement Solution


Let "Auto" mode do the magic - adapting seamlessly for optimal, pristine sound

- **Revive Your Audio:** Breathe life into tracks compromised by lossy sub-band coding like MP3, AAC, and Bluetooth. Say goodbye to dull, lifeless sounds.
- **Seamless Audio Transition:** Effortlessly switch between TV channels with different audio bandwidths as HFFx Auto ensures a full and consistent audio experience.
- **Unlock Full Audio Potential:** Enhance recordings limited by technology constraints, such as analogue tape or low sample rates, unveiling a richer sound.
- **Elevate Your Sound:** Up-convert standard audio (44.1k & 48kHz) to High-Resolution audio rates like 88.2KHz & 96KHz and even up to 176.4KHz & 192KHz for an incredibly open and immersive audio experience.

Download the product brief at:

<http://www.oxford-digital.com/products.htm>

Contact: E: info@oxford-digital.com T: +44 845-450-5664



Incoming audio with missing HF



Reconstructed Harmonics



Reconstructed Signal

www.oxford-digital.com

Excellence in audio technology





Figure 4: These graphs show the directional characteristics of a circular diaphragm as a function of the diameter and wavelength. (Olson, Harry F., Acoustical Engineering, New York, Van Nostrand, 1957, p. 46.)



Photo 14: Janszen 1-30 tweeters were first offered with AR-1 woofers in 1956.

Figure 5: The Heil Air Motion Transformer achieves superb performance through the use of accordion-like squeezing motion rather than pistonic action. (By Oskar Heil—US patent 3636278)



creating a push-pull driving force. The stators are perforated to allow sound to pass through. When optimally designed and constructed, electrostatic speakers can provide excellent sound with quite low distortion. However, it was not until the 1950s that suitable metallized plastic films became available for diaphragms.

The limited excursion of the diaphragm to avoid its bottoming out on the stators makes maximum SPL and bass response challenging. Overcoming this challenge usually involves large speakers, and as you can see from **Figure 4**, the size creates beaming as frequency increases. (The effect of diaphragm size on directivity is about the same for a square diaphragm. For a 1'x1' diaphragm, the graphs correspond to frequencies from 283Hz to 4500Hz.) Some manufacturers have chosen to pair an electrostatic mid-high-frequency speaker with a moving-coil bass speaker to provide adequate low-frequency SPL.

In 1949, Arthur A. Janszen patented an ESL based upon his work with Dr. Hunt. He used a speaker formed as a segmented arc of ESL panels to increase the high-frequency radiation angle (**Photo 14**). Perhaps the most impressive ESL system was the full-range KLH Nine, introduced in 1959.

ESLs are not the only planar speakers. If the stators of an ESL are replaced with thin, flat magnets and a voice coil formed of a thin aluminum ribbon is attached to the plastic film diaphragm, the result is a planar magnetic speaker. These were introduced by Magnepan in 1969.

The Heil Air Motion Transformer

In 1972 Oskar Heil invented a new type of loudspeaker that has been described as the offspring of an ESL and an accordion. Using the planar magnetic transduction principle, the Heil speaker's diaphragm displaces air using a squeezing motion; hence, the moniker "air motion transformer"... (**Figure 5**). Since the device is inherently dipolar, it is used primarily for mid- and high-frequency reproduction. The original Heil device was designed to operate at frequencies from 650Hz to 23kHz.

Summary

Based upon the work of scientists and inventors from 1915 until the late 1960s, foundations were laid for today's wide variety of home stereo (and home theater) speakers. In this article, I have made no attempt to trace the development of these speaker designs to the present. I hope the readers find the history of home speaker development as fascinating as I do!



Elevate your hearables & wearables design with CEVA's comprehensive Audio IP solutions

Immersive audio experience and the finest sound quality with CEVA's Audio Software and DSP solutions, to make your devices stand out

Audio Software Solutions for Immersive Audio Experiences



RealSpace: Spatial Audio & Head Tracking Solution

Embedded spatial audio rendering software engine that redefines immersion with precise head tracking technology



ClearVox: Speech Clarity Meets Innovation

AVS qualified audio front-end voice processing software, for speech clarity and clear voice communication



WhisPro: Your Voice, Recognized and Understood

Neural network-based speech recognition technology to empower voice assistants and IoT devices

Audio Platforms for the Smart, Connected World



CEVA-BX1: Transform Your Audio Experience

Audio DSP for battery operated, high-performance, audio and voice applications



CEVA-BX2: Redefine Audio AI Applications

Powerful Audio AI DSP, optimized for high-end audio and AI processing on edge devices



Bluebud: Complete Wireless Audio Solution

The ultimate DSP-enabled Bluetooth Audio IP platform for TWS earbuds and wearables

Request a demo

www.ceva-dsp.com

info@ceva-dsp.com



Easing Deployment of Automotive Audio Functions

By
John Whitecar

In this article, we delve into the signal processing challenges that designers encounter when aiming to create elaborate automotive audio systems that transform driving into a harmonious experience.

