

audioxpress

ADVANCING THE EVOLUTION OF AUDIO TECHNOLOGY

Market Update—DSP Products and Applications

Real-Time DSP, Edge Processing, and Artificial Intelligence

By J. Martins and Mike Klasco

The Evolution of Audio DSPs

By Youval Nachum (CEVA)

Communications

Testing and Optimizing Strategies for Conference Terminal Equipment

By Hans W. Gierlich (HEAD acoustics)

Fresh From the Bench

Earthworks M23R Measurement Microphone

By Stuart Yaniger

Practical Test & Measurement

What to Measure for Quality Control of Audio Devices

By Dan Knighten (Audio Precision)

Audio Scope

Simple Improvements to Legacy Audio DAC Output Interfaces Using Op-Amps

By Michael Steffes

Sound Control

Time-Domain and Distortion Testing

By Richard Honeycutt

You Can DIY!

Ultra-Low Distortion Microphone Buffer Amplifier (Part 2)

By Michel Nieuwenhuizen

R&D Stories


Getting Started with A²B System Design

By Brewster LaMacchia



AES70

network-controlled
DSP module



Resolute's **RA-ES70DSP04 DSP module** is easily programmable, so you can add AES70 networked control and monitoring to your next generation product, be that a powered speaker, microphone pre-amp, power amp or small mixer.

The algorithms available for the RA-ES70DSP04 include multiband EQ, Crossover Filters, Vocal Reverb, Compressor, Noise Gate, Spring Reverb and many more. An onboard web server provides rich controls for each algorithm.



You'll find all Resolute products and other audio parts at profusionplc.com

DSP Design Tools

The RA-ES70DSP04 comes with all the software tools to help you program the DSP and implement AES70. In addition, there is an evaluation board that includes 4 channels of analog audio in and out, and 2 channels of AES67 audio in and out, ideal for developing networked audio products.



1140 Series

Studio-Grade Impedance Matching Transformers

Rooted In Solid ENGINEERING

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

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





November 2020 ISSN 1548-0628

www.audioxpress.com

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

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

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

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

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

Advertising rates and terms available on request.

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

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

© KCK Media Corp. 2020
Printed in the US

The Multiple Facets of Audio Signal Processing

In a 2014 presentation at the 137th Audio Engineering Society (AES) convention, Paul Beckmann, the founder of DSP Concepts, revealed the results of a search within his LinkedIn network to determine how many people had “Audio Engineer,” “Embedded Software,” and “Audio DSP” listed in their profiles. He made a Venn diagram of the results, which showed that very few people included “Audio DSP” and even less (11 out of a universe of more than 100,000), actually claimed to do all of these three things. To illustrate the complexities and challenges of current audio product development and particularly in making an embedded audio product, Beckmann also identified skill sets in electrical, mechanical, acoustical, and user interface software development as basic requirements.

It is not surprising that digital signal processing (DSP) continues to be one of the best kept secrets in the audio industry. The people who actually push DSP forward are not many, while the required “math and numbers” demand a lot of intense work to translate into code by a whole lot more people, who simply aren’t there.

A growing number of companies in the audio DSP field deserve recognition, even if sometimes they don’t necessarily want to be in the limelight. Mostly, because their work is done behind the scenes in the industry, mainly focused on licensing IP. For that reason, it’s only natural that they don’t exactly publish tutorials explaining how they do what they do.

In professional audio and broadcast applications, leading companies evolved from the days of analog audio, later embracing digital signal processing, while others started with DSP. Still, many rely on selling hardware to package their processing tools, avoiding getting into too much detail about how things work—just offering a few basic knobs for the user to tweak to get the best result.

In contrast, software companies are extremely more visible, currently offering some of the most sophisticated audio processing ever available: removing background noises, replacing notes and whole polyphonic sequences, or doing sound source separation within existing recorded tracks. Of course, those companies don’t detail how their algorithms work, and they focus on explaining the user interface for their application or plug-in. Yet, that’s where we find today’s most vibrant community of audio developers, particularly since the availability of mobile platforms and the boom of the app stores. But again, not many are actually involved with DSP. Most understand how to implement DSP algorithms as long as they are offered in ready-to-use “blocks.”

Increasingly, on all product development fronts, companies are combining DSP algorithms with artificial intelligence (AI). That’s the vanguard of digital signal processing these days and lots of very clever things are already made possible because of AI.

Room correction is becoming a commodity in integrated home audio amplifiers, even if combining calculations performed remotely, in the cloud, and stored for active processing in the local device. Same with headphones and head-related transfer functions (HRTF) used for personalization. But soon, all that digital signal processing and AI will be able to happen directly on-device. Even in portable and wearable devices.

Our current mobile platforms are now becoming capable of running those extremely intense computational tasks for real time processing. That’s precisely why consumers are now able to benefit from digital signal processing being used so extensively. And thanks to AI, users no longer need to learn control interfaces for things they don’t need to know are running in the background.

Still, no matter how powerful the processing and number of sensors, more and more, we need “audio engineers” who are able to understand how signal processing algorithms benefit our experiences, and provide algorithms to feed all that AI.

With voice interfaces and voice assistants taking center stage across multiple industries, the demand for DSP coders and anyone with real knowledge in application development will expand even more.

Finding companies that not only deliver the solutions but also have the experience required to build hardware and software solutions for audio applications—professional and consumer—and are able to support custom development in audio processing and programming is not an easy task but is more requested than ever. In fact, that is why many brands settle for me-too designs, and are not able to create differentiated products.

J. Martins
Editor-in-Chief

The Team

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

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



OUR NETWORK



SUPPORTING COMPANIES

ACO Pacific, Inc.	61	Linear Integrated Systems	41
Acoustic Sciences Corp.	49	Menlo Scientific, Ltd.	37
All Electronics Corp.	59	Oxford Digital, Ltd.	25
Ambiq Micro, Inc.	55	Parts Express International, Inc.	68
Audio Precision, Inc.	29	Profusion plc	2
Clockworks Signal Processing, LLC	45	Purifi Audio	33
COMSOL, Inc.	9	Redco Audio	61
DSP Concepts	23	Syntiant	27
Hammond Manufacturing, Ltd.	3	Triad Magnetics	67
Hypex Electronics B.V.	53	Warkwyn	45
KAB Electro-Acoustics	59		

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.

COLUMNISTS

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

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

Mike Klasco is the president of Menlo Scientific, a consulting firm for the loudspeaker industry, located in Richmond, CA. He is a graduate from New York University, with post graduate work in signal processing, and he holds multiple patents licensed or assigned. For the past 35 years, Mike has worked on countless R&D projects for large and small companies. Mike specializes in materials and fabrication techniques to enhance audio performance.

Michael Steffes has worked in high-speed signal path IC development, spanning IC design roles for the BurrBrown high-speed amplifier products, and product launch/application manager roles through Texas Instruments/Intersil. He continues to focus on high-performance ADC interface solutions, DSL/PLC line driver design and optimization, and signal path design optimization tools. In a career spanning five different major IC vendors, he has participated in the introduction of 150 industry leading amplifier products (setting many of the now standard practices in high speed amplifier characterization curves), publishing more than 100 contributed articles and application notes from 1989 onward.



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



● Features

38 What to Measure for Quality Control of Audio Devices

By Dan Knighten

Audio Precision's Dan Knighten focuses on the practical question of what to measure for the quality control of audio devices, with an emphasis on manufacturing test strategies and objectives. He also shares what measurements reveal about the device under test, offers suggestions for specific measurements, and provides techniques for various devices and scenarios.

42 Testing and Optimizing Strategies for Conference Terminal Equipment

By Hans W. Gierlich

Hans W. Gierlich, of HEAD acoustics, provides some insight into how to evaluate and determine strategies to deal with speech transmission quality, which is especially relevant with most company employees working from home and the trend toward decentralized organizations, remote collaborations, and conferencing services.

50 Earthworks M23R Measurement Microphone

By Stuart Yaniger

Stuart Yaniger reviews the M23R Earthworks measurement microphone, which is the latest omnidirectional model in Earthworks' infinitely matched series of measurement microphones.

56 Getting Started with Automotive Audio Bus (Part 2) Understanding A2B Features



By Brewster LaMacchia

In Part 2 of this article series Brewster LaMacchia examines how certain optimizations for the automotive applications must be accounted for when targeting broader audio applications. He also includes a look at analog phantom power and how A2B provides that feature in a digital system.

62 Ultra-Low Distortion Microphone Buffer Amplifier (Part 2) Simulations and Measurements

By Michel Nieuwenhuizen

In the second part of his article, Michel Nieuwenhuizen shows the results of the simulations and the measurements. He also compares the performance of his own circuit with that found in many Chinese condenser microphones (the "Schoeps" based circuit).



● Market Update—DSP Products and Applications

8 Real-Time DSP, Edge Processing, and Artificial Intelligence

By J. Martins and Mike Klasco

This year's digital signal processing market update highlights solutions and platforms that are now able to run extremely intense computational requirements in real time, as required by consumer applications (e.g., voice detection and acoustics optimizations). Increasingly, these platforms are evolving in large scale, targeting the ability to run artificial intelligence inference directly on-device.

24 The Evolution of Audio DSPs

By Youval Nachum

CEVA's Senior Product Marketing Manager, Youval Nachum, discusses how DSP technology has evolved, its impact on the user experience, and what the future of DSP has in store for us.

● Departments

4 From the Editor's Desk

5 Client Index

● Websites

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

● Columns

SOUND CONTROL

28 Electroacoustical Testing Time-Domain and Distortion

By Richard Honeycutt

AUDIO SCOPE

34 Simple Improvements to Legacy Audio DAC Output Interfaces Using Op-Amps

By Michael Steffes



@audioxp_editor



audioxpresscommunity




linkedin.com/company/audioxpress



Real-Time DSP, Edge Processing, and Artificial Intelligence

This year's digital signal processing market update highlights solutions and platforms that are now able to run extremely intense computational requirements in real time, as required by consumer applications such as voice detection and acoustics optimizations. Increasingly, these platforms are evolving in large scale, targeting the ability to run artificial intelligence inference directly on-device.

By
**J. Martins and
Mike Klasco**



Developing real-time embedded audio systems involves a combination of software development and porting onto one or more hardware components, such as System-on-Chips (SoCs), microcontrollers, and specialized audio chips. It's a complex process that requires an audio development platform with several mission-critical features.

With audio strongly at the center of all the leading consumer electronics product categories, from smartphones to smart speakers and wearables—where voice interfaces and voice assistants are taking center stage—the audio industry is trying to balance embracing the latest, more efficient technologies, and keeping up with the need to differentiate their designs. And differentiated products, for most manufacturers, start on DSP implementation level, which in turn depends upon development tools and platforms. And that's the reason why most audio brands become dependent of their DSP partners and respective eco-systems.


For the reverse equation, companies supplying DSP services also struggle to match all the industry

requirements to port their software to different hardware platforms and fit the design and budget restrictions of manufacturer's systems. As a result, a frequent option is to simply allow other companies and engineering consultants to use their DSP technology, software, and tuning tools—which in turn makes it all the more difficult to recognize the actual technology providers.

DSP Blocks

Signal processing permeates an astonishing range of devices with which we interact. DSP is now heavily featured in the engine-revving sounds from the Porsche Taycan, despite the fact that it's an electric vehicle. And, in automotive, it is also showing up in everything from the signal processing audio systems

*Simulation + testing =
optimized loudspeaker designs*



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

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

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

comsol.blog/loudspeaker-design

and active noise canceling within vehicle interiors to driver awareness warning alerts.

Other signal processing tasks range from cleaning up noisy signals, suppressing echo with hands-free natural (full-duplex) conversations, and smart products that dynamically optimize to continuously do a better job as conditions change.

But for troubleshooting subtle audio and tuning issues and their interactions, some of us have learned that signal processing development requires a step-by-step approach to debug hardware and software issues at all iterations to avoid catastrophic production delays. While in the throes of an endless algorithm porting struggle, you may have heard a fellow engineer sarcastically mutter “SMS” (which means “small matter of software”).

Early projects with ambitious digital processing were plagued with high development costs, long and unpredictable time to market, lack of a cohesive design path, and endless remedial repetitive design efforts in general. To change the complexity and challenges of end-to-end signal processing, DSP vendors gradually provided developers with signal processing development tools that make it possible to actually listen to the results before building any prototypes, and respond in real time to device and environmental changes in the finished product.

Ideally, developers should be able to work and perfect their audio algorithms on a standard PC, before moving to an evaluation platform. That work stage requires development, measuring and

tuning, and optimizing features in real time, before generating any code.

Approximately 20 years ago, Analog Devices introduced the foundations for its SigmaStudio, essentially an operating system shell graphical tool that did not require coding. With some parallels to the breakthrough of Windows from DOS commands, these tools expanded software development for Analog Device’s SigmaDSP audio digital signal processor range.

Similarly, Texas Instruments (TI) developed a range of graphical development tools (including PurePath Studio, etc.), offering an easy-to-use environment to simplify software development for audio products. Programming is done graphically using a library of low level and high level audio processing components, with a vast library containing functions such as digital volume control, mixers, equalization, tone controls, loudness, dynamic range compression, bass boost, and surround effects. And from libraries supporting professional audio development, TI today offers full support for voice processing on DSPs, with algorithms that work with multiple microphone configurations to enable voice recognition in voice assistant applications.

Founded in 2003, DSP Concepts changed the equation with its Audio Weaver graphic environment launched in 2006. During its first nine years, Audio Weaver mostly remained an internal tool used by DSP Concepts to accelerate its consulting business. Finally, in 2015 the company productized it. Like other existing graphic development tools, Audio Weaver can be used at the research and development stages, and for producing optimized code for a target processor. The same tool can even be used to tune up the algorithms at the final stages of production. And across the process, there are regression tools to ensure that any change is accurately reflected and will not mess up previous codes.

But while Analog Devices’ or TI’s tools were created to support each company’s own DSP chips, DSP Concepts offered support for a large number of platforms. And existing designs created with Audio Weaver can even be easily ported to another target processor, if and when required. For an audio product company, not being tied to just one kind of target processor is a huge advantage.

Also, the no-code approach implemented in Audio Weaver paved the way to an improved development flow, with engineers specializing in audio tuning focusing on the algorithms while other members of the team with a deep understanding of the processing requirements and optimizations for real-time audio processing can supervise the system design.

Plus, Audio Weaver offers advanced audio DSP algorithms, with optimized Fast Fourier Transforms

(FFTs), Biquads, Butterworth, FIRS, sample rate converters, and custom modules for all that’s needed for audio designs. And the whole library of module “blocks” are optimized for the ARM Cortex-A, ARM Cortex-M, as well as the traditional DSPs (e.g., Analog Devices’ Sharc and Blackfin). This reduces the typical audio DSP software design process by 90%, compared with native DSP coding.

Platforms and Differentiation


As we discussed in our DSP Market Update last year, the reason why digital signal processing coding is increasingly one of the most requested skills by employers in the audio industry, is because these applications—which previously were in the realm of science fiction or required the most powerful mainframe computers to process—are now available in a variety of platforms. DSP is now used in portable professional tools, in our budget smart TVs and home audio amplifiers, in our pocket devices, and even in our ears—with smartphones and true wireless stereo (TWS) earbuds now leading the charge.

Whether assisted by remote computation (in the cloud) from connected devices, or performed locally on-device, these new audio products must be able to apply the most effective and sophisticated forms of DSP in real time when needed—increasingly in ultra low-power platforms.

Signal processing functions of filters, crossovers, noise and echo removers, equalization and dynamic processing, limiting, and so forth are all common features to be processed by edge devices (the extreme end of a network of connected devices). However, processing can now respond in real time to actual variations in the working environment (e.g., acoustic conditions), and change according to the actual user application. That, obviously also requires an important change in the platform’s ability to deal with variable conditions. Fortunately, the economies of scale for sophisticated mainstream products (e.g., smartphones and wearables) are allowing every other class of product to benefit.

While professional audio DSP applications typically run in dedicated hardware using application-specific integrated circuits (ASICs) and field-programmable gate array (FPGA) platforms, after a transition where companies explored Intel x86 CPU platforms and personal computers in general, the industry focus is increasingly centered on ARM processors, because of the large eco-system of resources that are available to developers.


Analog Devices’ Super Harvard Architecture (SHARC) processors also continue to be a popular platform for audio manufacturers, because their SigmaDSP tools don’t require coding and provide




Running on a standard Windows PC, Texas Instruments’ PurePath Studio graphical development suite made audio system design and development much easier.

a high integration level, with full audio System-on-Chips (SoCs).

And in general, the combination of ever-more powerful FPGA solutions from Altera, Xilinx, and Microsemi/Microchip are also a powerful ecosystem for DSP development, considering the many available development tools and that FPGAs are now frequently combined in SoC and System-on-Module (SOM)-types of solutions, with ARM-based and RISC-V processors. That combination also offers all the advantages of simplifying the system’s interface and peripheral implementations.



The Sonnox Elite plug-ins show how pioneering work on the development of digital signal processing, targeting studio and professional audio applications, inspired practically all the exciting developments behind the power of consumer electronics. Sonnox Oxford plug-ins trace their heritage back to the 1980s, to the development of the legendary Oxford OXF-R3 digital mixing console. As music production began shifting toward digital workstations, the console inspired a series of unique software developments starting with the Oxford EQ and Dynamics plug-ins. The original Oxford Digital company continues to offer those original developments for direct integration in DSP platforms.



The SigmaStudio graphic development platform was designed to support DSP programming for Analog Devices’ chips.

The argument of audio quality also supported the adoption of ARM-based solutions. While the big audio manufacturers quickly realized the advantages of replacing traditional low-cost DSPs with ARM Cortex CPUs, partly because of the added connectivity options, most manufacturers remained committed to using low-cost ASICs with basic audio functions and codec support on low-cost audio products.

Having an easily programmable platform enabled manufacturers to have better sounding solutions and have differentiated features, and the industry gradually migrated in that direction, largely also motivated by the massive economies of scale enabled by the mobile industry. The transition from cheap basic MP3 players (manufactured in the tens of thousands) to sophisticated smartphones (manufactured in the millions) made it possible to have a niche market such as dedicated portable audio players with high-resolution audio support and high-quality analog stages using the same basic “computers” inside.