In the world of automotive design and technology, the focus has traditionally been on performance, safety, fuel efficiency, and aesthetics. However, in recent years, the importance of audio systems in automotive products has gained significant weight. Audio systems are no longer just accessories. They have become integral components that contribute to the overall driving experience, ease of use, comfort, and safety of vehicles.

Delivering a multifunction audio system is rapidly growing more complex and challenging. There is now a transition from a hardware-based approach to a centralized design supporting the software-defined vehicle audio system. This design approach requires new tools and methods to deliver the required audio features and to support feature delivery to the customer over the life of the vehicle.

Automotive audio systems have come a long way since the first AM radios were introduced in cars. Initially, car radios were considered a luxury, but over time, they became standard features. With advancements in technology and the escalation of consumer expectations, audio systems expanded, evolving into sophisticated systems that incorporate various electronic components such as amplifiers, telematics units, chime modules, pedestrian warning modules and more. In a hardware-centric system design, independent electronic control units (ECUs) support a specific mixing of audio features. The downside is that once a vehicle leaves the factory, the feature set is locked, and independent features do not always work well together as a complete system.

As far as audio products go, an automobile is one of the most complex on the market. Automotive OEMs must choose audio features and functions from a wide array of technologies, suppliers, and brands. The vehicle, and especially the driver seat, serves as living room, communication hub, and vehicle control center requiring bi-directional interaction with vehicle through voice and

audio alerts, and active noise management to create a focused, relaxed environment. Immersive sound staging must work well with popular streaming multichannel and stereo entertainment sources and support an assortment of mixing and ducking functions depending on the priority of active sources, alerts, voice prompts, and vehicle chimes.

Automakers need the ability to easily customize audio features for each vehicle configuration and audio brand relationship. Having flexibility to independently select these features from different sources allows the OEM to choose the best of each, combining features in new ways that create unique interactions.

Automotive audio systems go beyond playing music and radio broadcasts. Communication and connectivity features are incorporated that enable hands-free calling, voice commands, and smartphone integration. Bluetooth technology allows drivers to sync their smartphones with the car's audio system, enabling them to make and receive calls without taking their hands off the steering wheel.

Voice recognition systems have improved significantly, allowing drivers to control various functions of the vehicle and audio system using voice commands. This enhances safety by minimizing distractions and keeping the driver's attention on the road. Many vehicles require multiple voice engines to support direct vehicle control, off-board navigation, or streaming services as well as applications such as Apple CarPlay or Android Auto. All these voice applications require embedded signal processing to provide clean audio from the onboard microphones in the vehicle.

The shift toward electric vehicles has brought about changes in automotive design and technology. Electric vehicles (EVs) are known for their quiet operation compared to traditional internal combustion engine vehicles. While reduced noise pollution is a positive aspect, it has also highlighted the importance of audio


systems in EVs in two significant areas: There is a need to add external sounds to alert pedestrians of the approaching vehicle. Also, the lack of an internal combustion engine exposes road noises to the cabin, so road noise cancellation technology is required to provide a quieter interior environment.

Signal Processing Challenges

In the realm of complex automotive audio designs, where the pursuit of auditory perfection meets the challenges of a confined and highly reflective space, signal processing emerges as both a critical tool and a formidable challenge. The intricate task of manipulating audio signals to achieve optimal sound quality within the unique acoustics of a vehicle demands a deep understanding of both audio engineering principles and the intricacies of automotive integration.

Signal processing is the art of manipulating audio signals in various ways to achieve desired outcomes. It is a multifaceted process that involves techniques such as equalization, crossover, time alignment, compression, and more. In the context of automotive audio, signal processing serves as a bridge between the raw audio output of the source unit and the optimized sound that reaches the listeners' ears.

Modern vehicles often feature advanced infotainment systems with user preferences, profiles, and seat-specific audio adjustments. This demands dynamic signal processing that can adapt to individual preferences and optimize the sound for each seat position.




Audio designs need to accommodate a variety of audio sources, from radio and CDs to streaming services and smartphones. Signal processing tools must be versatile enough to adapt the audio output to different sources and maintain consistent quality.