Even TI, a long-time leading vendor of digital signal processors with a vast portfolio of programmable DSP devices, migrated to support ARM-based systems, because of the architecture’s ability for embedded real-time signal processing applications. Currently, TI’s DSPs range from low-

power single-core processors to multicore and DSP+Arm SoCs supporting a wide, scalable range of signal processing needs.

The Power of RISC

A few years ago, it was common for the industry to believe that general-purpose processors found in personal computers (e.g., those from Intel and AMD), would continue to lead in scalability and evolution cycles. But the truth is that ARM-based solutions are now superseding the x86 CPU architecture and everyone believes that other reduced instruction set computer (RISC)-based platforms will also continue to evolve at a faster pace. And the main reason for that transition has to do with the “performance per watt” logic, which was a requirement for laptops and notebook computers and is now the number one priority in low-power mobile devices.

When it comes to power requirements, cost/performance ratio, and floating-point operations essential for DSP applications, RISC platforms continue to lead, particularly for real-time signal processing and leveraging a real-time operating system (RTOS) in timing-critical applications like audio. ARM processors now offer efficient high-level language programming tools and compilers, as well as friendly graphic-based dedicated tools for audio DSP. A notable example being the already mentioned Audio Weaver from DSP Concepts.

Audio Weaver has in fact expanded in popularity because it runs in popular ARM Cortex development boards, including those from STMicroelectronics and Atmel and even affordable open-hardware solutions (e.g., the BeagleBone Black or Raspberry Pi).

DSP Concepts also works with several chip partners to provide efficient audio product deployment, including for voice recognition applications, using different hardware platforms from MediaTek, Qualcomm, TI, NXP and others—all featuring ARM Cortex CPU cores with both 32-bit and 64-bit support, combined with their own application processors and modules.

Due to the large scope of applications and platforms supported, DSP Concepts is in a unique position to offer an end-to-end solution that addresses all the audio processing challenges, including voice-activated products—making it possible for manufacturers to create, integrate, and deploy custom audio features for consumer electronics and automotive applications, using both low-code/no-code designers and experienced engineers.

Chin Beckmann, co-founder and CEO of DSP Concepts, offers an interesting perspective on this when she states: “The development process and ability to innovate and iterate with audio is stuck in the 1990s,

but audio experiences have never been a more critical component of every technology product sold in the market today. The Audio Weaver platform is real time and takes care of optimization across multiple chip families so customers can focus on building product-differentiating features rather than low-level coding. With a unified audio engineering solution, there are no barriers to the development of any unique audio product, from design to deployment.”

To which, Paul Beckmann, DSP Concepts’ co-founder and CTO adds: “We like to think of Audio Weaver as building blocks for audio: We provide over 400 of these audio ‘Legos’ that include modules to support everything from hyper-efficient IoT devices to ultra-premium automotive systems.”

Another example of a popular RISC-based DSP platform is the Cadence Tensilica HiFi DSP series, which resulted from the 2013 acquisition of Tensilica by Cadence Design Systems. Tensilica, a company founded in 1997 by Chris Rowen (also one of the founders of MIPS Computer Systems) together with former MIPS engineers, developed the HiFi audio/voice DSPs with a software library of more than 225 codecs from Cadence and more than 100 software partners. This was an extension development from their Xtensa processors, which created the foundation for all Cadence DSPs. The Xtensa 32-bit architecture offers a user-customizable instruction set with a base of RISC instructions. The Cadence Tensilica HiFi audio and voice DSPs have been licensed by companies as diverse as AMD, Microsoft, VIA Technologies, Realtek, and Ambiq Micro.


Audio Weaver’s agile graphical design environment from DSP Concepts is also available and optimized for Cadence Tensilica HiFi Audio/Voice DSPs.

CEVA and RISC-V


Another example of a leading vendor in DSP is CEVA, which was created in 2002 to be the DSP IP licensing division of two leading chip companies: DSP Group (Israel) and Parthus (Ireland). The majority of CEVA’s SoCs combine a RISC microprocessor and DSP processor core in the same chip creating a complex design, development, and verification process. To reduce the risk, cost, and complexity of these systems, in 2003, CEVA and ARM worked together on a joint specification standard for a hybrid RISC/DSP based SoC, which covers the hardware interfaces between the cores and software APIs for inter-processor communications. The specification also includes circuitry for low-latency message passing and bulk data transfer integrated with powerful DMA capabilities. This approach created a standard integration between the two leading RISC and DSP architectures, reinforcing CEVA’s position, particularly for mobile applications.

Already a leading IP licensor of DSP technologies, in 2014, CEVA acquired RivieraWaves, a French company with a unique portfolio of Bluetooth and Wi-Fi technologies, enabling CEVA to quickly become a leading vendor for wireless audio products. With the RivieraWaves family, CEVA built a leading portfolio of Bluetooth and Wi-Fi SoCs and application-specific standard product (ASSP) integrated circuits, where the protocol software is able to execute on the manufacturer’s choice of embedded processor.

In recent times, CEVA is increasingly leveraging the RISC-V open-instruction set architecture, originally developed at the University of California, Berkeley, and based on RISC, just like ARM and MIPS. With the support of the open-source community, RISC-V evolved




An example of a real-time software plug-in from Waves Audio, which serves to show the latest trends of replacing audio signal processing interfaces with automatic (AI-enhanced) tools. Part of the Waves Content Creator Audio Toolkit software bundle for podcast and live-stream, the new Playlist Rider is a set-and-forget volume management plug-in that automatically adjusts the volume of multiple sources (host, guest, music playlist, and audio clips) to a selected volume target, automatically and in real time.



The NoiseAssist algorithm automatically and continuously monitors the background noise independently from the desired speech in the signal and runs directly on a Sound Devices location recorder. This enables user to record a clean voice signal even in the direct presence of a fan or noisy HVAC system. The algorithm is the result of two years of in-house research and development by Sound Devices’ Steve Popovich and CEO/President Matt Anderson, who collectively hold more than 20 patents and have several decades of experience in this domain.

Audio Weaver is a graphical configuration tool that enables the simple and rapid development of audio/voice DSP, eliminating the need for native coding and kernel optimization.



quickly, and CEVA is one of the companies determined to take advantage of the architecture.

In 2019, CEVA's RivieraWaves platforms became the first to be integrated with a RISC-V sub-system, including a comprehensive set of peripheral IP. CEVA has also been exploring the potential of RISC-V for applications that require a small processor and low-power operation, with designs that can be easily deployed in FPGA and in ASIC/ASSP systems.

And to further explore the RISC-V potential, in 2020, CEVA started a joint silicon development with SiFive (the leading provider of commercial RISC-V processor IP) to develop SoCs for artificial intelligence applications in edge-devices, responding to the growing demand from the smart home, automotive, and wearables markets—all high-volume

markets. The partnership is centered around RISC-V CPUs, CEVA's DSP cores, AI processors, and software, which will be designed into SoCs for the mentioned applications. As CEVA highlights, "Domain-specific SoCs that can handle machine learning processing on-device are set to become mainstream, as the processing workloads of devices increasingly includes a mix of traditional software and efficient deep neural networks to maximize performance, battery life, and to add new intelligent features. Cloud-based AI inference is not suitable for many of these devices due to security, privacy, and latency concerns. SiFive and CEVA are directly addressing these challenges through the development of a range of domain-specific scalable edge AI processor designs, with the optimal balance of processing, power efficiency, and cost."

And we have no doubt that this partnership with SiFive will enable innovations in machine learning with faster time-to-market. According to Dr. Naveed Sherwani, president and CEO of SiFive, "The rapid evolution of AI models combined with the requirements for low power, low latency, and high performance demand a flexible and scalable approach to IP and SoC design that our joint CEVA/SiFive portfolio is superbly positioned to provide. The result is shorter time-to-market, while lowering the entry barriers for device manufacturers to create powerful, differentiated products."

Multi-Core Development

To better understand how new embedded hardware platforms are totally changing the equation for signal processing in general, and very specifically for audio, we need to highlight DSP Concepts' recent announcement of AWE Core OS, a new variant of

the audio-processing engine powering its popular Audio Weaver platform—now optimized for multi-core development. This dramatically reduces the costs and components typically required to build new audio features and is particularly relevant for demanding applications (e.g., automotive brands) to address the complexities associated with creating new audio experiences in that space.

AWE Core OS gives audio engineers the freedom to move audio processing to the main SoC without worrying about running out of CPU on a single core. In the automotive industry, developers are already pushing the limits of quad and octo-core development to improve the in-cabin experience with new audio and voice components. This multicore optimization perfectly suits the ARM Cortex family of processors for both delayed and real-time applications.


Audio development in the automotive industry is not only incredibly complex, but it's also vital for product differentiation—especially as we move toward more electric models and widespread adoption of voice input. Audio Weaver is also already embedded in leading automotive audio providers Panasonic, Sony, Alpine, and more, and AWE Core OS enables them to move with next-generation in-cabin experiences like "Individualized Sound Zones," using directional audio to customize and localize sound.

AWE Core OS supports designs requiring multiple threads or various sample rates on the same SoC. For instance, a low-latency telephony process can run concurrently with a normal latency entertainment post processing process. In addition, AWE Core OS also supports distributed audio processing loads across multiple SMP cores.

Edge Processing and Learning Platforms

But the real motivation behind the adoptions of extremely powerful platforms that are highly beneficial for on-device audio processing, is also the strong industry push toward edge-intelligence. Artificial intelligence and neural networks (learning engines) are vital to enable not only voice commands and interactions, but also for contextual awareness and audio enhancement workloads.

In an article titled "Dedicated Audio Processors at the Edge are the Future. Here are the Reasons Why," John Young (then working for Knowles, and now Director of Product Marketing at DSP Concepts) describes "...as audio is finally advancing as a key interface for human-to-machine and even machine-to-machine communication and interactions, a comparable evolution is taking place where DSPs are being optimized to do more than just traditional audio processing. For example, by augmenting audio




Custom built to run neural workloads, the Syntiant NDP100 processor consumes less than 150 μ W, enabling devices to be approximately 200x more efficient, and provides 20x the throughput of the typical low-power MCU solution.


DSPs with the latest machine learning circuitry, they provide more value than just improving audio quality. They bring artificial intelligence (AI) to consumer devices, making them smarter and more powerful at the edge, without needing a network connection. The implications are powerful."

In his article, Young explains how "audio edge processors" are becoming critical requirements for designing "the next generation of soundbars, smart home devices, smart speakers, smart TVs, and a whole host of intelligent IoT devices."

The new "specialized audio edge processor with an audio-specific instruction set, machine learning



Syntiant has been named a CES 2020 Best of Innovation Awards Honoree for its Syntiant NDP100 microwatt-power Neural Decision Processor, which provides always-on AI processing for voice and other sensor applications in battery-powered devices.



DSP Concepts introduced AWE Core OS, which allows embedded audio features to be integrated, deployed, and managed with extreme efficiency with support for multiple cores or running multiple threads.



CEVA's DSP cores and WhisPro speech recognition software now also support TensorFlow Lite for Microcontrollers in edge devices.

optimizations, and a focus on the audio fabric," as Young describes, helps deal with latency, application processor overheat, and power consumption budgets issues. "These processors can deliver enough compute power to process complex audio with AI and machine-learning algorithms at a fraction of the energy of an application processor or generic DSP implementation. This results in high-quality audio performance and intelligent functions for smart devices such as the ability to identify or locate the person speaking or cancel out noise interference."

Basically all companies leading the way in DSP audio processing are now strongly committed to support those "intelligent functions" without sending out information for remote processing.

This is also motivating new players in the space, such as Syntiant, a technology company leading in AI voice solutions at the edge, which has shipped more than 1 million units of its Syntiant NDP100 and Syntiant NDP101 microwatt-power Neural Decision Processors (NDPs) in a short period of time.

The Irvine, CA-based company successfully completed new funding rounds, and continues to execute on its strategy—fulfilling growing demand for its "ultra-low-power, deep learning voice solution at the edge."

Syntiant's processors enable customized voice control interfaces implementation on-device, across multiple product and use cases, enabling wake words, command words, speaker identification, and event detection, all free from cloud connectivity, ensuring

privacy and security. Custom built to run neural workloads, the NDP100 and the NDP101 chips consume less than 140 μ W, while processing audio events.

Currently implemented in a wide range of battery-powered devices, from smartphones, to smart speakers and earbuds, and built from the ground up for these demanding applications, Syntiant's NDP processors offer approximately 100 $\times$ the power efficiency and 10 $\times$ the throughput at half the die size—presenting a solution with the highest performance, lowest power, and lowest cost compared to current MCU and DSP solutions.

Syntiant is also partnering with some of the leading companies to service the rising demand for edge processing. The company announced a partnership with Infineon Technologies to combine the NDP processors with Infineon's XENSIV microelectromechanical systems (MEMS) microphones, creating a platform that is able to deal with the challenges of low-latency voice interaction but also supports audio event and environment classification, and sensor analytics.

Earlier in 2020, Syntiant also announced a partnership with Sensory to deliver a deep learning, multi-lingual voice interface for battery-powered devices. The combined solution merges Sensory's TrulyHandsfree wake word engine and voice control with Syntiant's NDPs, bringing low-latency, real-time inference to edge devices. This cooperation enables manufacturers to implement seamless "voice" commands in dozens of languages.

Tiny Machine Learning

Other manufacturers are accelerating collaborative efforts to accelerate the integration of DSP and voice neural networks with power-efficient edge processors.

CEVA announced that its CEVA-BX DSP cores and WhisPro speech recognition software now also support TensorFlow Lite for Microcontrollers, a production-ready, cross-platform framework for deploying tiny machine learning on power-efficient processors in edge devices. WhisPro is CEVA's speech recognition software for voice wake words and custom command models. TensorFlow Lite for Microcontrollers is a solution from Google for implementing machine learning at the edge, already optimized and available for CEVA-BX DSP cores.

Tiny machine learning (TinyML) is an industry-wide effort to bring the power of AI to extremely low power, always-on, battery operated IoT devices for on-device sensor data analytics in areas such as audio, voice, image, and motion. CEVA's holistic approach to AI at the edge ensures that developers using TensorFlow Lite for Microcontrollers can utilize a unified processor architecture to run both the framework and the associated neural network workloads required to build these intelligent connected

products. CEVA's WhisPro speech recognition software and custom command models are integrated with the TensorFlow Lite framework, further accelerating the development of small footprint voice assistants and other voice controlled IoT devices.

"CEVA has been at the forefront of machine learning (ML) and neural networks inferencing for embedded systems and understands that the future of ML is Tiny going into extremely power and cost constrained devices. Their continued investment into powerful architectures, tools and software which support TensorFlow models provide a compelling offering for a new generation of intelligent embedded devices to harness the power of AI," commented Pete Warden, Technical Lead of TensorFlow at Google.

The CEVA-BX DSP family is a high-level programmable hybrid DSP/controller offering high efficiency for a broad range of signal processing and

control workloads of real-time applications. Using an 11-stage pipeline and five-way VLIW micro-architecture, it offers parallel processing with dual-scalar compute engines, load/store and program control that reaches a CoreMark per MHz score of 5.5, making it suitable for real-time signal control. Its support for SIMD instructions makes it ideal for a variety of signal processing applications and the double precision floating point units efficiently handle contextual awareness and sensor fusion algorithms with a wide dynamic range. It facilitates simultaneous processing of front-end voice, sensor fusion, audio processing, and general DSP workloads in addition to AI runtime inferencing. This also enables brands and algorithm developers to take advantage of CEVA's extensive audio, voice and speech machine learning software and libraries to accelerate their product designs.

DSP Concepts' Audio Weaver


A Comprehensive, End-to-End Audio and Voice Development Platform for Innovators

DSP Concepts is generally considered the leading provider of audio development tools and IP for creators of audio-enabled products, with most of the DSP hardware providers collaborating with their program. This position was further solidified by a recent \$14.5 million funding round led by Taiwan Capital with participation from BMW i Ventures, Sony Innovation Fund by Innovation Growth Ventures, MediaTek Ventures, Porsche Ventures, and the ARM IoT Fund. This latest investment brought the total of DSP Concepts' venture capital funding to more than \$25 million to date. DSP Concepts' core offering, Audio Weaver, is a graphical configuration tool that enables rapid development of audio/voice DSP software, eliminating the need for native coding and kernel optimization. Audio Weaver is easy to use and enables the design, debugging, and tuning of audio processing flows in real-time, reducing product development time. It also enables product designers to build complex digital audio applications by choosing the required audio modules from a library. Simply drag, drop, and connect.

Audio Weaver enables product designers and engineers to integrate, invent, and fine-tune audio-features at unprecedented speeds. The ability to change signal-processing layouts in real-time (without recompiling or reflashing hardware) yields fast iteration times.

Unlike earlier generation development packages, Audio Weaver is processor agnostic and portable across all processors. Product makers are able to separate software development from hardware design in the early prototyping phase, and then

re-integrate in the final design phase. This level of flexibility combined with a cross-platform development environment enables testing multiple features and designs on a PC before porting it to hardware. Additionally, Audio Weaver's Real Time Audio System Check (RTASC) testing and debugging interfaces identify potential issues such as microphone gain mismatch, mic




Audio Weaver is a graphical configuration tool that enables rapid development of audio/voice DSP software, eliminating the need for native coding and kernel optimization. It already powers more than 40 million devices worldwide.

isolation, CPU load, and system latency early in the development cycle. It's critical to verify hardware and software functionality at each step of the development cycle to mitigate the risk of missing product delivery deadlines, and DSP Concepts' RTASC is particularly valuable because it doesn't require any DSP expertise. It is specifically designed for hardware engineers and embedded software developers, and if all goes well you don't have to hire a DSP engineer to check system performance.