Signal processing tools must seamlessly integrate with the vehicle's existing electronics. This involves interfacing with infotainment systems, steering wheel controls, and vehicle communication networks (e.g., the CAN bus). Compatibility challenges can arise, requiring engineers to navigate the intricacies of automotive electronics.

Addressing the signal processing challenges in complex automotive audio designs necessitates a combination of technical expertise, advanced equipment, and a willingness to experiment and fine-tune. Engineers must possess a deep understanding of signal processing theory and techniques, as well as a practical understanding of the vehicle's acoustics and electronics.


The most complex audio product is the car.



Multi-function audio system features are delivered as software on a Domain controller utilizing Audio Weaver framework.



Users can make routing changes within the Audio Weaver Designer canvas without redesigning underlying DSP.



Digital signal processors (DSPs) have emerged as powerful tools for tackling these challenges. DSPs allow for precise control over various signal processing parameters, including equalization, crossover frequencies, time alignment, and more. These devices provide the flexibility required to tailor the audio output to the specific acoustics of a vehicle's interior and the preferences of the listeners.

In addition to the hardware, sophisticated software applications and measurement tools play a crucial role. Real-time analyzers, acoustic measurement systems, and spectrum analyzers aid in objectively assessing the audio system's performance and identifying areas that require adjustment.

Providing all these audio and voice functions on a single domain controller requires significant coordination and time-consuming, individual integration efforts for the many software features. A flexible audio framework is essential. Audio Weaver, available from DSP Concepts, provides a flexible, easy-to-use platform for developing audio processing signal flows and incorporating a wide range of audio algorithms from third parties and brands and can support in-house algorithm development to create new user experiences. On a single design canvas, audio system designers can view and edit all the required audio functions using a graphical interface. Whether the functions require a single DSP core or

multiple cores available on modern automotive infotainment SOCs, full transparency of the design allows easy feature inclusion or exclusion depending on the audio requirements for the particular vehicle.

Early user experience concept development can be created on a PC environment using standard sound cards connected to microphones and audio amplifiers. Once the design is verified, the audio feature can be immediately deployed to the embedded target hardware and software environment supporting on-system validation.

Audio Weaver also provides embedded debugging tools to visualize and instrument the audio system to assist in tracking down system issues.

Modern complex automotive audio systems can now migrate from a hardware-focused design approach to a software environment that supports software-defined vehicle architectures and powerful domain controller SOCs. Audio features can be implemented easily using the Audio Weaver graphical design environment and deployed on most commonly used SOC, DSP, and ARM processors.

Conclusion

Developing complex automotive audio designs that offer an unparalleled auditory experience is a journey that requires a thorough grasp of signal processing challenges. The art of crafting sound in the confined and reflective environment of a vehicle interior demands the application of equalization, crossovers, time alignment, and more to achieve an optimal soundstage. Overcoming these challenges requires not only technical expertise but also the right tools and a commitment to refining and tuning the system to perfection. By skillfully navigating the intricacies of signal processing, designers can transform every drive into a symphonic journey, where the harmonious fusion of technology and creativity creates an auditory masterpiece.

www.dspconcepts.com

About the Author

John Whitecar is the VP of Product Management, Automotive at DSP Concepts where he works with world-class automakers to design cutting-edge automotive audio experiences. John'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. Before joining DSP Concepts, John was Senior Manager Wireless Engineering, Principal Engineer at Tesla.



MENLO SCIENTIFIC

Engineering
Success
since **1983**

Engineering Consulting for the Audio Industry

We Know

- Product Development

We Provide

- Design and Technical Support

We Deliver

- Liaison, Sourcing, and Guidance

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

MENLO
SCIENTIFIC

Richmond, California, USA

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



Distortion Measurement of OPA1656 and OPA1612

This article discusses measuring distortion of the extremely popular OPA1656 and OPA1612 Burr-Brown audio operational amplifiers and alternative THD measurement approaches.

By
Vojtěch Janásek




Figure 1: A bridge was used for the THD measurement of an inverting amplifier.

According to the manufacturer's datasheet, the OPA1656 has a THD below 160dB [1,2]. Measuring such low distortion is a real challenge because the best THD meters available, such as the APx555, have a THD of -130dB. However, Texas Instruments (TI) uses a clever trick described in section 8.3 of datasheet OPA1612—they increase the noise gain (typically 100x) of the amplifier under test. This change also increases the THD to a level that can be measured with a standard distortion meter. This approach can work but has an important assumption that is not stated in the datasheet—the amplifier input must be distortion free. This may be true for an inverting amplifier, but it is not true for a noninverting amplifier, where the source of distortion is finite and variable common mode rejection ratio (CMRR). Therefore, a different approach to THD measurement was developed and tested for the two amplifier circuits.