The technology is especially useful for OEMs building voice-enabled devices that integrate both microphones and loudspeakers. Voice UI performance depends on all parts of the system functioning perfectly, and the complex nature of these devices makes real-time system checks even more important. RTASC works within Audio Weaver Designer to simulate and test the performance of every element of the device, including microphone matching and sensitivity, loudspeaker harmonic distortion (a major issue for keeping the acoustic echo canceler happy), reference signal to microphone latency, and more. Once your system passes RTASC for voice-enabled products, you can drop in TalkTo—an audio front end (AFE) that enables voice-control in noisy and unpredictable environments. Its suite of proprietary microphone processing algorithms uses adaptive learning technology to cut through the noise. Audio Weaver's AWE Core, is an embedded, dynamically-configured audio-processing engine. The hardware's processor software interacts with the AWE Core by exchanging three types of data: Audio, Tuning, and Control.

Industry acceptance is huge and includes STMicroelectronics, Analog Devices, NXP, Qualcomm, Amlogic, MediaTek, Texas Instruments, and others. Audio Weaver gives product makers access to nearly 500 production-tested modules through a drag-and-drop graphical interface. To speed proof-of-concept and short run products, DSP Concepts is partnering with hardware module makers to provide turnkey audio processing solutions.



For voice applications, DSP Concepts' TalkTo audio front end is a complete software solution, supporting two-mic, four-mic, six-mic, and custom configurations with Acoustic Echo Cancellation (AEC) for up to seven channels.

StreamUnlimited and LinkPlay both provide Audio Weaver-enabled "smart" audio solutions with voice control. Stay tuned for additional announcements over the next few months.

DSP Concepts is also starting to extend beyond purely audio applications and has added support for sensors. For example, by using aggregated data from more sensors (e.g., a motion sensor), a smart speaker "knows" that it has been moved and should generate a new calibration process to optimize the speaker's response for the new position. Sensors are also important for immersive headphone applications that track head motion and headphones/earbuds in order to implement gesture control.

There are countless ways sensors could be used to enhance the performance of audio products, particularly in the smart home space, so it will be exciting to see where DSP Concepts goes with it.

www.dspconcepts.com

Waves Audio

Waves MaxxAudio Sound Enhancement Technology Now Available On Knowles AISONIC Audio Processors

Waves confirmed that its MaxxAudio Sound enhancement technology is now available on the newest Knowles AISONIC Audio Edge Processor, the IA8201. The joint solution makes it easy for design engineers to integrate studio-quality sound into consumer electronics products. Knowles is the latest platform vendor to join an expanding list of chip partners for Waves, supporting developers working to enhance sound on mobile and premium audio devices, including those featuring augmented reality (AR) and gaming features.

The IA8201 processors, the latest product in Knowles' AISONIC family of audio edge processors were launched in 2019 to enable robust voice activation and multi-microphone audio processing optimized for power-sensitive applications. The processors were also specifically designed for advanced audio and machine learning



applications on-device, offering the compute power to perform advanced audio output, context awareness, and gesture control for today's most advanced consumer electronics.

Knowles IA8201 processors include a 128-bit core (DMX) using Knowles' proprietary instruction set and Tensilica HiFi3 cores (HMD), both with Knowles audio enhancements. The DMX is a four-way floating-point SIMD processor targeted toward efficient performance computing (e.g., beam-forming, barge-in, and AEC), while the HMD is targeted toward efficient, low-power, wake-on-voice applications with a two-way floating-point SIMD processor.


The IA8201 combines two Tensilica-based, audio-centric DSP cores: One for high-power compute and machine learning, the other for very low power, always-on processing of sensor inputs. Either can also function as an open DSP controller using Tensilica-based instruction sets. A rich set of audio and general-purpose high-speed interfaces enable flexible interfacing with digital microphones and motion sensors.

Both cores contain dedicated accelerators for FFT, peak finding, and DNNs. The family also offers generous memory capacities that can be used to store various keyword models, engines, and other algorithms.

Rich interface options include 2x I²C, 2x SPI, 2x UART, and 24 GPIOs with interrupts to support various sensor and data connections, plus 4x PDM in, 2x PDM out, 3x I²S/TDM (four channel in/out) to support up to four microphones and various audio connections.

Knowles' open DSP platform facilitates partnerships with third-party algorithm providers and already include machine learning platforms, such as TensorFlowLite, the open-source deep-learning framework for on-device inference.

With the introduction of an IA8201 reference design, Waves



This block diagram shows the AISONIC IA8201 audio edge processor, specifically designed for advanced audio and machine learning applications.

MaxxAudio becomes an ideal solution for output processing and loudspeaker enhancement for smart speakers and headphone development. MaxxAudio provides significantly improved audio quality using technologies that are directly derived from the company's popular software studio tools and technologies. The MaxxAudio suite includes extended bass using psychoacoustics, dynamically boosted treble, sonic maximizing to reveal details, 20 bands of EQ to optimize output, and dynamic tuning that decreases device vibrations.

Utilizing the AISONIC HMD core, Waves MaxxAudio code takes full advantage of the high compute efficiency, specific Knowles instruction set for audio, and low power capabilities of the Knowles AISONIC Audio Edge Processor.

With the Waves-Knowles reference solutions, both companies significantly streamline the design efforts to integrate complex audio output into consumer products, enabling design engineers to reduce overall effort and effectively launch new products.


www.waves.com | www.maxx.com
www.knowles.com | www.aisonic.com

Klippel

Nuvoton Unveils Smart Amplifier Solution with Integrated KCS Speaker Control Algorithm

Klippel announced an exciting cooperation with Nuvoton Technology, a leading provider of consumer-electronics and computer ICs from Taiwan, to release a new type of smart amplifier with integrated Klippel Controlled Sound (KCS) technology that compensates for nonlinear speaker responses. As a result of the project, Nuvoton developed the NAU83G10 and the NAU83G20, two highly efficient mono Class-D audio amplifiers that use a low-latency Tensilica HiFi Audio DSP, combined with high-quality 24-bit audio ADCs for current and voltage sense, and a battery tracking and brownout prevention ALC, all integrated on-chip and optimized for KCS technology.

The powerful Cadence Tensilica HiFi Audio DSP core is able to fully support the complex calculations required for the KCS algorithms, and both devices can support I²C control and I²S/PCM/TDM audio interfaces. The NAU83G10 device delivers an output power up to 8 W into a 4 Ω load, and 6.5 W for an 8 Ω




load, using a highly efficiency Class-G boost converter. The NAU83G20 device delivers up to 20 W for a 4 Ω load, and 1 W for an 8 Ω load, using an external battery supply around 12.6 V.

Running KCS Adaptive Nonlinear Control System


The Klippel and Nuvoton integrated solution enables manufacturers to cope with the undesired effects of signal distortion, heat, and external influences that limit the maximum level and the quality of the reproduced sound. The development efforts between the two companies started in 2017, aiming to offer an integrated, versatile solution for the industry to fully explore the potential of the adaptive and nonlinear speaker model implemented in the KCS algorithm, to improve speaker performance and sound quality. The result is now available in two ultra-low-noise and yet powerful mono boost Class-D amp devices.

The nonlinear speaker model describes speaker states, such as voice coil excursion and velocity, not only at low levels (small signal domain), but also in the large signal domain where the voice coil excursion is high. This enables the mechanical protection system to work very precisely over the full operating range, which is not possible with linear models. In addition, a nonlinear model can identify and actively shift the voice coil to its optimum position using a small DC voltage. This is achieved by exploiting nonlinear distortion found in the current signal, which is caused by nonlinear transducer parameters.

These features allow the transducer to operate over a wider working range without damaging the speaker and improve the sound quality and acoustic echo cancellation (AEC) barge-in performance by reducing linear and nonlinear speaker distortion. Furthermore, the possibility of describing the speaker's behavior in the nonlinear range of operation allows the use of transducers with more nonlinear motor topologies to focus on maximizing efficiency (Green Speaker Design).



Nuvoton NAU83G10 evaluation system block diagram



The Nuvoton KCS development platform

The adaptive identification of the speaker's resonant frequency, damping and other time-variant parameters allows KCS to compensate time-variant behavior and hence ensures that the performance stays constant over the product's lifetime. By doing an automatic system alignment, KCS eases system tuning as the user only has to specify a target response (e.g., a sixth-order high-pass filter at a certain cut-off frequency) and KCS automatically calculates a filter to achieve this target.

The adaptive, nonlinear control system cancels the DC displacement, harmonic and intermodulation distortion generated by speaker nonlinearities. Furthermore, the KCS operates the voice coil at the optimum rest position without mechanical sensor

and compensates for speaker parameter variation due to aging and climate changes.

Other KCS integrated software features include predictive speaker protection against thermal/mechanical overload, automatic speaker equalization to a desired bass response, and online speaker diagnostics to improve product reliability.

The new hardware and software capabilities in the NAU83G10 and the NAU83G20 enable users to safely operate smaller speakers at the physical limits to generate more acoustical output at higher sound quality, while providing reliable protection against thermal and mechanical overload compared to conventional linear algorithms.

Development Platform

An evaluation system for the NAU83G10 and the NAU83G20 smart amp devices is now available, including all the necessary hardware and software to demonstrate the performance of these novel amplifiers together with KCS. The evaluation platform ships with reference speakers, initial KCS data sets, and equalizer settings so that manufacturers can immediately evaluate the system's performance. Users of the evaluation system can get a free temporary license for the Klippel QC measurement system for performing objective tests on request.

Developers will be able to connect and tune their own designs with this evaluation system for their target application. The initial identification of KCS speaker parameters is done almost automatically using the Klippel Analyzer 3 measurement system. The identification of a new speaker only takes a few minutes and the maximum working range of the device under test is identified automatically, providing comprehensive information about the suitability and potential quality of a speaker system.

NAU83G10 and NAU83G20 engineering samples and evaluation systems are available now through all Nuvoton sales channels.

www.nuvoton.com | www.klippel.de

Oxford Digital

Designing and Implementing Audio DSP Algorithms to Optimize Product Performance

Flying below the radar of almost everyone in the audio industry—except for those deeply committed to audio development and in particular with DSP—Oxford Digital is an established reference in audio processing.

The origins of the company—already lodged in a long tradition of audio companies from the Oxford region—date back to the creation of the famous Oxford OXF-R3 digital mixing console. One of the pioneering efforts of the digital audio age, the OXF-R3 resulted in a treasure trove of technology, which is still generating returns today, including one of the very first Ethernet-based audio interconnect and routing protocols over Cat-5/Cat-6, supporting high-resolution audio: SuperMAC, also known as AES50-2005, and HyperMAC.

As Sony divested from its "Oxford Pro-Audio Lab" in June 2006, Oxford Digital was created and sold those networking assets to Klark Teknik, another British company with a remarkable heritage.

Another important product of the Oxford OXF-R3 digital mixing console was the very impressive collection of DSP audio processing and effects algorithms—the foundation for the whole system, painfully coded in dedicated ASICs and subsequently to FPGA cores. That technology served as the foundation for the creation of another company called Sonnox, fully focused on expanding those developments as software plug-ins for the studio and live sound markets.

The remaining operations are what is now Oxford Digital, a company with a respected and experienced team, developing audio solutions for professional and consumer applications, and able to provide engineering consultancy services. Oxford Digital continues to license its legacy DSP cores and sonic tuning tools, provide contract research and development work, deliver software (including real-time embedded) and firmware, and support manufacturers

in system architecture/system level design, and hardware design. Over the years, Oxford Digital supported development and delivered signal processing algorithms for Sony, Yamaha, and Denon/Marantz, among many others, while also working closely with technology and semiconductor manufacturers, such as SRS Labs, Cypress, Wolfson, and Texas Instruments.


Coding for Any Platform

Even though the company today delivers code for FPGA and even ARM platforms, Oxford Digital's main asset was another key hardware development from the Oxford OXF-R3 era. The "Tiny" DSP core, a highly efficient, low-cost, stream processor for audio applications, which Sony and other pro audio companies continued to use, was incorporated in multiple semiconductor chips and is complemented with graphical programming tools dedicated to audio processing and effects, including support for licensed audio algorithms from third parties. While specializing in delivering sound optimization for audio designs, with the Tiny DSP core, Oxford Digital was able to deliver a high-performance unified hardware/software package.

The Tiny DSP Core is sample-rate independent, can be implemented in 16-bit, 24-bit, or 32-bit (or any other bit-width) configuration, and can be integrated into chips (e.g., ADC/DAC front-ends, multimedia processors, SoCs, etc.). It uses a reconfigurable RTL model that permits synthesis of FPGA and ASIC gate-level netlists with a wide range of configurable parameters that can be selected for the needs of the application, and configured to the required performance. As an example, Oxford Digital licensed this technology to ODMs such as China Hansong Technologies and Harman Professional.



John Richards, CEO of Oxford Digital, explains some of the company's black magic from which most of us benefit when enjoying music from affordable audio products.



An evaluation platform that uses a simple Raspberry Pi board, enables testing algorithms developed by Oxford Digital in the target product, before committing any DSP code.



While Oxford Digital is now predominately porting code to a variety of platforms adopted by large-scale CE manufacturers and OEMs—improving sound quality for TV sound, soundbars, and portable Bluetooth speakers—the team’s experience at the DSP core level allows it to embrace projects that require highly optimized code. The company’s graphic development tools are being constantly expanded, also embracing headphone and earphone correction, or enhancing dialog intelligibility, adding to its established acoustic tuning systems to improve sound quality in consumer equipment, with automated frequency correction, bandwidth extension, or noise reduction.

Oxford Digital offers popular tools, such as its EQ-master, for preset effects and speaker/earphone correction; Loud-master, to make the audio seem louder without significant effect on battery life and also useful as an MP3/ATRAC/lossy-compressed audio restoration tool; Peak-master, to enhance listening and intelligibility in noisy environments, removing the need for frequent adjustment of local volume controls; Edge-master, to add “punch” to percussive effects or to remove transients to avoid overload and save battery drain or for background listening; Bass-master, to improve the bass response of speakers and earphones beyond the range that they can naturally reproduce; and Wide-master, to enhance listening and give the impression of a bigger space. In any combination, these tools provide audibly impressive results.


More recently, Oxford Digital implemented a new generation of tuning tools that totally decouple the assembly code from the evaluation platforms, using a simple PC application and the processing of PCs to run the tuning process. The resulting code can be exported as an embedded application, which will enable testing any target system. The company also created a development platform leveraging a simple Raspberry Pi board (with an ARM Cortex processor), in a ready-to-use compact chassis with physical buttons to turn the processing on/off.



The user interface of the NMFx processing PC app, and before and after loudness plots of TV material, show reduced loudness-unit peaks and amplified low-volume sections reducing range by ~10.8 dB loudness units.

The new approach enables loading code into a client’s existing product in a matter of weeks. After writing a simple PC application with the basic parameters and sending the result to an ODM (this is delivered as a time limited PC Application under NDA), in just a few weeks the client can evaluate and approve the results. Evaluation tools such as these, which can be optimized for any platform, can save months or even years of development work.

As John Richards, CEO of Oxford Digital, explains, the most important thing in this type of development is always to establish the algorithm and define the functional parameters, allowing optimization of the processing blocks, and hence simplifying the resulting code. Embedded deliverables can vary depending on



The Oxford Digital HFFx II algorithm is beneficial where lossy, sub-band codecs have damaged the high-frequency content resulting in a dull, lifeless sound and also in restoring and/or enhancing audio from band-limited material.

requirements. The standard system is obfuscated C code, but Oxford can port to low-cost Qualcomm CSR8670 or MAPX audio processors and many other systems.

Latest Developments

When we have the opportunity to try out the algorithms that Oxford Digital perfects for real-world clients, particularly those in the consumer market, one can only be impressed. The company continues to leverage all the vast work of its original processing tools, also reinforced with licensed processing tools from third-parties, and continues to develop the most effective algorithms for always changing requirements.

Significant developments in those tools include the second generation of the HFFx high-frequency restoration algorithm, which formerly was targeted mostly at reconstructing high frequencies lost by poor sub-band audio coding (e.g., MP3 and Bluetooth), but can also be applied to extending the bandwidth of material that was recorded band-limited (analog tape) or broadcasted with low-pass filtering in the channel encoder.

In the HFFx II, Oxford Digital also introduced an “Auto” mode, in which the channel bandwidth is continually monitored and the frequency range over which reconstruction takes place automatically adapts to give a full bandwidth reconstruction.

Another recent example in Oxford Digital’s arsenal of processing is the NMFx (Night Mode effect) loudness gain algorithm. While dramatic sound with wide dynamic range is very enjoyable in a theater, it is not always practical in the home due to the risk of disturbing sleeping children, other family members, or neighbors. Turning the sound level down usually results in difficulty in hearing the quiet passages of dialog leading to a frustrating and unenjoyable experience. NMFx solves both these problems by managing the sound level to reduce very loud sounds while boosting speech intelligibility and the ability to hear quiet sounds. In addition, NMFx can prevent hearing damage while listening in a noisy environment by providing a more consistent sound level. This is a useful pre-tuned algorithm to be applied in active speakers and soundbars, TV sound, and AV receivers, but it is also great for earbuds and headphones.

Finally, Oxford Digital offers the updated SmartFx tuning tool, successor to the company’s acclaimed FixFx. SmartFx incorporates a fixed configuration of audio DSP effects that allow limitations in a device’s acoustic performance to be overcome and features sound processing derived from Oxford Digital’s pro-audio heritage. The tuning tool now can be used to tune 1.0, 1.1, 2.0, and 2.1 products (not just 2.0), now using a single set of controls to achieve the desired objective—hence the name SmartFx. No evaluation board hardware is required. SmartFx includes eight parametric equalizers, a switchable crossover network used for speaker response correction and pre-set effects; a Bass Enhancer; a Low-Frequency Boost (a combination of a Bass Enhancer with Dynamic Gain Control to make the signal louder); a High-Frequency Boost, which uses a combination of a Limiter section to make the signal louder; and LoudMaster, an effect that can increase the volume without increasing peak signal level.

www.oxford-digital.com



John Richards at CES 2020 demonstrates the SmartFx speaker tuning tool, now available as a PC application that can tune for all platforms.

DSP
CONCEPTS

voice control that works

Enable high-performance voice control in your embedded products in record time.

dspconcepts.com


The Evolution of Audio DSPs

By

Youval Nachum

(Senior Product Marketing Manager, CEVA)

We're all familiar with the buzzword DSP, which stands for digital signal processing or processor, and it seems to have been around forever in the audio world. But how has DSP technology evolved to where we are today? What has its impact been on the user experience? And what does the future hold for DSP?