THD Measurement of an Inverting Amplifier

Figure 1 shows a bridge used for the THD measurement of an inverting amplifier. IC1 keeps the voltage at summing node B at virtual zero, so if $R1/R2 = (R3+R6)/(R4+R7)$, the voltage at node A is also zero, a small mismatch can be compensated by R7. This is true for an arbitrary voltage U_g and an ideal amplifier IC1. Using a real amplifier with finite gain, bandwidth and some distortion, the balancing is not perfect. The limited bandwidth causes some frequency dependent phase shift, so perfect balance is not possible, especially for the higher harmonics of U_g , but an attenuation of about 94dB can be achieved, see **Figure 2**.

The amplifier distortion can be modeled as an additional voltage source U_d . This additional voltage appears at node A divided by the ratio $(R4+R7)/(R3+R4+R6+R7)$, with the same resistors, divided in half. U_g itself and its distortion is greatly attenuated at node A, and the amplifier distortion is much more visible. As mentioned, the attenuation of the bridge is about 90dB so U_g with distortion lower than 80dB to 90dB is necessary.

The ultra-low distortion oscillator with distortion below -140dB was used here [3]. The sense amplifier with gain 500x consists of two stages—the first stage is OPA1656 with gain 50x, the second one is implemented with LME49710. The output of the amplifier was digitized using a 14-bit/400kS/s module, the digital output was averaged and down-sampled to 6.25 kHz, and then the fast Fourier transform (FFT) was calculated with 128kSamples.

The long-sample FFT is essential for obtaining very low noise floor—distortion of -160dB or less means that the amplitude of higher harmonics is tens of nanovolts (nV) and distinguishing such low-level signal from noise is a real challenge! Figure 2 shows the measured spectrum with R5 disconnected and an output voltage of $20V_{pp} = 7V_{RMS} = 17dBV$. The noise floor is -155dBV, the fundamental signal




Figure 2: The limited bandwidth causes some frequency dependent phase shift, so perfect balance is not possible, especially for the higher harmonics of U_g , but an attenuation of about 94dB can be achieved.

is -77dBV, and the higher harmonics are invisible in the noise floor, so they are below -149dBV. (Keep in mind the 6dB attenuation from the R3/R4 divider.) A direct measurement without R5 only shows that the higher harmonics are below -166dBc, so even in a bridge configuration it is necessary to increase the noise gain by adding some R5, a value of 820Ω was used here. Even with this resistor, the distortion remains invisible (i.e., the higher harmonics are below -175dBc). Total load $R1||R3 = 3.4\text{ k}\Omega$, with reduced Rload the distortion increased rapidly (see **Table 1**), but even with a Rload of 1.5kΩ, HD2 and HD3 are -161dBc and -164dBc, respectively. These values agree with the OPA1656 datasheet [1].

THD Measurement of a Noninverting Amplifier

A schematic of the distortion measurement of a noninverting amplifier is shown in **Figure 3**. Instead of measuring voltage directly at the output of the amplifier under test, the differential voltage between the input and output of the amplifier is measured by a differential amplifier with a gain of 500x. The distortion of this amplifier suffers from variable CMRR, so its distortion was measured first with the shorted inputs connected to the output of the source signal. There is a high probability that the distortion of the differential amplifier will be better than that of the follower itself due to the CMRR matching of the amplifiers on the same chip and the corresponding suppression of the nonlinearity.

The measured results are shown in **Table 2** and prove that this assumption is correct. The table shows that the distortion performance of the noninverting amplifier is significantly worse than that of the inverting amplifier and the increased noise gain approach is wrong here and gives incorrect results. Table 2 also shows that increasing the supply voltage reduces the distortion.

Another excellent low distortion amplifier is the OPA1612. Its main advantage against OPA1656 is its bipolar input stage with lower voltage noise, so this amplifier is better suited for amplification signals from sources with low output resistance. Distortion performance for both the inverting and noninverting configuration is shown in **Table 3**.

Conclusion

The distortion of inverting and noninverting amplifiers with OPA1656 and OPA1612 was measured. Measured data show that the method using higher noise gain gives correct results for an inverting amplifier but is incorrect for a noninverting amplifier because it ignores the effect of the variable CMRR, which is a significant source of distortion. 

	Uout	Rload	HD2	HD3
Inverter with no R5	7V _{rms}	3.4kΩ	≤166dBc	≤166dBc
Inverter with R5=820Ω	7V _{rms}	3.4kΩ	≤175dBc	≤175dBc
Inverter with R5=820Ω	7V _{rms}	2.3kΩ	-163dBc	-163dBc
Inverter with R5=820Ω	3.5V _{rms}	1.5kΩ	-161dBc	-164dBc

Table 1: This table shows the distortion of the OPA1656 inverting amplifier.




Figure 3: This is a schematic of the distortion measurement of a noninverting amplifier.

	Upower	Uout	HD2	HD3
Only differential amplifier	±15V	3.5V _{rms}	-162dBc	≤173dBc
Only differential amplifier	±17V	7V _{rms}	155dBc	-168dBc
Follower OPA1656	±15V	3.5V _{rms}	-155dBc	≤171dBc
Follower OPA1656	±15V	7V _{rms}	-148dBc	-159dBc
Follower OPA1656	±17V	7V _{rms}	-150dBc	-168dBc

Table 2: Increasing the supply voltage reduces the distortion.

	Uout	HD2	HD3
Inverter, Rload 3.4kΩ	3.5V _{rms}	<165dBc	<165dBc
Inverter, Rload 3.4kΩ	7V _{rms}	162dBc	165dBc
Inverter, Rload 1.5kΩ	3.5V _{rms}	157dBc	159dBc
Inverter, Rload 1.5kΩ	7V _{rms}	158dBc	159dBc
Follower, Upower ±15V	3.5V _{rms}	154dBc	173dBc
Follower, Upower ±18V	7V _{rms}	147dBc	164dBc

Table 3: This table shows the distortion performance for both inverting and noninverting configuration.

About the Author

Vojtěch Janásek has been designing and manufacturing precision analog and mixed-signal custom electronics for 30 years for various customers such as ASML, Framatome, and the Weizmann Institute of Science.



References

- [1] OPA1656 datasheet, Texas Instruments, www.ti.com/lit/ds/symlink/opa1655.pdf
- [2] OPA1612 datasheet, Texas Instruments, www.ti.com/lit/ds/symlink/opa1611.pdf
- [3] V. Janásek, "An Ultra-Low Distortion Oscillator with THD below -140dB," JanasCard, <http://janascard.cz/PDF/An%20ultra%20low%20distortion%20oscillator%20with%20THD%20below%20-140%20dB.pdf>

Tube Filament Supplies

Delve into the reasons to use DC filament supplies, advantages over using an AC heater supply, and alternative methods to eliminate induced hum.

By

Richard Honeycutt

From my teenage years, I read occasional mentions in various audio and electronics magazines of the use of DC for filament supplies, especially in preamp stages. Since the articles I read had been written (presumably) by experts, I just accepted the idea that hum could be reduced by using DC rather than AC for the filaments. Recently, I have become aware of several online discussions debating the merits and details of DC filament supplies, and I decided to delve more deeply into the matter.

Viewed simplistically, a filament carrying hundreds of milliamps (mA) of 50Hz or 60Hz AC, located very close to a control grid, would seem to be a good recipe for electrically (via capacitive action) or magnetically (via transformer action) coupled noise to mix with the desired signal at the control grid. But then we consider that some tubes use directly heated cathodes (the filament is the cathode) while others—probably most—use indirectly heated cathodes. The cathode is a metal sleeve surrounding the filament and insulated from it (**Figure 1**), so we can see that more analysis is needed.

Referring to the venerable *RCA Radiotron Designer's Handbook*, we find that the only relevant considerations mentioned concern making sure that the quiescent control grid bias is appropriately adjusted if a tube designed for a DC filament supply is used with AC, or vice-versa. For example, if a power triode designed for a 2.5V DC filament supply is used as designed, and the filament terminals are connected to ground and 2.5 VDC, respectively,

the average (midpoint) voltage on the filament will be 1.25V. If the filament is supplied with AC, the average voltage will be 0V. Clearly the control grid bias will have to be modified to maintain the same quiescent current in the tube.

Hum

But the AC vs. DC question typically is not centered around grid bias concerns. It has to do with hum. In years past, tubes used in auto radios often had directly heated cathodes because of heat and efficiency considerations, as well as the desire to minimize the current that the vibrator-type 12V-to-B+ converter had to supply. Power triodes such as the 2A3 and the 300B have directly heated cathodes, as did the small-signal triodes 01, 201, and so forth. The control grid signal voltage of a power triode is typically much larger than the hum voltage induced on the control grid from AC on the filament, minimizing any possible hum concerns.

Virtually all tubes used in modern audio preamps have indirectly heated cathodes, and the conductive cathode provides a sort of Faraday cage surrounding the filament. Thus, induced hum can appear on the cathode. But does this cause a problem? Given that a Faraday cage around the filament shields the control grid from electromagnetic fields, one's first answer might be "no." But remember that the filament is emitting electrons that can constitute a current from filament to cathode. With an AC filament supply, during the half-cycle when the filament is negative with respect to the cathode, diode action can create a hum signal on the cathode




Figure 1: Note the filament (heater) centered within the metal cathode sleeve.

at 120Hz and harmonics thereof. Since this signal is shielded from the grid, it appears as an AC voltage between the grid and cathode. Except....