The first dedicated audio digital signal processor (DSP) chip was Texas Instruments' TMS5100—first used in its Speak & Spell in 1978.

Let's start our evolutionary journey by briefly reviewing some ancient history: The genesis of DSP is arguably in 1948, when Claude Shannon published his article, "Mathematical Theory of Communication," which explained the key concepts of how signals are transmitted and received as digital bits. The earliest practical DSP applications in audio were in speech processing for telephony in the 1960s, and the first dedicated audio digital signal processor chip was Texas Instruments' TMS5100, first used in its Speak & Spell in 1978.

For audio processing of music, we need to fast forward a few years, to where CD players joined the party in the 1980s, and then the MP3 players arrived a little later at the end of the 1990s. As digital audio became mainstream, DSP became increasingly important—enabling high-quality music and sophisticated processing features to become available in low-cost consumer devices.

The Classic DSP Era

In the era before smartphones, DSPs were relatively simple by today's standards. A typical DSP device (e.g., the CEVA TeakLite family) would have an architecture that enabled it to directly access the digital audio data in memory, perform a mathematical

operation on it, and write the result back to memory. A very common operation in DSP applications is the Multiply-Accumulate (MAC), where two numbers are multiplied together, and the result added to an accumulator, or running total—the MAC must often be performed very frequently, so the processor must complete it as quickly and efficiently as possible. The TeakLite DSP provides the capability for a 16 × 16-bit MAC, with either one or two simultaneous MACs supported per device.

By performing such tasks repetitively on every sampled digital value, the DSP could encode and decode the audio signal (e.g., retrieve it from a CD), and then perform processing tasks such as filtering. The user may have noticed improved voice and speech quality, but most of the DSP activity was hidden under the hood.


The DSP device would have used simple fixed-point arithmetic, and ran at a clock frequency of likely somewhere between 50 MHz and 200 MHz—very slow by today's standards. This meant that software had to be hand-coded in low-level assembly language to get the required performance, making programming difficult, although vendors did provide some software libraries for common tasks to make things easier.

Everything Connected

Moving on to the iPhone era, DSP processors are starting to get much more powerful, which means they can be used for increasingly more sophisticated tasks. Smartphones are taking over from MP3 players, and mainstream consumers are also starting to experience ever-higher sound quality, due to DSPs in their TVs, computers, and video games. Just as importantly, we're moving to a connected world, with Bluetooth, Wi-Fi, and cellular everywhere, and audio streamed between multiple devices—all requiring extra processing.

Noise control for voice interaction is becoming cleverer, with beamforming techniques using DSP to combine the signal from multiple microphones in devices such as smart speakers. Speech recognition is bringing sci-fi voice interfaces to life, in the car and on our phones. If you're watching a movie at home, you may well be listening to a 5.1 multi-channel system, and cinemas are investing in high-quality sound to keep a step ahead.

Behind all of these developments is the DSP. But as the user experience is enhanced, there's an ever-growing need for more and more sophisticated processing. At the same time, most of these



Some of our readers might remember this Nokia 3310 mobile phone from 2000.

applications are in portable devices, so battery life is vital—meaning that DSP devices need to have very low power consumption to be viable. There's a constant trade-off between features, sound quality and power usage. If users have to charge a gadget too often, they're unlikely to care how great it is in use. In fact, engineers using DSPs often talk about

HFFx Auto

High frequency restoration effect

OXFORD
DIGITAL

New 'Auto' mode dynamically tracks channel bandwidth and automatically adjusts for optimum reconstruction:

Targeted at embedded processing in Soundbars, Speakers, Headphones, Earbuds, TVs and AVRs for:

- Restoration of high frequency content that has been removed by lossy sub-band coding (e.g. MP3, AAC, Bluetooth) producing a dull, lifeless sound
- Switching between channels of varying audio bandwidth (e.g. on a TV)
- Extending the bandwidth of material that was recorded band-limited (e.g. recorded to analogue tape, low sample rate, or other restrictions)
- Up-conversion of audio from 1Fs (44.1KHz and 48KHz) to High Resolution audio rates such as 2Fs (88.2KHz, 96KHz), 4Fs (176.4KHz, 192KHz) to create a more 'open' sound

Incoming audio with missing HF

Reconstructed harmonics

Reconstructed signal

Download the product brief at
<http://www.oxford-digital.com/products.htm>

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

www.oxford-digital.com

excellence in audio technology

These are Google Pixel Buds 2.



devices in terms of “performance per watt,” or how much a processor can do for a given amount of power, rather than its overall performance.

These new demands meant that DSPs added more powerful number-crunching engines, and more ability to handle data in parallel, using Very Long Instruction Word (VLIW) architectures. They mostly moved to floating-point, not fixed-point, arithmetic, and clock speeds increased to around 500 MHz to 600 MHz. For example, the CEVA-X and the CEVA-BX families provide high performance with support for up to eight 16 × 16-bit and four 32 × 32-bit MACs. Overall, these technology improvements meant that the hardware could keep up with users’ expectations, overcoming the performance limitations of older DSPs.

On the software side, it was no longer practical to program the DSPs in low-level assembly language. Instead, software engineers wrote their code in the high-level “C” language—which was not quite as



Augmented reality (AR)/virtual reality (VR) add smart features, such as 3D audio and sound sensing. (Source: CEVA)

efficient for the processor, but far quicker and easier to develop in. Device vendors also provided more comprehensive libraries and pre-written modules, to cut development time.

Experience and Interaction, Driven by AI

Bringing things up-to-date, today we’re more surrounded by DSP in the gadgets around us. We can consider ourselves as moving into a new “Experience and Interaction” stage, with technology in transition, and new applications appearing and becoming mainstream.

For example, smart speakers (e.g., the Amazon Echo) are in many of our homes, and rely on DSP to process the audio their microphones capture. They are also increasingly adding Artificial Intelligence (AI) capabilities, using neural networks to process our speech, and to learn how to respond.

This enables the devices to perform Natural Language Processing (NLP)—which is a jargon-heavy way of saying you can talk to a gadget like it was a person. You don’t need to remember exactly what wording Alexa needs to hear, just ask it a question and you’ll get an answer. It seems simple when you use it, but NLP requires a lot of clever DSP, and makes the user experience far better. To achieve the required performance, vendors such as CEVA now often include dedicated speech recognition engines and NLP capabilities in their DSPs that are targeted at voice applications.

We’re seeing a step up in the audio quality provided in our cars, with automakers recognizing that an audiophile music experience can be a big selling point. Reducing noise (e.g., from the road and engine) is another automotive audio application that is getting increasing attention.

Many new gadgets are also starting to become mainstream, with different degrees of success, such as smartwatches, virtual reality (VR) glasses, and hearables that add smart features such as 3D audio and sound sensing.

All of these new capabilities and devices rely on DSP for audio processing, and in many cases, this means DSP devices need to be extremely high performance, while still keeping power consumption to a minimum. With more and more features expected of our gadgets, there’s always pressure to do more with less power—both to extend battery life, and to allow a smaller battery to be used in space-constrained portable products. Weight, size, and time between charges are important differentiators for consumers, so manufacturers need the right DSP to stay competitive.

Processing is also starting to be handled more and more by edge devices, such as the Echo speaker,

rather than data being sent to the cloud—thus improving security, removing latency and eliminating the need for an ever-present Internet connection.

On the other hand, many consumer products will now include a DSP “core” as part of a larger silicon chip—companies such as CEVA are developing power-efficient DSPs that can be integrated into other chips, which drives cost and complexity down.

The high performance needed by many applications can be provided by “vector processing” DSPs, which act on multiple pieces of data at the same time to increase the throughput of the processor. Clock speeds have now reached 1 GHz or faster, and multiple processing capabilities are combined in a single device. For example, the CEVA SensPro™ is a vector DSP, which also includes scalar components, as well as support for neural networks. This means it can handle huge amounts of data from real-world sensors, and has the flexibility needed to meet the requirements of diverse applications.

Software has become more complex, with new support needed to configure neural networks. For example, a modern DSP such as the SensPro is shipped with software development tools—including a compiler, a debugger, and a real-time operating

system (RTOS), as well as framework support and compute libraries for neural networks.

Conclusions

Whatever your view of digital audio vs. analog, it’s fair to say that the DSP has transformed our audio world for the better. With new applications and products appearing all the time, the future will only bring more demand for DSPs—stretching their hardware and software capabilities to the limit to deliver great performance from tiny amounts of power, and in particular, employing AI to improve the user’s experience.

We’ve come a long way with DSP, but there’s no limit in sight as to the technology’s capability. It’s going to be an interesting journey.

About the Author

Yuval Nachum serves as CEVA’s Senior Product Marketing Manager for the company’s audio and voice product line. Nachum brings more than 20 years of multi-disciplinary experience, spanning marketing, system architecture, ASIC, and software domains at leading technology companies. He is passionate about anticipating long-term trends and leading technical programs to their successful completion. He is highly proficient in combining market requirements, product definitions, industry standards, and design innovations into breakthrough products. Nachum holds a B.Sc. and M.Sc. in Electrical Engineering from the Technion – Israel Institute of Technology.

SYNTIANT®

Syntiant's **Neural Decision Processors™** enable customized **Always-On Voice** experiences at the edge for devices of any size.

- consuming <140 microwatts
- 100x the efficiency, 10x the throughput
- increasing privacy, reliability and responsiveness
- smallest footprint

www.syntiant.com
info@syntiant.com

Electroacoustical Testing Time-Domain and Distortion

By
Richard Honeycutt
(United States)




Photo 1: The sound from the main speaker arrives at 1.12 seconds, and the un-delayed sound from the fill speaker arrives at 1.261 seconds.

The Sound Control article published in the October 2020 issue and this article were written for readers who sometimes need to measure the performance of live-sound electroacoustical systems.

Although frequency-response testing is the most common form of electroacoustical measurement, there are others—one of which is delay. In large venues, there may be main speakers covering most of the venue, plus “fill” speakers covering the rear seats. In order to acoustically align the fill speakers’ sound with that from the main speakers, electronic delay must be introduced into the fill speakers’ signal chain to compensate for the “time of flight” of sound originating at the main speakers.

In years past, those who had a storage oscilloscope could play an impulsive sound through the speakers and measure the time between the “main” and “fill” spikes on the oscilloscope trace. Today, it is easier to record the sound as a .wav file and note the time between spikes on the screen of a .wav editor (see **Photo 1**). In this photo, the delay between sound arrivals from the main and fill speakers is 141 ms. By adjusting the fill’s delay, you can synchronize the sounds to the extent desired.

Measuring the delay makes setting the speakers’ delays in a DSP or digital delay easier, but you can accomplish the same thing by listening to the playback and adjusting it until you hear only one impulse. The best signal for this kind

of aural delay test is a rimshot on a snare drum, although a group of anechoically recorded bongo strikes spaced about 5 seconds apart also works well. Even though audio professionals often state that the Haas effect merges sounds within a 50 ms window, the truth is that the actual Haas window width depends upon the signal used and the reverberation time (RT) of the room.

In a room having a really short RT, a snare drum rimshot can enable a good listener to time-align speakers within 10 ms. The 50-ms value applies for speech in a normal room with an RT of a second or two. For most music, the Haas window is about 80 ms, although it also depends upon the type of music. Some sound-system engineers prefer that the sound from the main (“precedence”) speakers lead the sound from the fill speakers by up to 10 ms, in order for the acoustical image to be pulled to the front of the venue rather than to the fill speakers. There is no consensus on this practice, and excellent engineers may prefer the sound delay of the fill speakers to be zero, 5 ms, 10 ms, or someplace in between.

Speaker Polarity Testing

Although mis-polarized speakers are not as

APx Power, Unbelievable Price

APx500 software paired with 3rd party audio interfaces



APx500 Flex Audio Analyzer

Need mics?
GRAS
www.gras.us

**Audio
precision**
www.ap.com

common as they were years ago, there is still an occasional need to test speakers for proper polarity. Mis-polarized full-range speakers will give a weak bass response and can produce a honky sound in the voice range due to destructive interference. In one civic center for which I consulted some years ago, the cause of the owners' complaint—lack of bass response—was traced to the subwoofer polarity being reversed with respect to that of the main speakers. This system was installed by a highly respected contractor, showing that even the best of us is subject to making mistakes.

A variety of speaker polarity testers are available—basically they produce an impulsive sound that is fed to the system, then analyze the result picked up by a microphone (see **Photo 2** and **Photo 3**). If you're testing a passive speaker with a cone-type driver, you can connect a 9 V battery's positive terminal to the speaker's positive input terminal to check polarity. If the cone moves forward (away from the speaker magnet), the polarity is correct.

For active speakers or speakers whose cones cannot be seen (including horn drivers), you can record the sound made when the battery is connected and disconnected, then look at it on a wave editor. The results from the screen of the wave editor will also appear like those shown in Photo 2 (in phase) or Photo 3 (out of phase). Note that when the polarity is correct, the first (upper) pulse is positive.

Acoustical Imaging

Although I am unaware of rigorous testing to support this, many audio engineers believe that audiences enjoy a drama, a lecture, or a musical performance more if the sound seems to come directly from the actor, lecturer, or musician, respectively. The so-called "acoustic image" is identified by the directional response of the human hearing system.

The sound source is aurally perceived to be in the direction of the first-arriving sound (the "law of the first wavefront"), although which ear experiences the louder sound is also instrumental in the brain's aural localization. The interaural phase difference affects localization as well, and interaural spectral differences particularly affect vertical localization. "Stairstep," "footlight," or "stage-lip" speakers are often used to pull the acoustical image down or to the center in cases where the main speakers are suspended high above the audience, or are located on each side of the stage (see **Photo 4**).



Photo 2: The impulses from the two speakers are in phase, so the speakers are wired correctly.



Photo 3: The out-of-phase impulses show that the speakers are wired out of phase.

Since adjustment of delays and relative speaker levels can affect acoustical imaging, it is important to be able to measure it. There is no standard instrumentation-based method for doing this. Therefore, the best practice is to play a recording of speech and listen while walking down an aisle with your eyes closed, occasionally stopping to concentrate on where the sound seems to be coming from. Of course, you'll have to face where you think the sound source is and open your eyes periodically in order to see whether the acoustical image is where it should be.

Distortion

Usually we only need to test sound-system distortion if someone hears a problem. Distortion can be generated in the electronic parts of the system, or in the speakers. If there are multiple speakers in the system, and distortion is heard in all of them, there is likely a problem in the electronics. By monitoring the signal at various points in the system—mixer output, DSP output, power amplifier, you can use the divide-and-conquer method to find the defective piece of equipment.




Photo 4: This section of a computer model of a sanctuary shows the main speaker and the stairstep speaker used to bring down the acoustic image.

The "Must Have" reference for loudspeaker engineering professionals.

Back and better than ever, this 7th edition provides everything you need to become a better speaker designer. If you still have a 3rd, 4th, 5th or even the 6th edition of the Loudspeaker Design Cookbook, you are missing out on a tremendous amount of new and important information! **Now including: Klippel analysis of drivers, a chapter on loudspeaker voicing, advice on testing and crossover changes, and so much more!**

A \$99 value!
Yours today for just **\$39.95.**

www.cc-webshop.com



Photo 5: Piezo tweeters such as these can be connected across pretty much any audio line.

One way to listen for distortion in electronic equipment is to use clip leads to connect a piezo tweeter, such as one of those shown in **Photo 5**, across the output of the equipment under test, be it a mixer, DSP, CD or MP3 player, computer, or whatever. The piezo tweeter is very sensitive and has a high impedance (~ equal to the reactance of a 2 μ F capacitor), so it draws very little current. If a sine-wave test signal is used, only a small amount of clipping is needed in order for it to be audible on the piezo.

The frequency of the test signal is not very important, because what you are listening for is the harmonics produced by distortion. Unless the test signal is above about 1 kHz, you may hear none—or very little—of it, but that's not important. If you hear distortion at the output of a piece of equipment, but not at the input, you've located the problem.

If the distortion is in a speaker, you can use a manually swept sine wave as a test signal to

find the frequency range in which the distortion appears. This leads you to identifying whether you have a defective woofer, midrange (in three-way systems), or tweeter. Be careful not to overpower the high-frequency drivers, or you may create a problem where there was none.

Although distortion testers are available, and many Fast Fourier Transform (FFT) test platforms will measure distortion, your ears are usually the easiest and least expensive instruments to use for this purpose. Play music on the system and listen. Rock or similar music with distortion in the recording is useless for this purpose. Cleanly recorded flute or piano music (for mid- and high-frequency drivers) or smooth male vocals (for low-frequency tests) work well. Janis Joplin recordings are not recommended for this application!

Speech Intelligibility

The other important electroacoustical test is speech intelligibility: either %ALcons or the Speech Transmission Index (STI). This test requires special instrumentation, and the protocol depends somewhat upon what instrument is used.

The first speech intelligibility metric in common use was called "percentage Articulation Loss of Consonants" (%ALcons), and was introduced by V. M. A. Peutz in 1971 [1]. The basis of determining intelligibility was the ability of listeners to correctly distinguish consonants pronounced by a talker in the room. A standardized list of test syllables was used. Peutz developed a method of predicting %ALcons from the RT and background noise level in the room.

Houtgast, Steeneken, et al. introduced a method of specifying speech intelligibility using the Modulation Transfer Function (MTF), which quantified the effect upon speech articulation produced by any defect in the electroacoustical and/or acoustical path between the talker and the listener [2]. The metric is called the STI. The underlying theory is that speech can be considered to originate as a continuous wave created by the vocal folds, and then modulated by the throat, lips, and tongue. The "intelligence" or information is carried by the modulation, and the depth of the modulation correlates with speech intelligibility. Reverberation, background noise, and other

acoustical or electroacoustical effects can decrease the modulation depth of the acoustic wave reaching the listener's ears, hurting speech intelligibility.