As illustrated in **Figure 2**, most audio triode preamps use either grid-leak or cathode bias. The cathode of a grid-leak biased stage is grounded from both an AC and a DC standpoint. AC grounding of a cathode-biased stage involves the cathode bypass capacitor. With grid-leak bias, any AC leakage from filament to cathode is grounded out, but with cathode bias, the AC grounding depends upon the bypass capacitor. Since the AC leakage from the filament consists of the second harmonic of the power frequency, plus abundant higher even-order harmonics, the dielectric absorption and equivalent series resistance of the bypass capacitor must be low, in addition to the capacitance being large enough. (Typically, the reactance of the bypass capacitor should not exceed 10% of the value of the cathode resistor at the power frequency.) If the capacitor does not meet these conditions, a DC filament supply can reduce hum and buzz in a preamp stage.

AC leakage of the same sort can occur between the filament and the grid, even in tubes with indirectly heated cathodes. This can occur when portions of the filament are directly exposed to the grid, without the cathode intervening. Tubes using folded rather than coiled filaments are more susceptible to hum from this cause. In such cases, DC filament supplies will reduce hum.

In an online discussion of DC vs. AC filament supplies ("Vacuum tubes heaters supply, AC or DC?",




Figure 2: The most common bias methods for audio triodes are grid-leak bias and cathode bias.

see Resources), "user 16324" makes the following statement with regard to hollow-state guitar amplifiers:

Now with AC heating, there will be some small variation in cathode temperature across the AC cycle, and this modifies the emission characteristics. In a push-pull amplifier, everything cancels out under small signal conditions and the amp sounds quiet and clean, but under heavy load (where one valve is conducting hard and the other is off) this contributes some intermodulation products.

Hi-fi amplifiers minimise [sic] the audible effect using negative feedback, and modern ones can afford (especially at those prices!) to use regulated DC heater supplies; but that's not what you are designing : these imperfections are partly responsible for the grungy guitar amp sound.

So I would use an AC heater supply as a deliberate part of the design, especially if you are trying to achieve something like the original guitar amp sound.





Figure 3: The bottom power transformer winding feeds a rectifier and filter providing DC for the filaments.

Figure 4: The power supply of a Heathkit W5M power amp illustrates the use of a center-tapped filament winding to reduce hum.



The question naturally arises, “What examples can we find in which well-respected manufacturers of hollow-state audio equipment have used DC filament supplies?” While my knowledge is not comprehensive, at least one hi-fi preamp does so: the McIntosh C22, whose power-supply schematic is shown in **Figure 3**. Other manufacturers of the same era do not seem to have chosen this route.

Hollow-state mic preamps used in recording studios virtually always use DC filament supplies, but doing so is unnecessary because these devices are powered by 48V phantom power, so a switch-mode power supply (SMPS) is always used. Three

examples of current tube microphone preamps are the RMS5A7 from Roll Music Systems, the 583s from LaChapell Audio, and the Aphex Dual RPA 500.


An interesting discussion of DC filament supplies can be found on the diyAudio website (see Resources).

Low-Noise Design and Wiring Methods for Tubes with AC Filament Supplies

Tube filaments powered by AC in audio equipment should always be grounded in such a way that the electromagnetic fields produced by the AC are cancelled out. The preferred method, shown in **Figure 4**, uses a center-tapped filament winding. An alternative, shown in **Figure 5**, uses a “hum balance” potentiometer (pot) connected across the filament winding, with the pot’s slider grounded. For most audio amplifiers using 6.3V filaments, a 100Ω, 2W wire-wound pot is used. The pot can be adjusted by observing the amp’s output on an oscilloscope or listening to it through speakers while tweaking the pot for minimum hum. Either the center tap or the pot will reduce the filament-to-cathode AC voltage compared to grounding one end of the filament, thus lowering hum. In addition, either method prevents large chassis currents resulting from filament currents passing through the chassis. Such AC currents can cause very significant hum.

As mentioned above, using grid-leak biasing reduces hum from filament-to-cathode leakage, as does the use of sufficiently large cathode bypass capacitors. This latter option is not applicable in some cases: for example, in Williamson amps, negative feedback is often applied to the unbypassed cathode of the first preamp tube. In such cases, the

Figure 5: A hum balance potentiometer can be used instead of a center tap.