Speech intelligibility in %ALcons can be measured by any of a variety of test systems, usually using a swept-sine or maximum length sequence (MLS) test signal. STI requires a specific test signal usually provided by the manufacturer of the test system. Two variants of STI measurements are RASTI (a faster adaptation of STI standardized in 1988 but now obsolete), and STIPA (a modified STI measurement designed for implementation on small portable instruments).

Figure 1 shows the correlation between STI score or Common Intelligibility Scale (CIS) score and descriptions of perceived intelligibility of speech. (CIS is mathematically derived from STI measurements.)

Wrap-Up

Most acousticians and sound engineers began their careers with an interest in home stereo or cinema sound, where the important measurements pretty much consisted of frequency response,


harmonic distortion, intermodulation (IM) distortion, perhaps transient IM distortion (TIM), and noise. Most of these measurements are not useful in evaluating live-sound electroacoustical systems. Instead, we have to look at other performance metrics. Hopefully, this Sound Control article and the previous one have been helpful for readers who sometimes need to measure the performance of such systems. 

Author's Note: This series of two articles (Sound Control's October and November articles) on electroacoustical testing is respectfully dedicated to the memory of my late friend and colleague Fred Schafer, who helped me hone my electroacoustic testing abilities over a period of many years.

References


- [1] V. M. A. Peutz, "Articulation Loss of Consonants as a Criterion for Speech Transmission in a Room," *Journal of the Audio Engineering Society* (JAES), December 1971, Volume 19, Number 11.
- [2] T. Houtgast, H. J. M. Steeneken, and R. Plomp, *Acta Acustica United with Acustica*, Volume 46, Number 1, September 1980, pp. 60-72 (13).

Figure 1: The Speech Transmission Index (STI) and the Common Intelligibility Scale (CIS) are two current ways in which speech intelligibility is quantified.



The Electrostatic Speaker Design Cookbook

is a complete guide for the novice and the experienced builder to successfully build this unusual sound reproduction device.



THE ELECTROSTATIC LOUDSPEAKER DESIGN COOKBOOK
by ROGER R. SANDERS

cc-webshop.com



PURIFI

www.purifi-audio.com

PURE SOUND



1ET400A



PTT6.5W PTT4.0W



sales@purifi-audio.com

Simple Improvements to Legacy Audio DAC Output Interfaces Using Op-Amps

By
Michael Steffes
(United States)

In this article, our audio electronics expert discusses how to achieve improved performance out of audio digital-to-analog converters by choosing different resistors.

The commonly recommended output amplifier solutions for audio digital-to-analog converters (DACs) are perhaps leaving a little performance on the table through poor resistor selections. A few simple value changes to account for some second-order effects will move in the direction of reduced harmonic distortion. The complementary or differential outputs of all audio DACs can be either a current or voltage output. Both can benefit from the enhancements described here.

Probably the most common recommendation for a complementary current source output DAC shows termination resistors to ground then a simple differential to a single-ended op-amp stage to remove the input DC common mode and produce a single-ended output (see **Figure 1**).

This recommendation has been updated (from Reference 1) to use the unity gain stable, wideband

OPA820 (see Reference 2) voltage feedback amplifier (VFA), where the full-scale DAC current is assumed to be 20 mA coming out of the DAC with each side swinging ± 10 mA around a 10 mA midscale.

This legacy solution suffers from an imbalanced input impedance on the two sides. One fundamental theme for differential outputs is to match their loading (automatically achieved for differential loads) to get the best possible even-order harmonic suppression. The impedance looking into the inverting op-amp input of Figure 1 is actually an active impedance. As the DAC output voltage is going positive into the R_G element, a portion of the negative going voltage into the non-inverting side appears on the inverting summing junction. This acts to make the R_G element appear smaller than its physical value. The required adjustment term is $1+\alpha$ where:

$$\alpha = \frac{R_2}{R_1 + R_2} \quad \text{Equation 1}$$

Using this simple factor, the apparent active input impedance looking into R_G is given by Equation 2:

$$Z_A = \frac{R_G}{1 + \alpha} \quad \text{Equation 2}$$

The example of Figure 1 shows an $\alpha = 2/3$. Using that in Equation 2 solves for an active input impedance on the inverting side of:

$$500 \Omega / (1 + 2/3) = 300 \Omega = Z_A$$

This differential to single-ended circuit operates correctly if the resistor ratios are equal—it is

not necessary that the “values” be equal. It is a simple matter to adjust R_1 and R_2 to the same ratio, but also to match the 300Ω loading seen on the inverting side. Since we will now have 300Ω in parallel with each resistor to ground, it is also a simple matter to adjust those two resistors up slightly to meet the desired 25Ω DAC output loading. Making these resistor value changes to Figure 1 gives the improved circuit shown in **Figure 2**.

The ideal “gain” for this circuit would be 100 Ω (or 40 dB Ω). The original circuit of Figure 1 was a bit off, while the simple improvements in Figure 2 hit that gain exactly as shown in **Figure 3**. This example from Reference 1 far exceeds the usual bandwidth requirements for an audio interface, but these concepts apply to any lower speed audio op-amp solution as well.

Getting the gain correct is one thing, getting the voltage swings at the DAC outputs exactly matched (180° out of phase) should also move in the direction of incrementally lower even-order harmonic distortion. Setting the input current source to a ± 10 mA 1 kHz signal gives the DAC voltage swings into the unbalanced circuit of Figure 1 shown in **Figure 4**. These are clearly not matched in amplitude, whereas the adjustments in Figure 2 match these up perfectly. With only two resistor values changes, this simple improvement to the legacy circuit of Figure 1 can be easily tested.

The improved matching circuit of Figure 2 will slightly degrade the apparent common-mode rejection ratio (CMRR) and create slightly higher DC offsets. If the application is intended for a complementary output DAC, and the input impedances are matched, there is no common mode signal to reject, so lower CMRR is inconsequential. More general-purpose audio differential receivers




Figure 2: Matched impedance and gain adjusted complementary to single-ended op-amp solution




Figure 3: Complementary input to single-ended output response shapes




Figure 1: Typical differential current output to voltage circuit (see Figure 3 from Reference 1)




Figure 4: Input voltage swings for the simple unbalanced circuit of Figure 1

overcome this active input impedance imbalance using a second op-amp (see Reference 3). That is not needed for the audio DAC application and only a more judicious resistor value selection will move to a more balanced design.

Voltage Output Audio DAC Impedance Balancing

These issues also apply to the typical circuit of




Figure 5: Example differential voltage output DAC reference design




Figure 6: Voltage output DAC differential to single-ended design with balanced input impedances

References

[1] W. Kester, "DAC Interface Fundamentals," Analog Devices Tutorial MT-019.

www.analog.com/media/en/training-seminars/tutorials/MT-019.pdf

[2] "Unity Gain Stable, Low Noise, Voltage Feedback Operational Amplifier," Texas Instruments OPA820 Datasheet, revised October 2016 www.ti.com/lit/ds/symlink/opa820.pdf

[3] W. Kirkwood, "Audio line receiver impedance balancing using a 2nd diff amp," EDN, January 2013 www.edn.com/audio-line-receiver-impedance-balancing-using-a-2nd-diff-amp

[4] J. Caldwell, "Headphone Amplifier for Voltage Output Audio DACs," Texas Instruments reference design, www.ti.com/tool/TIPD189

[5] "SoundPlus High Performance, Bipolar Input Audio Op Amp," Texas Instruments OPA1611 Datasheet, revised August 2014 www.ti.com/lit/ds/symlink/opa1611.pdf

Celebrating
37 YEARS
1983-2020

MENLO SCIENTIFIC

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

**MENLO
SCIENTIFIC**

audioexpress.com | November 2020 | 39



Photo 1: The GRAS 45 CC Acoustic Test Fixture is a flexible and configurable platform for dual-channel testing of headphones and headsets. It has multiple configuration options: microphones, IEC 60318-1 ear simulators, and mouth simulators.

compared (e.g., because the acoustic test fixtures are different), it may be possible to develop transfer functions that adjust for the difference between different fixtures. Note, a transfer function for comparison between test fixtures is only valid for a given device.

What to Test: Electronic Devices

When testing a purely electronic device (e.g., a power amplifier, headphone amplifier, DAC, ADC, or really any device with an electrical input or output), one thing to note is that even the simplest test will yield valuable data. That is, just measuring level, distortion, and noise with a single tone (e.g., 1 kHz) effectively yields a lot of information. A single-tone test can be performed easily and quickly, likely in as little as one second, and will functionally verify most of an electronic signal chain. It verifies that the power supply is functional, that the signal path and all electronic components in it are correctly installed and, by measuring the added spurious noise, verifies that many adjacent devices are not defective. Adding a few more test tones or points for practical purposes provides almost complete validation.

A classic approach for a full-bandwidth electronic device is a test with points at 20 Hz, 1 kHz, and 20 kHz. Testing at additional frequency points will verify that output coupling filters are correct. In electronic devices, beyond basic functionality, a classic cause of fallout is variability or tolerance in the passive analog components (i.e., the capacitors, resistors, and inductors), particularly in the output filters used in power and headphone amplifiers. By

choosing test frequencies at the corner points of the filters, you can easily verify that these components are correct without having to perform a complete frequency sweep of the device. In addition, if you offset the frequencies in each channel (e.g., test at 19.5 kHz in one channel while testing at 20 kHz in the other), you can simultaneously identify any crosstalk.

What to Test: Speakers

For speakers, earphones (see **Photo 1**), and balanced armature drivers, it is usually only minimally useful to test the device at one frequency. Usually these devices are swept across a range of frequencies of interest. However, since uncorrelated noise is not a useful measurement on a speaker driver, the sweeps can be very fast, with a few second-long sweep usually enough. When testing a speaker driver, it is useful to examine the acoustic response of the device and the impedance function. Combined, the two will essentially reveal whether the system is performing correctly or not.

For an easy incoming inspection of raw speaker drivers, I would like to suggest the value of an electrical-only measurement of impedance. Because any deviation in mass of the cone or compliance of the suspension will cause a shift in the impedance curve—particularly at the frequency of the resonant peak—it can be used to quickly weed out devices with common assembly errors. In addition, measuring just impedance considerably simplifies the test requirement. No acoustic test box or measurement microphone is required, considerably simplifying the test setup. Finally, for the best measurement repeatability, impedance should be measured at small signal levels, that is a signal level where the driver's motion is within the linear, piston range of motion. This leads to a common disconnect between manufacturing and R&D. In manufacturing, speaker drivers are usually measured at high signal levels, to test the driver through their full range of motion. However, the impedance curve of a driver at high signal levels cannot be compared to the typical small signal used when calculating Thiele-Small (T-S) parameters, the more common measurement made in R&D.

Modern audio analyzers can usually measure all the important parameters of a speaker in a single fast sweep. With a measurement microphone and electrical connection, a speaker's acoustic frequency response and distortion can be verified along with the impedance and resonance, Rub & Buzz, or high-order harmonic distortion, and even polarity can all be checked in a single test.

What to Test: Microphones

A microphone is often tested just for frequency response, usually with a swept signal across the frequency range of interest. This can be used to digitally calibrate the sensitivity of the microphone for noise cancelling application and phase data can also be captured for microphone beam forming application. A second test is a seal test. For this test, the microphone is swept twice, with the acoustic port for the microphone occluded the second time. This is to verify that the microphone is appropriately acoustically isolated from other parts of the device.




Figure 3: If measured data is collected and stored, it is possible to examine it for several purposes.

Tracking Data

One area that is often overlooked in manufacturing tests is the tracking and analysis of test results beyond simply recording pass and fail results. If measured data is collected and stored, it is possible to examine it for several purposes (see **Figure 3**).

First, a field failure can be audited or traced against the measurements made during manufacturing. Perhaps the manufacturing tests need to be modified to check for a failure mode that is otherwise being missed. Or perhaps the device passed QC, but marginally, and the pass/fail limits should be adjusted. Perhaps the device passed but only after extensive retesting (i.e., an operator simply retested a device until it passed), indicating the device had erratic behavior and policies on retest need to be evaluated. Even as devices pass, it is useful to examine trend data to detect drift in materials and tolerances in order to preempt future failures.

Conclusion

While it may seem like an abstract connection, the technical details of

the measurements made in manufacturing are directly related to an organization's overall quality standard and a product's value proposition. How is the product priced? Is it positioned as a high-end or luxury item? Or a value-oriented product? What is the cost of a field failure? (Either the direct cost of managing return and replacement or in damage to the brand.) From these questions, a budget can be set for test and quality control. Will every unit manufactured be subject to full functional test, will basic functionality be checked, or will only items subject to common failure be inspected? Will every unit be tested or will manufacturing just be sampled?

With the test objectives and the budget set, it is then possible to outline the technical objectives of the test. These may be to just probe a device with a single tone and record output level and THD+N. With a larger test budget, it may be to use multitone, continuous sweep, or another fast measurement technique to characterize a device over a more complete operational envelope. Ultimately, the technical requirements for quality control tests are the product of overall quality standards. 

About the Author

As Vice President of Product Development at Audio Precision (AP), **Daniel Knighten** is focused on the definition of the features, functions, and price of new products while managing AP's engineering team and shepherding products through their complete life cycle. He also supervises the AP technical support team, a rich source of firsthand experience on how customers use, and aspire to use, AP analyzers, software, and test accessories.




LSK389

The Industry's First
100% Noise Tested
Ultra Low-Noise
Monolithic Dual
N-Channel JFET



Ultra Low-Noise:
 $e_n = 1.3\text{nV}/\sqrt{\text{Hz}}$ (typ),
 $f = 1.0\text{kHz}$ and NBW =
1.0Hz

NEW! test capability
screens to a billionth of
a volt ($\text{nV}/\sqrt{\text{Hz}}$) of noise to
ensure all LSK389s produce
no burst (RTN or popcorn)
noise and meet or exceed 1/f
and broadband noise
specifications.



Contact us today
for samples and
pricing!

Linear Systems, Inc.
(510) 490-9160
sales@linearsystems.com
www.linearsystems.com

Testing and Optimizing Strategies for Conference Terminal Equipment



With most company employees working from home and the trend toward decentralized organizations and remote collaborations, webinar and conferencing services are the new way of communication. Once these communication solutions are adopted, we all have the same problems: The speech transmission quality varies. Sometimes it is so bad that good communication is not possible. But the reasons for this are complex. This article addresses how to evaluate and determine strategies for these challenges.

By
Hans W. Gierlich
(HEAD acoustics GmbH)

The more complex new applications and services for webinar and conferencing services, the more difficult testing gets. While in the past a connection in communication systems could always be subdivided in the components near-end terminal, network, and far-end terminal, this separation is not always possible for new types of services. The separation of the components in a connection provides the huge benefit of qualifying each component separately, and ideally guarantees seamless interworking of all components in a connection regardless of which device or network is used. This approach has proven its validity since the beginning of telephony networks. However, it requires the clear definition of interfaces, and it requires the use of a reference interface in testing. The reference interface is part of the test equipment and provides reference access to the device under test (DUT) for certain access technologies.

Due to non-linear time invariant (non-LTI) transmission systems found today (VoIP, VoLTE, etc.), the transmission quality may change even during a call. Therefore, a certain amount of typical network impairments (e.g., jitter and loss) should be included in the device testing of terminals. Another area is the interfacing of components (e.g., headsets or smart speakers), connected wirelessly or via USB. Similar to the tests of terminals, a reference interface is required for testing. For proprietary access technologies in modern conferencing services, a reference interface-based access may not be available. Different testing and optimization strategies have to be pursued. Here, the focus is on end-to-end quality and testing only at the acoustic interfaces.

This article describes the requirements and benefits of reference interface testing as well as end-to-end testing. Using some device examples, important test methods and results are shown.

Reference Interfaces

The definition of interfaces between interconnection points in networks, between terminals and networks, and between user equipment and terminals is key for a successful and seamless transmission of speech. Besides the definition of the electrical or the radio interface, the protocols, and so forth, the definition and knowledge of all communication-relevant parameters is important. When defining a reference interface for testing communication equipment several parameters must be precisely known.

Level

The correct knowledge of reference signal levels is required to adjust and measure levels in sending and receiving correctly. In telecommunication, levels are typically defined in dBm (0 dBm equals to 1 mW typically with 600 Ω termination impedance). In digital systems, dBm0 (represents at the reference frequency (1020 Hz) an absolute power level of L dBm measured at the transmission reference point (0 dBr point), and a level of L + x dBm measured at a point having a relative level of x dBr) is used. Although more easy to understand, the use of dBfs (dB re full scale) is not typical in telecommunication.

Frequency Response Characteristics

The linearity in frequency has to be a magnitude higher than the frequency response characteristics to be measured. This is required for signal generation as well as for signal reception. When the frequency response characteristics are determined by the speech codec used in the transmission system, the accuracy has to be within the limits given by the codec.

Codec Implementation

Ideally, the codec implementation is the reference codec implementation for a given codec. When using reference implementations, any accuracy requirement is fulfilled. If something other than reference implementations are used, the conformance of this implementation to the reference implementation has to be demonstrated.

Clock Accuracy and Clock Adjustment

The clock accuracy of a reference implementation should be high and follow the requirements for the individual interface definition. For measuring clock skew of a DUT, the reference interface should be at least 10 times higher than the clock accuracy of the DUT. For the test of devices with non-synchronized clock and frame interfaces (e.g., VoIP, VoLTE, and Bluetooth), a clock adjustment of the reference




Figure 1: Example of a complete test setup for conferencing and smart home testing




Figure 2: Typical delay drift due to unsynchronized clocks




Figure 3: Constant delay after clock drift compensation

interface is required. A clock adjustment enables users to adapt the clock of the reference interface to the clock of the DUT and eliminates the delay drift, which otherwise always occurs during testing.

Delay

The delay introduced by a reference interface must be precisely known. The delay must be stable. With the exception of a procedure for clock adjustment to the DUT, any other adaptive processing (e.g., adaptive jitter buffers), must be deactivated.

Signal Processing and Jitter Buffer Handling

A reference interface shall not include any signal processing other than speech coding and associated filtering for proper band limitation. Signal processing often found in interfaces, such as noise cancellation, echo cancellation, or signal level adjustments, must not be present. As an example, the definition of a VoIP reference interface can be found in the ETSI Guide (EG) [1].

Testing with Reference Interfaces

The most common tests in telecommunication are




Figure 4: Delay variation vs. time with jitter profile compared to clean condition (a) and MOS with jitter profile compared to clean condition (b)

About the Author

Hans W. Gierlich started his professional career in 1983 at the Institute for Communication Engineering at RWTH, Aachen. In February 1988, he received a Ph.D. in electrical engineering. In 1989, Gierlich joined HEAD acoustics GmbH in Aachen as vice president. Since 1999, he has been head of the HEAD acoustics Telecom Division and in 2014, he was appointed to the board of directors. Gierlich is mainly involved in acoustics, speech signal processing and its perceptual effects, QOS and QOE topics, measurement technology, and speech transmission quality. He is active in various standardization bodies such as ITU-T, 3GPP, GCF, IEEE, TIA, CTIA, DKE, and VDA. He is also chairman of the ETSI Technical Committee for "Speech and Multimedia Transmission Quality."

based on the use of reference interfaces. A typical setup is shown in **Figure 1**.

The tests in such a setup are generally independent of the interface used and consider all perceptual relevant parameters in the talking, listening, and communication situations as well as user behavior and environmental conditions (e.g., room reverberation and background noise). The basics and details can be found in my white paper on the subject [2]. A variety of relevant standards are listed there.

The basic measurements for telecommunication equipment (e.g., Loudness Ratings, response characteristics, distortion, and echo loss) are not considered in the following. These are well known. With the example of some devices, some advanced and important measurements are shown, which cover the user situation and user environmental conditions in a better way.

Delay and Delay Variation—Jitter Buffer Handling

Most systems today work IP-based, and therefore do not work on a synchronized clock. Devices operate on their own clocks and compensate clock deviations from interconnected devices by providing buffers, which may overrun or underrun and discard speech packets from time to time.

Unsynchronized clocks lead to delay drift between devices, which makes it impossible to determine the delay of a device without special measures. In order to reliably measure the delay of a DUT, the clock of the test equipment needs to be synchronized to the clock of the DUT. This is a function that has to be provided by the reference client. **Figure 2** and **Figure 3** show the result of clock drift compensation after determining the clock deviation by the reference client.

Besides the handling of any clock drift, the buffer in receiving in IP-based systems serves another purpose—the handling of jitter and packet loss. In non-managed networks, the packet arrival time may change during a call due to time-variant network load resulting in packet jitter.

A de-jitter buffer is used to align the packet stream and allow constant playout of the audio data. The adjustment strategy of the de-jitter buffer is crucial for speech quality. Depending on the amount of jitter, the buffer size needs to be increased in order not to lose too many packets resulting in poor listening speech quality. Increasing the jitter buffer size too much will result in too high delays and greatly impair the conversational quality [2]. A good compromise has to be found. Therefore, it is advisable to test

communication systems under jitter and loss conditions. It is possible to record such conditions in advance and insert those time-synchronous to the audio data [3], [4]. This way different implementations can be tested and optimized under identical conditions. For such a test, a series of sentences lasting typically 160 s pairs (according to Recommendation ITU-T P.501 [5]), are used as a test signal. During the 160 s, two measurements are made. For each sentence pair, the delay is calculated and the Mean Opinion Score (MOS) based on Recommendation ITU-T P.863 [6] are calculated. Both are compared to the clean channel condition. Certain limits (e.g., maximum delay increase and maximum decrease of the MOS value) under clean conditions are evaluated.

A detailed example is given in **Figure 4**. Jitter obviously increases the delay introduced by the DUT, while speech quality remains almost identical except for one sentence at about 110 s, where a packet loss coincides with a speech sample. Finding the right compromise with low delay keeping good listening speech quality based on a proper packet loss concealment in conjunction with the speech codec used is the art of design.




Figure 5: Echo and the location of its origin

Echo and Perceptual Echo Evaluation

Echo is a very annoying impairment and in contrast to any impairment in the listening situation, it is a talking-related impairment. Echo always originates from the far-end devices, and it is due to the acoustic coupling between loudspeaker and microphone as shown in principle in **Figure 5**.


Echo tests are well known in telecommunication and the underlying measure is Terminal Coupling Loss (weighted) TCLw [7]. The measurement is simple: A test signal is inserted into the DUT and the loss provided by the terminal and its echo canceller is expressed in decibels (dB) [7]. Typically a requirement of TCLw > 50 dB has to

Independent Transducer Test & Measurement • Audio Systems
Prototyping • EASE & CLF Speaker Data • Klippel Sales,
Calibration and Service • Microflown Sales & Support

WARKWYN

Now offering **MICROFLOWN**
Sales and Support

Warkwyn is proud to offer Microflown Technologies' highly innovative testing equipment for sound visualization, sound power quantification, noise control and troubleshooting. We use Microflown's acoustic particle velocity and sound intensity probes in our own labs.



KLIPPEL
Exclusive North American Sales Representative

Microflown Technologies
Sharing world wide
Sales and Support

Warkwyn.com • Info@Warkwyn.com • 612.759.4042

A²B Audio Modules

32 channels of
digital audio over
twisted pair

Accelerate Design Time:

- Development tools
- OEM modules
- I/O modules



CLOCKWORKS
Signal Processing

www.clk.works info@clk.works

+1 978 258 5402

A²B is a trademark of Analog Devices Inc.

be met in order to provide echo-free connections under all circumstances, regardless of the delay in the transmission system. However, this measurement is perceptually not correct. The talker and the listener at the near-end terminal perceives the echo more or less masked by her/his own voice. Depending on transmission delay, the echo is more or less audible. With low transmission delays, the echo is not very perceivable, and as a consequence, the loss inserted by the DUT may be less than with high delays. This effect is not covered by traditional tests.

A new procedure developed by HEAD acoustics covers this effect and takes into account the masking

Device	Room	Echo Path	Volume Control	TCL (dB)	EMOS Average
1	Anechoic	Fixed	Maximum	64	3.87
1	Office	Variable	Nominal	64	2.86
1	Office	Variable	Maximum	64	2.36
1	Anechoic	Variable	Maximum	64	1.63
2	Anechoic	Fixed	Nominal	53	2.08
3	Anechoic	Fixed	Nominal	60	3.68
2	Anechoic	Fixed	Maximum	53	2.90
3	Anechoic	Fixed	Maximum	60	3.85

Table 1: TCLw and Echo-MOS in comparison for various terminals




Figure 6: Echo signal with time varying echo path (Device 1 in an office room and maximum volume setting)




Figure 7: Test signal for attenuation during double-talk testing, consisting of competing speech samples [5]

effect of the user’s voice. The amount of annoyance caused by the echo is not expressed in decibels but in MOS focusing on the amount of annoyance on a scale from 1-5. 5 represents non-audible echo, while 1 represents a severe echo impairment. This method is called EQUEST. **Table 1** provides some results of terminal tests using both methods. Different devices were tested with different settings and for some tests a time-variant echo path simulation was used simulating, for example, user movement in a room. A rotating plate close to the device, which is exactly controlled by software during the measurement, realizes this time-variant echo path. Furthermore, an office room was used for some tests and the device volume control was adjusted to different positions.

All devices listed in Table 1 would pass the typical TCL requirement of $TCL > 50$ dB. However, EMOS shows quite different results. Under steady-state conditions, EMOS is > 3.5 , which is acceptable echo behavior. However, for some conditions, especially with time varying echo path, EMOS drops drastically. The reason is re-convergence of the echo canceller in the devices leading to short temporal echo peaks. These are just averaged out by TCL, but subjectively these are quite annoying as shown by EMOS. **Figure 6** provides a time signal of the echo loss showing these temporal echo peaks.

Conversational Performance—Double Talk

The performance of devices in double-talk situations (conversational partners talking simultaneously) is a very important performance parameter for conversational quality [2]. A double-talk test requires the simultaneous generation and evaluation of at least two test stimuli inserted from both ends of a communication system or a DUT. Typically, the main impairments result from the echo canceller in the device. The echo canceller may either not provide sufficient echo cancellation in double talk or the non-linear processor (NLP) required to remove remaining echo may introduce switching impairments. The near-end speech signal is temporarily cut off. Testing double talk requires synchronized speech test signals. Testing requires the measurement of the DUT delay in order to exactly synchronize the test signal at the DUT, in such a way that the signal overlap in the test sequence is always identical at the DUT. Such test sequences are specified in Recommendation ITU-T P.501 [5] and the appropriate analysis is described in my white paper [2].

This test is completely independent of the type of DUT. Speech samples (short and long duration) with varying overlapping are inserted from both sides (see **Figure 7**). The test is basically checking how

much attenuation is applied to the various speech segments and calculating a double-talk type [8]. The double talk tests characterization of a device range from 1, 2a, 2b, 2c, to 3 where 1 is a full double-talk implementation, 2a–2c characterize partial double-talk capability, and 3 is a non-duplex device. In a similar way, echo during double talk is evaluated. If the echo attenuation remains high during double talk, the characterization is a Type 1 device. Less sufficient implementations producing more echo are characterized accordingly from Type 2a to Type 3. The test signal is different, however.

A type of voiced sound, producing a comb-filter spectrum, is created. The signal used in the opposite direction is generated basically the same way but with a different fundamental frequency and harmonics, which do not overlap with the signal inserted in the other transmission direction. The signals are inserted simultaneously from both sides. The separation is made in the frequency domain and this way the amount of echo during double talk can be evaluated.

Multi-Talker Testing

A special type of test is the multi-talker test, which is especially important for all types of conference solutions. A multi-talker scenario is simulated by using at least two head-and-torso simulators (HATS), see Figure 1. The focus of such tests is the complete transmission of simultaneous or alternating talkers in a room for all possible talker positions. Some basic tests for this situation are described in Recommendation ITU-T P.340 [9]. One example is the concurrent talker test. The test signal is the same as the one used in the double-talk echo test. Voiced sounds with different fundamental frequencies and non-overlapping comb-filter spectra are produced by the artificial mouths of two HATS (see **Figure 8**).

The output signal G_{out} is analyzed by separating the two signals and determining the frequency dependent gain of each talker separately. When comparing this result to the individual transmission of each signal without concurrent talkers, we can determine the amount of attenuation inserted by the device in case of concurrent talkers. **Figure 9** shows a result of this measurement for one device.

For talker A (GA), it can be seen clearly that talker B is attenuated by about 8 dB to 10 dB in the concurrent talk situation, while the gain for talker B (GB) remains almost unchanged. For a conferencing device, this is typically not a preferred behavior.

End-to-End (Acoustic-to-Acoustic) Testing

With no reference client available, virtually only acoustic-to-acoustic tests are possible. Instead of a second artificial head, sometimes a Bluetooth

connection or an electrical connection to the second terminal is used. However, since the characteristics of the terminal used for assessing the DUT are unknown, such measurements are only of limited value since it is generally unknown whether effects observed are caused by the DUT or the interface used for connection.




Figure 8: Test signals for a concurrent talker test




Figure 9: Frequency dependent gain of two concurrent talkers compared to the gain inserted for each talker when talking individually (talker A—blue, talker B—red), single talk: solid lines, concurrent talk: dashed lines





Figure 10: Test setup for end-to-end testing

Figure 11: Sending response characteristic (a), overall frequency response characteristics calculated with sending and receiving measured individually (b, red) overall frequency response characteristics measured end-to-end (b, blue)



When testing from mouth-to-ear, the same problem is present. However, the benefit of this scenario is that we get the complete user experience of a complete connection. A typical test setup for this scenario is shown in **Figure 10**.

Overall Frequency Response vs. Frequency Response in Sending and Receiving

Although neither sending nor receiving frequency response characteristics can be separated, you get a complete picture of a connection by end-to-end measurements. Moreover, if the connection between the devices is free of artifacts, the result is just the combination sending and receiving response characteristics as shown in **Figure 11**.

Performing the measurement with different devices can give you an educated guess of the individual sending and receiving characteristics by evaluating different device combinations.

Compensation of Unwanted Signals in End-to-End Testing

Since most of the testing procedures are targeted to the test of one device only, it is required to remove unwanted signals in order to apply traditional testing procedures. This is needed since perceptual based methods are mostly not available for tests where the signals of both conversational partners or the signals in the talking and listening situation are present simultaneously. For such tests, the cancellation of the unwanted signal is required. A very elegant procedure about how to perform such cancellation is described in References [10], [11]. The principle is rather simple. You first measure the terminal with the signals not wanted for the analysis. In case of acoustical echo tests, the pure sidetone time signal used is measured and stored using exactly the same test sequence, which is used for the echo test. In a second step, the terminal is activated and the same measurement is performed

but now with an activated terminal. In a third step, the prerecorded sidetone signal is subtracted from the measurement performed in Step 2. This requires the exact time synchronization of both signals. The remaining echo signal is then analyzed in Step 4.

Figure 12 illustrates this testing with the example of one terminal test.

Without compensation, the signal shown at the left side of Figure 11 is measured. If this would be used (e.g., for TCL calculation) a wrong TCL of 8 dB would be measured. The measurement is completely falsified by the overlaying sidetone signal. If this is removed, the real echo signal (see the middle of Figure 12) is visible and can be used for a reliable analysis of TCL and EMOS.

Conclusions

A large variety of tests and procedures, test arrangements, and simulations of user environments and user behavior is available for optimizing the communication quality of terminals, components, and complete connections from end-to-end. Some important requirements, some test setups, and some test procedures were presented in this article. Of course, this just covers a fraction of




Figure 12: Echo signal extraction by subtracting the pre-recorded sidetone signal (right) from the acoustically recorded echo plus sidetone signal (left) leading to the echo signal only (middle)


the procedures and setups available. When fully applied, the procedures available are well suited to optimize all types of devices and connections up to the limitations given either by the physical setup of a device or by the technical possibilities built in the individual design. If applied early in the development process, these procedures and setups enable users to create more superior devices than sometimes found in the market today. In the light of "work from home," suppliers might revisit their strategies and focus more on conversational speech quality and less on the costs only. The users will certainly appreciate such solutions.

References

- [1] EG 202 425: Speech Processing, Transmission and Quality Aspects (STQ); Definition and implementation of VoIP reference point, The European Telecommunications Standards Institute (ETSI), 2007
- [2] H. W. Gierlich, "Communication quality in conferencing scenarios - Principles and optimization potential," white paper, www.optimize-audio-conferences.com
- [3] ETSI ES 202 737 – ETSI ES 202 740 series: "Speech and multimedia Transmission Quality (STQ); Transmission requirements for narrowband @wideband VoIP terminals (handset, headset & hands-free) from a QoS perspective as perceived by the user," The European Telecommunications Standards Institute (ETSI), March 2020.
- [4] 3GPP TS 26.132, "Universal Mobile Telecommunications System (UMTS); LTE; Speech and video telephony terminal acoustic test specification."
- [5] Recommendation ITU-T P.501, "Test signals for use in telephony and other speech-based applications," International Telecommunication Union (ITU)-Technical Committee (ITU-T), May 2020.

- [6] Recommendation ITU-T P.863, "Perceptual objective listening quality," International Telecommunication Union (ITU)-Technical Committee (ITU-T), March 2018.
- [7] Recommendation ITU-T G.131, "Talker echo and its control," International Telecommunication Union (ITU)-Technical Committee (ITU-T), November 2003.
- [8] Recommendation ITU-T P.502, "Objective test methods for speech communication systems using complex test signals," International Telecommunication Union (ITU)-Technical Committee (ITU-T), May 2000 and amendments 2010, 2014.
- [9] Recommendation ITU-T P.340, "Transmission characteristics and speech quality parameters of hands-free terminals," International Telecommunication Union (ITU)-Technical Committee (ITU-T), May 2000 and amendments 2014, 2019.
- [10] U. Muesch, F. Kettler, S. Bleiholder, "Applications for Time-synchronized Noise Compensation (TNC)," DAGA Conference Proceedings 2016, Aachen, Germany.
- [11] J. Reimes, "Instrumental Assessment of Near-end Speech Quality," DAGA Conference Proceedings, 2017, Kiel, Germany.

A Comprehensive Course in Speaker Design




Get your copy today!

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


cc-webshop.com

IsoThermal TUBETRAP DESIGN KIT #2


A: TubeTrap: Start with an Acoustic RLC Series Parallel Circuit



B: IsoThermal Upgrade: Add IsoThermal Matrix



C: TubeTrap + IsoThermal = IsoThermal TubeTrap



Save 10% on IsoThermals Use Coupon Code: **audioXpress10**
Acoustic Sciences Corporation | shop.acousticssciences.com

Earthworks M23R Measurement Microphone

In this review, Stuart Yaniger details his findings about the M23R Earthworks measurement microphone. The M23R is the latest omnidirectional model in Earthworks' infinitely matched series of measurement microphones. It is designed to be an affordable reference solution, delivering an IEC 61672 Class 1 frequency response before applying any equalization.

By
Stuart Yaniger
Photos by Cynthia Wenslow

With the burgeoning of easy-to-use and powerful acoustic measurement software, inexpensive measurement microphones have become wildly popular. The proliferation of electret capsules used in other applications (e.g., smart speakers and home automation) has caused the prices to go down, and quality and consistency to go up.

The typical inexpensive mic will have a batch-calibrated electret capsule, a simple source-follower interface, and a single-ended or USB output. And for much of their usage, they're more than adequate. However, they have clear limitations—dynamic range, noise floor, consistency from unit to unit, stability over time and environmental conditions, and distortion will not be inspiring. They are good enough to get some basic speaker and room measurements, but beyond that, one has to invest in a better mic.

Traditionally, high-quality instrumentation-grade measurement mics have been condenser types requiring polarization and having analog balanced outputs. And expensive—a good lab mic will typically run in the \$1,000 to 3,000 range. In the past few years though, small labs, studios, live

recordists, sound reinforcement engineers, and serious hobbyists have been offered electret-capsule mics of better quality construction, higher dynamic range interfacing, and individual calibration for intermediate prices between the popular cheapies (e.g., miniDSP's UMIK) and lab-grade condensers. For \$300 to \$600, one can get an electret-based mic with relatively low distortion and individual calibration that meets IEC 61672 Class 1 frequency response conformance at room temperature when used with the calibration file.