About the Author

Dr. Richard Honeycutt bought a *Popular Electronics* magazine from the rack at a bus station in 1960. It contained an article on building a transistorized audio amplifier, which captured his interest. He followed up by subscribing to several electronics magazines. To avoid the cost of parts, he built up a “junk box” by disassembling trade-in TVs from his uncle’s furniture store. His dad bought him the Van Valkenburg, Neville, and Nooger Basic Electronics book set, so he learned about tube circuits as well as transistors. He started repairing electronic devices about 1965, earning his First-Class Commercial FCC license in 1969. He has worked part-time in broadcast radio engineering and audio electronics repair since 1968, in addition to full-time work in acoustical and audio system design, plus a 20-year stint teaching electronics at the college level.



heater supply of the preamp tube can be isolated from ground and biased at about +15V DC to 50V DC, thus attracting leakage electrons back to the filament rather than to the cathode. It is important not to exceed the tube's maximum rated heater-cathode voltage (+100V average for a 12AX7).

DC Filament Supplies

At first glance, one might think that implementing a DC filament supply would be a no-brainer. However, if we simply rectify the AC from a 6.3V transformer winding, we'll get 6.3V minus the ~0.6V drop across the diode, plus a lot of noise at power frequencies and their even-order harmonics. So, we add a capacitor filter. If it's large enough to kill the ripple, it will charge to the peak value of the AC, $6.3V \times 1.414 = 8.9V$. This will be true for either a half-wave rectifier or a bridge. If we use a non-bridge full-wave rectifier, we'll get half that: 4.45V

In either case, we'll need a voltage regulator. IC 3-terminal regulators are very good at not only stabilizing the DC output voltage at the required value, but also at keeping their output voltages




Figure 6: An adjustable three-terminal regulator can be used to provide the correct DC filament voltage with low ripple.

clean and free from hum. **Figure 6** shows an LM317 adjustable three-terminal regulator circuit that can be used for this purpose. Resistor R2 should be adjusted for the proper filament voltage, commonly 6.3V. This means that the filaments of 12.6V tubes will be wired with one supply lead connected to the filament center tap and the other, to both ends of the filament. Most audio preamp tubes (12A_7, 5879, etc.) require no more than 0.3A filament current. An LM317 regulator in a TO-220 case can put out 1.5A DC. Maximum power dissipation is 15W. A capacitor-filtered bridge rectifier fed from a 6.3VAC filament winding will produce about 8.9V.

Since most power transformer windings output higher voltages than the nominal rating, we can assume an average DC input to the regulator of about 10% higher—maybe 9.8 V. This means the regulator will drop about $9.8V - 6.3V = 3.5V$. With

Unleash Your Brands' Sound Identity



Elevate your audio experience with Hypex Electronics **FusionAmps**. We offer the power to customize our audio solutions like never before...

*"It's not just about amplification;
it's about amplifying your brand."*

Tailored Perfection

- Choose your custom faceplate color to match your brand's aesthetic.
- Showcase your logo proudly, making every amplifier uniquely yours.
- Contact us for additional customization options.

Precision Sound, Great Performance

- Unlock the potential of FIR (Finite Impulse Response)
- Features that set the standard for quality and performance.
- 8 different models ranging from low power 1ch to high power 3ch solutions.
- The perfect solution for active speakers and subwoofers.

Make your mark in the world of audio with Hypex Electronics **FusionAmps**.

Contact us at sales@hypex.nl



Hypex Electronics B.V.
Kattegat 8
9723 JP Groningen
The Netherlands
www.hypex.nl

"Your brand, your sound, your signature."

300mA per tube, each tube will cause the regulator to dissipate about 1.05W. Since the maximum power output of the regulator depends upon the ambient temperature and the thermal resistance of the regulator and heat sink, these factors must be considered when one designs a DC filament supply using a three-terminal IC regulator. Normally, only one or two tubes need their filaments powered with DC, so often, a single LM317 with a reasonable heatsink will do the trick, assuming the regulator is not located in a very hot area of the chassis. To keep the instantaneous voltage at the regulator's input from ever going too low, we'll need to use a filter capacitor at the bridge's output. The minimum value of this capacitor is $C=i/228$, where i is the regulator's output current. For a single 6.3V, 0.3A filament, this value is about 1400 μ F.

Comparing the DC filament supply of the McIntosh C22 (shown in Figure 3) with an LM317-regulated supply, we find several differences:

- The C22's unregulated, RC-filtered supply produces 27.6V DC, implying a 9.5V AC transformer winding. The LM317 produces 6.3V DC, using a 6.3V AC transformer winding.
- The C22 provides DC for the filaments of six tubes: V through V6. The filaments are wired in series-parallel. The LM317 circuit is designed for a single tube's filament, or for parallel-connected filaments of multiple tubes.