These "intermediate class" electret-based mics use phantom power (typically 48 V) and have professional XLR outputs, so they require a mic preamp to accommodate them as if they were condensers. Examples include the Earthworks M30 and iSEMcon's EMX-7150, the latter of which I routinely use for speaker measurement. PCB Piezotronics offers the 130-series array mics, which are likewise excellent, but they use an ICP interface (current source), which is more common in the test and instrumentation world than in audio. The Earthworks M23R is the latest entry in this intermediate market segment (see **Photo 1**).



Photo 1: The M23R is packed in a rather plain but well-padded box, with its calibration sheet, a mic clip, and a calibrator adapter.

Just to put things in context, Earthworks was started by David Blackmer as a measurement microphone company back in the late 1990s before branching out to other markets, developing proprietary acoustics and electronics technology. Blackmer's goal was to produce the most accurate measurement microphones in the market, since he felt that most of the then-existing products did not give reliable results. To this end, Earthworks has assembled a line of products, all of which have in common attention to time domain response, full individual acoustic calibration, and rugged mechanical design to survive portable and touring use.

The Product

So what differentiates this product from the other excellent offerings in its price range? There are two things that stand out. First, the M23R has an IEC 61672 Class 1 frequency response before applying equalization. Hence, for most purposes, a calibration file is not needed and the frequency response will be flat enough for measuring room acoustics, machine noises, front-of-house (FOH) sound checks, and even loudspeaker measurement.

If your application requires even tighter frequency response conformance, Earthworks offers the user a digital calibration file in CSV text format, which can be imported into most measurement software. This file is delivered via email to any requesting user after the product is registered.

Figure 1 shows the provided frequency response of the mic tested in my lab, obtained by measuring a driver on a large baffle. It is indeed quite flat in response, comfortably threading a 1 dB needle. Before any correction, is already better than the corrected response of most inexpensive test mics and easily meets Class 1 standards.

The second interesting feature is that the mics can be pre-matched on request—not by the usual method of selecting two or more mics out of a lot with relatively close responses, but using a bit more elegant and versatile method. As with most electret mics, a field-effect transform (FET)-based amplifier is built in but will not be optimal for the lowest possible distortion, noise, and the highest dynamic range. The usual circuit is a simple source follower, which works reasonably well but will not be optimal for lowest possible distortion, noise, and dynamic range.

More sophisticated circuits can be employed to improve an electret mic in all of these areas. Several examples are given in the Scott Wurcer articles listed in Resources. Earthworks uses a circuit similar to one of Wurcer's. Because of the use of a




Figure 1: The unequalized measured response for the M23R, provided by Earthworks, shows a smooth and flat response.




Figure 2: Each M23R is packed with an individual sensitivity and frequency response calibration sheet.

Photo 2: To measure sensitivity and frequency response, the mic was used to measure a monitor speaker at 1 m distance, then the results were compared to a reference mic in the same position.





Figure 3: The M23R's on-axis frequency response, equalized using its calibration file, is impressively flat, showing less than 0.1dB deviation from flat across the midrange and treble.

About the Author

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

more sophisticated circuit than a simple follower, Earthworks can match mics for frequency response and sensitivity using some simple RC trimming of the electronics.

Because of this ease of adjustment, Earthworks can do something unique with the M23R. If you want to purchase a mic that is matched to one you already have for sensitivity and frequency response, Earthworks can trim a new mic to match the older one if you provide them the serial number. No need to buy matched pairs from the get-go. And of course, with this matching flexibility, you're not limited to pairs. If you want to build up, say, a tetrahedral array for Ambisonics or a 16 mic planar array for acoustic holography, every mic can be tightly matched to the others.

One other interesting feature is the mic's sensitivity. Most 1/4" electret or condenser mics I have used have sensitivities well under 10 mV/Pa. For example, the iSEMcon mic mentioned earlier has a sensitivity of 6 to 7 mV/94 dB. This makes the noise requirements of the following mic preamp quite critical. The M23R has a nominal sensitivity of 36 mV/Pa, more typical of mics with larger diameter diaphragms. This higher sensitivity translates to a lower noise floor from the microphone, as well as effectively dropping the mic preamp's noise contribution by 20 dB or so.

The M23R in the Lab

The M23R arrived packed with calibration data (see **Figure 2**), a mic clip, and an adapter for attaching the mic to a standard 1/2" calibrator (as shown in Photo 1). The mic case's form factor is consistent with other Earthworks 1/4" mics: a smooth tapered contour between the capsule and the preamp, minimizing diffraction artifacts.

I began with some basic measurements. As usual, my measurement system comprised an Audio Precision APx525 with an Audio Precision (AP) APx1701 acoustic interface. Reference mics included the PCB Piezotronics 376A31 1/4" and 376A32 1/2" phantom power lab mics and the iSEMcon EMX-7150 electret mic, all with calibration files loaded into the APx525's software.

First, I checked the sensitivity. The provided calibration sheet listed it as 35.5 mV/Pa. Unfortunately, I no longer have a NIST-traceable SPL calibrator, but the substitution method gives nearly the same accuracy if performed carefully. The reference mic (in this case, the PCB 376A31) was set up approximately 1 m from a small monitor speaker. I played a 1 kHz sine wave adjusted to 85 dB SPL shown on the AP Monitor/Meter panel. The capsule position was measured with respect to the mic clip, then

the PCB mic was replaced in the clip with the M23R positioned so that the capsule was in the same spot that was used for the reference measurement (see **Photo 2**). The rated sensitivity was 35.5 mV/Pa, and my measurement indicated 35.8 mV/Pa (about 0.07 dB hotter), which was well within both measurement error limits and environmental variables.

Next, I checked the mic's self-noise. My noise measurements are performed by wrapping the mic in a plastic bag, sealing the bag by wrapping completely with duct tape, and burying the mummy in about 200 lbs. of a mixture of cat litter and sand (the BLANK chamber). The baseline noise from this method over a 22 kHz bandwidth, A-weighted, unwindowed (or rectangular window), is under 5 dB SPL. The M23 showed a self-noise under those conditions of 19.6 dB-A. Although this was not an extremely low number in an absolute sense, it was quieter than any other 1/4" mic I have had in my lab and slightly better than Earthworks' 20 dB-A spec. It was certainly quiet enough for typical mike positions in studios.

The M23R's frequency response was then measured versus the equalized 376A31 reference mic via the substitution method described above, with the M23R's calibration file loaded into the AP.



Photo 3: I use a near-field method to hit the mic with high SPL at low distortion.

hypexelectronics
www.hypex.nl



NCAS500MP

400 + 100W in 8 ohm

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

New family Ncore modules

NCAS500MP is the first to be released

Future members:

NCAS250MP: 200 + 50W

NCAS1000MP: 800 + 200W

For further information:

website: www.hypex.nl

mail: sales@hypex.nl

Ncore
audio power conversion

© Hypex Electronics B.V. All Rights Reserved

The test signal was a swept sine wave (Farina chirp), which was converted in software to the source's impulse response. The measured impulse response was then gated to remove room reflections, and windowed (AP Equi-ripple) before performing an FFT.

The two equalized quasi-anechoic responses, reference and M23R, were then ratioed to derive the M23 on-axis response with respect to the equalized reference. Because of the gating at about 4 ms, the response was only valid above about 250 Hz, but generally, microphone response variations for electret capsules are pretty small at low frequencies.

The response from 400 Hz to 20 kHz is shown in **Figure 3** and was quite flat, well within 0.1 dB of perfect; any deviations were well within measurement variability. The slight ripple in the response above about 4 kHz may be from a minor reflection from the mic stand, since I did not use one of the long tube adapters common in calibration labs nor did I put foam between the capsule and the mic clip. At any rate, this was excellent performance when compared to the reference, a much more expensive lab-grade metal diaphragm mic.




Figure 4: A graph of distortion vs. SPL at 1 kHz shows low distortion up to 130 dB SPL test limit and likely well beyond that.

Resources

A. Khenkin, "How Earthworks Measures Microphones," Earthworks Audio, www.earthworksaudio.com/wp-content/uploads/2012/07/how-earthworks-measures-mics.pdf.

C. Walber, "Acoustic Methods of Microphone Calibration," *Voice Coil*, September 2017, <https://audioexpress.com/article/acoustic-methods-of-microphone-calibration>.

S. Wurcer, "Low Noise Microphone Preamplifiers," Part 1, *Linear Audio*, Volume 1, www.linearaudio.net.

S. Wurcer, "Low Noise Microphone Preamplifiers," Part 2, *Linear Audio*, Volume 3, www.linearaudio.net.

Next, I measured the distortion as a function of SPL. In order to get high SPL signals, one nice trick is to find a low distortion high output midrange driver with a concave or flat cone profile, then bring the mic to within a few millimeters of the surface (see **Photo 3**). A frequency should be chosen that is high enough to avoid large excursions but low enough for the driver to still be pistonic. For the particular driver I've been using for this duty, 400 Hz to 1,000 Hz seem to work well. This near-field method has the further advantage of the direct signal overwhelming any room noise or reflections.

For this THD measurement, I placed the driver face-up on a work surface and isolate it with a chunk of open-cell foam. The driver was connected to the APx1701 and driven with a 1 kHz sine wave up to the point where my personal tolerance hit its Xmax. My results up to that 130 dB SPL point are shown in **Figure 4**. As a baseline, driver distortion was well under 0.1% at this frequency and level in the nearfield. The decently low distortion from the mic was just beginning to rise at the end, and if extrapolated, would be comfortably under 3% at 140dB SPL. This is a tribute to the design of the electronics.

One of the advantages of a small diameter capsule like this is that it's easier to achieve a good omnidirectional polar pattern. To measure this, you can use a very nice software driven carousel, which will automate the process and get great angular resolution. Lacking that, you can do the Impecunious Experimenter method of marking angles on a Lazy Susan, then placing the mic stand on it so that the mic capsule is vertically aligned with the rotation axis.

I measured gated frequency response at 30° angular increments, with each measurement being a power average between three measurements at roughly ±4° about the nominal angle plus the nominal (e.g., the 30° measurement is a power average of about 24°, 30°, and 34°). The resulting data at the spot frequencies of 1 kHz, 5 kHz, 10 kHz, and 16 kHz were then transferred to a conventional polar pattern plot (see **Figure 5**). Note that I used a more expanded radial axis scale than usual, so take this into account when comparing to polar patterns of competitive mics. The M23R really held up well to the sides, being only -2 dB down at 10 kHz and -5 dB down at 16 kHz for a 90° incidence.

At the rear of the mic, we started to see some interference effects from the mic cable at the upper two test frequencies (10 kHz and 16 kHz), so take the absolute numbers at 180° with a grain of salt. But qualitatively, they are a good indicator.

Summing Up

It's always a delight to get a product to test that

isn't quirky, does exactly what it claims, and has a clear purpose in life. The M23R performed well in all of the basics: accurate and flat response, good polar pattern, low distortion, and for its size, very low noise and high output. In all of these measures, it performed about on a par with the similarly priced Earthworks M30, though it is not specified to the same bandwidth—I strongly suspect that the guts of the two models are the same, with just a few bits rearranged in the electronics.

Comparing the M23R to the M30, what you get here is something that has flatter frequency response in the audible band right out of the box, conforming to IEC 61672 Class 1 without further equalization. It can be put into pairs or arbitrarily sized arrays with tight matching, even if new mics are bought later on to expand the array. 

Author's Acknowledgements: Megan Clifford and Taoufiq Bellaime of Earthworks really went above and beyond in helping me out, being patient with my less intelligent questions, thorough in their explanations, and working me through problems I ran into along the way.




Figure 5: The polar pattern of the M23R shows it having a very omnidirectional pattern up to very high frequencies.

ULP Apollo MCU

for Always-On Voice Command Applications

Download the Whitepaper to Learn More


ambiq
endpoint intelligence

www.ambiq.com

Getting Started with Automotive Audio Bus (Part 2)

Understanding A²B Features

By
Brewster LaMacchia
(United States)



In Part 2 of this article series Brewster LaMacchia examines how certain optimizations for automotive applications must be accounted for when targeting broader audio applications. He also includes a look at analog phantom power and how A²B provides that feature in a digital system.

For applications where the nodes will be AC (line) powered or have common connections from shielding, there are considerations for grounding that can be different from automotive applications and understanding that difference is important for system design.

Phantom Power

Before looking at the electrical interface for A²B data, we'll review how A²B phantom power works as it's different from phantom power systems with which audio engineers may be familiar [1]. We'll look at two common analog systems for phantom power to help understand where A²B is different and what issues that could lead to in your own designs, particularly for non-automotive use cases.

The simplest phantom power circuit is found on PCs and similar low-end applications that use an electret-type microphone that includes a junction field-effect transistor (JFET) in a common source configuration. In this design, half the amplifier (the drain resistor) is in the equipment that the microphone plugs into, as shown in **Figure 1**. The coupling capacitor blocks DC. Typical values for R_x are 2.2 k Ω and 5 V for the supply. There are several variations on this basic design—some using a mini-phone TRS jack to split the power from the signal connection.

In a single-ended system, the questions of grounding and shielding exist, particularly if there's no metallic chassis to connect the shield to. As drawn in Figure 1, the shield is being used as the return, which is typical but not desirable.

More familiar to many readers is the 48 V phantom power standardized by IEC 61938 [1]. **Figure 2** illustrates the basic concept of phantom power supplied over balanced lines. The positive voltage is applied to both the positive (+) and negative (–) signals, the supply return current is carried by the ground wire. Since the two sides are matched, the common mode noise rejection of the input device's transformer removes power supply noise.

The two resistors R_p , which are 6.8 k Ω , need to be well matched to maintain the overall common mode noise rejection, along with the output transformer on the microphone side having matched halves. The two supply currents cancel in the microphone transformer, as otherwise the core would have a DC bias and increased distortion.

Figure 2 was drawn with the requirements of AES-48 [2] included, which means the shield is used to connect directly to the chassis ground and the "signal" ground functions as the supply return path. Not shown is the connection between the chassis and signal ground.

Understanding Figure 1 and Figure 2 will aid in the discussion in the next section on how A²B phantom power works and how it can and cannot be used in certain applications.

A²B Phantom Power

Having been refreshed on phantom power and grounding issues with analog audio, we'll look at the unique aspects of A²B's phantom power design. The A²B master node can supply 300 mA, which at typical A²B voltages of 5 V to 8 V is enough power for operating a microphone or even a small DSP with ADC and DAC. As with all things that seem too good to be true, this may not be as helpful as it first seems.

A²B systems are designed to be rugged—it's very expensive to have your vehicle serviced. Cable opens, shorts to ground or battery (24 V in some vehicle systems) must not release the magic smoke from the ICs. Having the system indicate where the failure occurred is important to keep costs down as ripping apart an entire cabin to find a failed connector after the vehicle was purchased is not going to make a good customer impression.

The A²B features that are helpful for automotive applications must be accounted for in other applications. **Figure 3** provides detail on how a locally powered node supplies phantom power to one or more bus powered slaves in an A²B network. Figure 3 also illustrates the master node as supplying phantom power—any locally powered slave node can also serve that role. It's allowed for an entire A²B network to be locally powered at each node.

Before looking in more detail at the system level phantom power issues in the simplified diagram of **Figure 4**, we need to understand how the currently available A²B ICs work. We start by contrasting the requirements for A²B with those of analog audio




Figure 1: Basic electret microphone interface provides phantom power for a simple microphone amplifier




Figure 2: Simplified functional schematic for phantom power for balanced lines used with microphones. In many cases, the transformers are replaced by capacitors to block DC. Not shown is protection circuits to deal with shorts or transients.




Figure 3: Simplified view of A²B phantom power between a locally powered node and a bus powered node. Master nodes are always locally powered; slave nodes can also be locally powered to overcome the total bus power limit (300 mA) from a local powered node. We label the master node as a transceiver because it does send data downstream and receive data upstream, but unlike a slave node does not pass data further upstream.

phantom power schemes previously described.

The first obvious difference between Figure 2 (phantom power for microphones) and Figure 3 is that there are only two wires available, as with the system shown in Figure 1. One wire becomes the positive phantom supply and the other becomes the negative. In the electret microphone case, only a few hundred microamps (μA) of current is needed—here we want to supply up to 300 mA.

With one line as positive and the other as negative, the system must consider what happens if a connector is plugged in flipped (any keyed connector can be overcome with sufficient force) as that is sure to release the magic smoke from ICs when the power is backward.

In the case of phantom-powered analog microphones, the signal bandwidth goes from a few Hertz out to 50 kHz. Not introducing noise or distortion is critical. By contrast, the A²B system uses Manchester encoding for a basic data rate around 50 Mbps and a bandwidth out to 200 MHz or more. There will be lower-frequency components present compared to continuous Manchester encoding as there's no clock once the data packet ends, though those dead times can be thought of as DC when considering splitting power from data. Digital system distortion doesn't apply in the way it would in an analog system. Noise would be a problem but it's not difficult to filter out supply noise in the frequency range of the A²B signal.

In Figure 3, the two coupling capacitors, acting as a high-pass filter, block the DC from the data ports (A_N , A_P , B_N , B_P) of the transceivers. The two inductors, acting as a low-pass filter, block the LVDS

signal from the power supply connections. Typical Analog Devices, Inc. (ADI) suggested values for the capacitors are 10 nF and 3.3 μH for the inductors, which are physically very small parts. By contrast at audio frequencies, capacitors in the dozens of microfarad (μF) range are needed for AC coupling, and a choke to block audio from the supply would be impractically large.

The choke to isolate the DC supply serves two purposes. It keeps the signal out of the supply and offers a high impedance to the signal so that it is not attenuated. The supply has low impedance at high frequencies. In an A²B system at the 50 MHz primary frequency, the impedance of the 3.3 μH inductor is about 900 Ω , which is sufficiently above the 100 Ω system impedance to be easy to accommodate (we will look at that in more detail later). To be completely correct, the 100 Ω is the differential impedance, the single-ended impedance from the data signal to ground is 50 Ω . Care must be taken to not pick a choke with a large parasitic capacitance that lowers the impedance at the high frequency end of the bandwidth used by A²B.