Photo 1: Twisting filament leads together and routing them close to the chassis can reduce hum.

Resources

"Running tubes' filament on DC," diyAudio, September 2015, www.diyaudio.com/community/threads/running-tubes-filament-on-dc.280209

Switching power supply for tube filaments possible???" diyAudio, December 2007, www.diyaudio.com/community/threads/switching-power-supply-for-tube-filaments-possible.113972

"Vacuum tubes heaters supply, AC or DC?," Electrical Engineering Stack Exchange, June 2013, <https://electronics.stackexchange.com/questions/73459/vacuum-tubes-heaters-supply-ac-or-dc>

- The C22 uses a three-stage RC filter employing 3400 μ F of capacitance. The LM317 can provide more than 80dB of ripple rejection if the ADJ terminal is bypassed to ground with a 10 μ F capacitor. As discussed above, a single filter capacitor of about 1400 μ F is required.
- The C22 circuit has a calculated ripple about 40dB below the DC voltage, whereas the LM317 circuit's ripple is over 80dB down.

Another method of reducing heater-to-cathode leakage without using a DC filament supply is to reduce the filament voltage of the critical preamp tubes by about 10%. This approach also extends tube life. However, reducing the filament voltage of power tubes or other tubes being operated at a substantial plate current (such as reverb tank drivers in guitar amps) is not recommended, as it can increase distortion or create other undesirable effects.

One way to obtain good voltage regulation without having to pay so much attention to heat dissipation is to use a SMPS for the filaments. However, these tend to have lots of AC noise in their outputs. In fact, I have had to use an SMPS as a pre-regulator, followed by a linear (three-terminal IC) regulator to filter out the horrid buzz from the SMPS. Using ferrite beads and a shunt ceramic capacitor near the tube's filament terminals can help, but it is very difficult to get the output from an SMPS to approach the cleanness of a linear regulator. Besides the noise issue, the ultrasonic noise from an SMPS can intermodulate with audio frequencies, causing a nasty sound quality. So, I avoid the SMPS approach for audio circuits whenever possible. More information can be found at on the diyAudio website (see Resources).

The Bottom Line

For low hum in a hollow-state audio amplifier, all AC filament leads should be twisted so that the AC electromagnetic fields from the two wires will cancel. Filament leads should be routed as far as possible from low-level signal leads and should be located as close to the chassis as possible (**Photo 1**).

If these guidelines are followed, a DC filament supply is unlikely to be necessary. As mentioned earlier, lowering the heater voltage on the first preamp tubes by 10% will reduce hum. This can be done by inserting a 2 Ω , 1/2W resistor in series with the filament of the first preamp tube. If a DC filament supply is found to be necessary, the easiest and best-performing design is probably to use an LM317 three-terminal IC regulator as previously discussed.

J SERIES

THE LEGENDARY J SERIES IS BACK!



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

SHIELDED AUDIO IMPEDANCE MATCHING
DESIGN FOR A WARM SIGNAL.

SINGLE-HOLE MOUNTING ALLOWS ROTATION
FOR MAXIMUM HUM REDUCTION.

VINTAGE CIRCUITRY, PREMIUM
MATERIALS, HAND-CRAFTED QUALITY.

VISIT ONLINECOMPONENTS.COM TO
ORDER OR LEARN MORE.

MADE IN THE U.S.A.



onlinecomponents.com
THE ONLINE DISTRIBUTOR OF ELECTRONIC COMPONENTS

Dayton Audio BST-300EX **Extreme High-Power** with **6X** the Watts!

Dayton Audio BST-300EX Extreme High Power Pro Tactile Bass Shaker 300 Watts

The BST-300EX is a powerful bass shaker designed for serious audio enthusiasts. With an impressive 300W of power, it delivers extreme bass that will immerse you in your favorite movies, games, and music. The inaudible operation ensures a seamless experience, while the easy mounting makes installation a breeze.

#295-243



Dayton Audio BST-1 High Power Pro Tactile Bass Shaker 50 Watts

With 50 watts of power at 4 ohms, the Dayton Audio BST-1 Bass Shaker serves as an efficient exciter speaker, delivering an immersive tactile bass experience.

#295-244



Feel the Tremors!



DAYTONAUDIO
daytonaudio.com

Distributed by:

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



Scan for more
Information!
Dayton Audio
BST-300



With a Dayton Audio 5 year warranty, it's more than a guarantee, it's a promise. We stand behind our products and our warranty is a demonstration of our commitment to quality and reliability.