Locally Powered Node

The node supplying power has features to enable detection of a range of bus faults that might damage the hardware. This is achieved in several ways. One is the SENSE signal to the IC. The P-type metal-oxide-semiconductor (PMOS) high-side switch allows the node to control the application of the phantom bus power. If an incorrect voltage is detected, the IC can shut the PMOS switch off. This sense capability is also used to determine if a downstream node is locally

powered by having a small diode/capacitor/resistor network (not shown in the bus powered slave of Figure 3) across the power lines on the slave node. The SENSE signal can also detect the presence of voltage on the B_P line when there shouldn't be any.

The diodes prevent pathological wiring faults from propagating. For A²B applications where all nodes are locally powered, the circuit is theoretically not needed if cable fault detection is dropped. The default software stacks expect to probe for cable faults so such an application would entail some additional software work and a deep understanding of the dozens of control/status registers in the transceiver IC.

Faults with the B_N line can be detected because the return line from the downstream node does not directly connect to the local ground signal on the locally powered node (the Master node shown in Figure 3). Instead it connects to the V_{SSN} pin, which connects through an internal N-type metal-oxide-semiconductor (NMOS) switch to ground.

In automotive systems, power and multiple data connections are combined on to a single connector so protecting against shorts to +12 V is advantageous. Outside of automotive use, audio connectors between

boxes will generally not be carrying power so the short to positive supply fault is of less concern. Finding shorted connector pins to each other or to ground and open wires can certainly make life easier for quickly isolating problems.

This robustness to opens and shorts comes at a cost—Figure 4 should make that more obvious. The positive supply line will drop some voltage at each node due to the saturation voltage of the PMOS FET and the parasitic resistance of the cable and connectors. V_+ and V_{GND} are the power supply connections, and it's reasonable to assume that V_{GND} connects to a common earth ground via local supplies or common shielding connections between audio boxes.

Following the procedures in Analog Devices' datasheet [3], an example of the three nodes shown in Figure 4 with 5 m total of 22 AWG, the resistance of the wire is about 0.26 Ω and the four connectors (assuming the DuraClik used on the EVM boards) equal 0.04 Ω for a total of about 0.3 Ω . The largest impedance comes from the V_{SSN} to V_{SS} NMOS transistor at about 1.2 Ω , from the second node this is about 2.7 Ω total. If two slaves were drawing 50 to 100 mA each, then we could




Figure 4: Simplified system view of phantom power and visualization of the multiple ground voltage levels. Connecting each node's Local GND to a common (earth) ground would bypass the A²B protection features.

ALL ELECTRONICS CORPORATION

Thousands of Electronic Parts Available Today!

LEDs • CONNECTORS • RELAYS
SOLENOIDS • FANS • ENCLOSURES
MOTORS • WHEELS • MAGNETS
PC BOARDS • POWER SUPPLIES
SWITCHES • LIGHTS • BATTERIES
and many more items...

We have what you need for your next project.

www.allelectronics.com
Discount Prices • Fast Shipping

SpiralAir™ 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

approximate the voltage at Local GND 2 as around 200 mV relative to the master node's GND signal (V_m in Figure 4).

If the slave nodes are all microphones, then the noise on the Local GND will not affect operation. If a slave node has ADCs or DACs, they will be referenced to the Local GND. If those connect to external devices that have a (system) ground connection, a ground loop is created, leading to increased noise. The A²B fault detection is also compromised by the alternate ground path.




Figure 5: Illustration of an application using an A²B backbone with branch nodes to avoid daisy chaining devices like microphones and headphones. The backbone to branch links can be up to 15 m in length. Total system cable length from the farthest node to the master is 40 m.

References

- [1] IEC 61938:2018, "Multimedia Systems—Guide to the recommended characteristics of analogue interfaces to achieve interoperability," International Electrotechnical Commission, <https://iectest.iec.ch>
- [2] AES 48-2019, "AES standard on interconnections—Grounding and EMC practices—Shields of connectors in audio equipment containing active circuitry," Audio Engineering Society (AES), www.aes.org
- [3] Adapted from AD2428 datasheet, Revision B, January 2020, Analog Devices, Inc. (ADI), AD2420(W)-AD2426(W)-AD2427(W)-AD2428(W)-AD2429(W), <http://analog.com>

Resources

A²B information, Analog Devices, Inc., www.analog.com/en/applications/technology/a2b-audio-bus.html

In Figure 4, we see that the available voltage at each node drops as the positive supply line will also have some parasitic resistance, though the PMOS switch has much lower impedance than the NMOS V_{SSN} one in the transceiver IC. Analog Devices offers a spreadsheet to aid in calculating the power node power and voltage needs (see Resources).

Another complication is possible if shielded cable is used. Shielded cable is avoided in automotive use because of the extra weight and cost. In the automotive case, there are tweaks that can be made as the node configurations are fixed and precise wire routes are known. That will not be the case in other applications.

While automotive electromagnetic compatibility (EMC) is more stringent than Federal Communications Commission (FCC)/CE Mark requirements, audio applications need to avoid creating even small amounts of EMC as at some point low-level analog signals will exist and be very susceptible to small signals (think radio stations miles away coming through a mic pre-amp).

Connecting shields to local grounds at each node creates an alternative ground path and a way for the system to radiate. The cable shield should connect to the metallic chassis (if present) and for where the local ground DC value is not the system 0 V value, connect to the local ground via a small value capacitor, possibly in parallel with a resistance of a few mega-ohms.

This capacitor provides an AC ground—if everything is at the same potential it can't radiate. Minor circuit imbalances can degrade this approach. If the shield starts carrying high-frequency signals (with A²B primarily 50 MHz and with harmonics stretching out to the 200+ MHz range), the relatively long lengths of interconnect cable act like an antenna. These problems are not new to high-speed digital system designers, but they may not have been experienced when working with analog-only audio systems.

A²B System Application

There are two obvious applications for A²B where phantom power is very helpful and the system grounding issues won't be problematic like they might be between boxes of audio electronics. Specifically microphone and headphone applications do not involve interconnections to multiple endpoints and don't present a grounding problem.

A single microphone is perhaps overkill to throw at A²B, however, the use of VR audio recording via high-order Ambisonics is an excellent fit for sending multiple channels of microphone audio through a single twisted pair and not worrying about noise pickup or a massive multichannel XLR snake.

A conventional stereo headphone over A²B is easy enough, but overlooks the new capability that A²B can offer. With A²B, multiple channels of audio can be sent broadcast-style and a small DSP that the user controls can select levels and adjustments to create a unique local listening experience from a common set of audio streams routed to multiple headphone endpoints.

Combining both the microphone and the headphone concepts opens up interesting possibilities for field recording, measurement, and monitoring. For both of these applications A²B can be used as a point-to-point link to some head end box (up to 15 m cable length), as to run it in a daisy chain two cables (four wires) are needed so that the chain can continue to the next device. Alternately, we can take advantage of an A²B capability that hasn't been discussed yet: the ability to have a branch off of the main A²B backbone (see Figure 5).

There are some topology limitations that are beyond the scope of this introductory article, but six nodes in the backbone is a reasonable assumption. Latency will increase to around four sample periods to the end points (100 μ sec) and the backbone nodes will probably need to be locally powered. The total

distance from the farthest node can't exceed 40 m, but 5 m spacing between the backbone nodes and 5 m cables to microphone arrays or locally controlled headphone endpoints fit a range of applications (e.g., specialized conference and media rooms, and small studio spaces for recording or streaming).

Summary

Part 2 of this article provided a closer look into A²B system design, bus topology, and phantom power features. Next month we'll examine A²B data transport and how to interface between peripherals and the A²B bus using parts made by Analog Devices, Inc. 

About the Author

Brewster LaMacchia started in signal processing working on military satellite communications systems. Since then he has worked on a wide range of signal processing systems, including a number of years on video processing and surround sound audio for consumer home theaters. Recently, he started Clockworks Signal Processing to develop open-source hardware for multichannel signal processing. He also dedicates time to public astronomy events to expose people to science and critical thinking. In the past 10 years, he's helped more than 7,000 people look through a telescope for the first time.

Celebrating 40 Years! ACOustics begins with ACO™



Measurement
Microphones
Polarized and Electret

Community/Industrial
Noise Monitors/Alarms
Indoors and Out



Microphone Systems
Standalone
IEPE Powered
Phantom Power
Custom



Windcreens
3 to 18 inch & Custom
Weather/UV Resistant



ACO Pacific, Inc

www.acopacific.com sales@acopacific.com

Tel: +1-650-595-8588

DO IT YOUR WAY

Cables/Panels/Boxes

WITH REDCO AUDIO'S

"Build Your Own" web pages

Full control for design, editing and more.

AND WE'LL BUILD IT JUST THE WAY YOU WANT!




REDCO AUDIO

www.redco.com 1-800-572-7280

Ultra-Low Distortion Microphone Buffer Amplifier (Part 2)

Simulations and Measurements



By
Michel Nieuwenhuizen
(The Netherlands)




Figure 1: This is the complete circuit including bias oscillator.

In Part 1 of this article, Michel Nieuwenhuizen presented a new microphone buffer amplifier that avoids most of the problems found in existing buffers. The circuit is shown in **Figure 1**. In this second part, he shows the results of the simulations and the measurements. He also compares the performance of the circuit with that found in many Chinese condenser microphones (the “Schoeps” based circuit).

For the simulations, I used MicroCap 12, which is an integrated schematic editor and mixed analog/digital simulator that provides an interactive sketch and simulate environment for electronics engineers (see Reference 1). I tried a number of well-known field-effect transistors (FETs): 2SK117, 2SK170, 2N3819, J201, BF862, and others. In all cases the impedance convertor was loaded with a 5 kΩ resistor.

The bandwidth of the new circuit is 7 Hz to 100 kHz ± 0.5 dB (using 10G resistors in the input circuit), as shown in **Figure 2**. Although I measured both total harmonic distortion (THD) and intermodulation distortion (IMD), I concentrated mainly on the IMD tests since IMD is far more important for complex sounds than THD. All the simulations and the measurements were done with an input of 1 kHz 100 mVeff + 100Hz 25 mVeff. This corresponds with an audio level of approximately




Figure 2: Bandwidth of the new circuit




Figure 3: Reference circuit adjusted for maximum allowable input level

100 dB to 120 dB in a real microphone (depending on the sensitivity of the capsule). This is about the level of a full orchestra some meters away.

The reference circuit has quite some distortion but with proper adjustment (see Part 1, October 2020) the second harmonic components can be brought down significantly. However, simulations show that a change in ambient temperature of 10°C changes the suppression of the second harmonics with 10 dB to 20 dB. So it is not clear how effective this compensation is in practical use (see **Figure 3** and **Figure 4**). The IMD simulations of the new circuit are a lot better (see **Figure 5**).

I found most of the JFET models in MicroCap were wrong. So what I show here is my best effort of selecting FET types and adjusting their parameters to better fit the published specifications. The performance differences between the reference circuit and the new circuit are quite clear. Second harmonics are at least 10 dB down and higher harmonics at least 20 dB, even when compared with the optimally adjusted reference circuit.

Building the Microphone

I only have approximately 67 mm x 15 mm space inside the microphone housing for the electronics.




Figure 4: Reference circuit adjusted for optimal second-harmonic distortion




Figure 5: New circuit with cascode FETs




Photo 1: The PCB with components (apart from the T2 and XLR socket)




Figure 6: Circuit for matching n-channel JFETs



Photo 2: The complete new circuit before I assembled the entire microphone




Figure 7: The test setup for distortion and noise measurements

The PCB measures 57 mm × 15 mm, and the bias oscillator coil is another 10 mm × 15 mm. I used surface-mount technology (SMT) components where possible, without hurting the audio performance (see **Photo 1**).

As previously mentioned, the FETs have to be paired. I used a very simple circuit for this (see **Figure 6**). Pairing is done by selecting FETs that have the same V_{dg} . The beta of the output transistors must also be matched to within 1%. **Photo 2** shows the completed new circuit prior to assembling the microphone.

Measurements

I used a PC with a decent sound card (the M-Audio Delta 2496) as an analog-to-digital converter (ADC) and a digital-to-analog converter (DAC). I then generated test tones with Audition as 44.1 kHz/32-bit floating point wave files. The output level of the sound card is set to 100 mV 1000 Hz + 25 mV 100 Hz (same levels as during the simulations). The audio files are played back with VLC media player and I use Adobe Audition 1.0 as my FFT analyzer. In order to be sure that the ADC is not overloaded, I set the input amplitude of the sound card at -6 dB by changing the gain of the microphone preamp. **Figure 7** shows my test setup for the distortion and the noise measurements.

Noise

The blue track is the noise with the input resistors shorted, red is with 1G resistors and yellow is with 10G resistors in the input circuitry. At 20 Hz there is a difference of about 12 dB between zero and 10G and another 12 dB between 10G and 1G. At higher frequencies the difference becomes smaller and it disappears at approximately 2 kHz (see **Figure 8**).

We can compare this with the noise of a complete microphone (i.e., the Brüel & Kjær 4165 capsule) in a very quiet room. Below 1000 Hz, the capsule noise dominates. Above that, frequency the electronics play a role (see **Figure 9**).

Distortion

I measured the IMD of both the reference circuit and the new circuit. The levels were the same as during the simulations: 1000 Hz 100mVeff + 100Hz 25 mVeff. IMD distortion is well visible both around 1kHz and 2 kHz and the THD (mainly second harmonic) is about 0.08% (see **Figure 10** and **Figure 11**).

When the reference circuit is adjusted for minimum distortion, second-order harmonic distortion drops dramatically. Higher harmonics




Figure 8: Noise level with 0 Ω (green), 1G (red) and 10G (orange)




Figure 9: Noise with the Brüel & Kjær 4165 capsule




Figure 10: IMD of the reference circuit at 100 mVeff, with the input adjusted for maximum output




Figure 11: IMD of the reference circuit at 100 mVeff, with the input adjusted for optimal distortion




Figure 12: IMD 1000 Hz 100 mVeff + 100Hz 25 mVeff, with the new circuit




Figure 13: Extended circuit with extra transistors for lower distortion

About the Author

Michel Nieuwenhuizen studied electronics at the Eindhoven Technical University in The Netherlands. He then worked for 27 years with Philips Electronics on audio and video processing, resulting in a number of patents on video image improvement. After his retirement, he concentrated on audio projects such as a software package for automation of local radio stations called Caban and microphone techniques.




Figure 14: IMD of the circuit with extra transistors

remain the same. The performance of the new circuit is considerably better. The distortion is not visible above the noise (see **Figure 12**).

The problem with a low input impedance of the preamplifier as previously discussed was confirmed. For optimum results, the input impedance should be at least 5 kΩ. The maximum output level is 8 Vpp or 2.5 Veff. Much more than most preamplifiers can handle.




Figure 15: Measurement of IMD for the new circuit with "super FET"

Conclusion

The new impedance converter shows an extremely low distortion and sufficient low noise. The circuit draws a very constant current from the phantom power avoiding "compression artefacts." The maximum output voltage is in excess of 1 Vpp, which is sufficient for almost all normal use. The circuit is flexible and mid frequency and high frequency adjustments are possible.

Further ideas

Is even lower distortion possible? Looking for an even better performance, I did some tests with extra transistors (T5 and T6) in the differential amplifier. This increases the gain and lowers the distortion further. Two extra capacitors are needed for stability. The extra resistors regulate the current through the FETs. **Figure 13** shows the extended circuit with extra transistors for lower distortion.

The simulations give an almost perfect result, but the measurements could not confirm this since all the distortion products are now hidden by the noise of the circuit (see **Figure 14**). If the simulation is to be believed, the IMD is now about 0.0001%!

The measurement shown in **Figure 15** confirms the simulations for as far as possible (ignore the high-frequency disturbances, they are caused by external interference). The noise obscures any distortion. 

References

[1] Cyril Bateman's Capacitor Sound articles, Linear Audio website, <https://linearaudio.nl/cyril-batemans-capacitor-sound-articles>

Sources

Brüel & Kjær 4165 capsule

Brüel & Kjær Sound & Vibration Measurement | www.bksv.com

M-Audio Delta 2496 PCI audio card

M-Audio | <https://m-audio.com>



ELECTRONIC TRANSFORMERS

- Original Equipment
- Replacement
- Amateur

TRIAD MAGNETICS, INCORPORATED

460 HARLEY KNOX BLVD., PERRIS, CA 92571

www.TriadMagnetics.com

THE ULTIMATE^{IN} AFFORDABLE SIGNAL PROCESSING



The Dayton Audio DSP-408 is a feature-rich 4 input, 8 output digital signal processor that was designed to give you the tools to contour your audio system to fit your tastes. The DSP-408 allows the user to choose the application of their device. With two methods of powering the DSP, this processor is the perfect addition to your home or car audio system.

Mobile App



Windows Program



The most unique feature of the DSP-408 is the ability to operate in a vehicle or home audio system right out of the box. This adaptation allows users to utilize a single model of DSP for all their audio processing needs.

The ability to control this DSP is provided via a windows application or via Bluetooth connection to an Android or iOS device when paired with the Dayton Audio DSP-BT4.0 (#230-502). The volume and preset selection for the DSP can also be controlled by the optional wired remote (#230-504).

For many individuals, the pursuit of perfect sound quality is a never-ending quest. When utilizing all the features the DSP-408 offers, the ability to manipulate and improve the sound quality of your audio system is simply astonishing. Connect to your DSP-408 and get to work!

Visit parts-express.com.
Custom OEM inquiries welcome.

For optimal performance, check out these accessories (sold separately).

**Dayton Audio
DSP-BT4.0 Bluetooth
Data and Streaming
USB Interface for DSP-408**



#230-502

**Dayton Audio
DSP-RC Wired
Remote Control
for DSP-408**



#230-504

**SAVE \$12
off \$100 order!**

Code: AXMQ42020
Expires 11/30/2020

DAYTONAUDIO
Your path to great sound.™

Distributed By:

**PARTS
EXPRESS**
AUDIO COMPONENTS & SOLUTIONS

Tel: 800-338-0531
725 Pleasant Valley Dr.
Springboro, OH 45066