www.audioxpress.com

audioxpress

ADVANCING THE EVOLUTION OF AUDIO TECHNOLOGY

Market Update
Improving Wireless Audio
By J. Martins

Digital Signal Processing
How to Choose the Right Signal Processing Technique
By Dave Betts

Solutions for Building Voice User Interfaces at the Network Edge
By Matthew Mitschang

Show Report
Parts Express 2022 Speaker Design Competition
By Thomas Perazella

Hollow-State Electronics
Voltage Regulator Tubes
By Richard Honeycutt

You Can DIY!
Making a DC Power Supply with Audible Benefits Part 2—Real Improvements
By Frans de Wit

Lundahl Transformers
How High Performance Transformers Are Made
By Scott Dorsey

Audio Electronics
Power Transformer Parameters, Selection, and Testing
By Chuck Hansen

$9.00US $10.00CAN

NOVEMBER 2022

ANATOMY OF A 300B

PRESENTED BY *Western Electric*®

New, patent-pending graphene-coated anode improves thermal emissivity, electron affinity, and eliminates secondary emission

Molybdenum grid creates negative coefficient of resistance

Filament core's proprietary oxide coating allows for 40,000 hour lifetime and superb impulse response

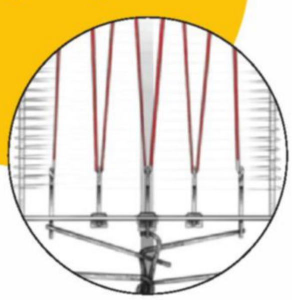

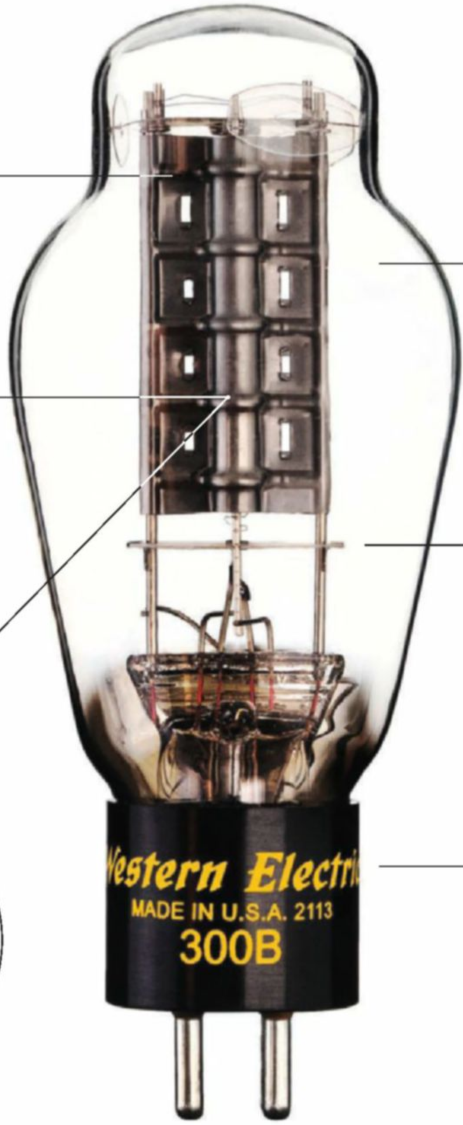

Ultra-high vacuum finishing at 2×10^{-7} Torr, made possible with state-of-the-art turbomolecular pumping systems

Proprietary mica coating increases dielectric strength, optimizes higher voltage operation

Quality guaranteed with a five-year extended warranty

There's more to it than cathode, anode, and grid. Each Western Electric 300B is assembled at our new facility by the hands of our highly trained staff using the highest quality materials and modern technology—all while honoring the original 1938 specifications and tooling.

What's the result? What's in it for the listener? Unmatched manufacturing quality and the most powerful, full-bodied sound there is. When you're through compromising, order a matched pair online or from Authorized Dealers worldwide.

AVAILABLE NOW

westernelectric.com/300b

WESTERN ELECTRIC • MAKER OF ELECTRON TUBES & HIGH FIDELITY • 201 W GORDON AVE • ROSSVILLE, GA • USA • +1 404.352.2000

MANUFACTURING / DESIGNING TRANSFORMERS & ENCLOSURES SINCE 1917

Hammond Manufacturing | www.hammondmfg.com

November 2022 ISSN 1548-0628

www.audioxpress.com

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

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

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

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

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

Advertising rates and terms available on request.

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

Wireless and DSP Evolution

The Market Update article about Wireless Audio technologies featured in this issue revisits some of the extraordinary changes that have taken place in that space. But to tell all the stories and detail the technical progress, more than one article is needed.

It's particularly impressive how things have evolved in less than 10 years, and also the number of companies and technologies that have been left behind without ever gaining any consumer awareness. The distinctive factors and key features that distinguish winning approaches in wireless audio are not always obvious. Audio quality, the main argument for most Wi-Fi based solutions, explains why the technology remained at the core of the home audio segment. While convenience explains in great part why Bluetooth technology has evolved into the behemoth we know today—benefiting also from the massive size of the mobile market, where it originated and evolved. Bluetooth has meanwhile expanded to the home audio segment and is about to become even more recognized. For this issue, there was a lot to say about the evolution of Bluetooth, and next we will expand on solutions for wireless audio based on Wi-Fi and other technologies that deliver the much desired "wired-quality without the wires" and well beyond.

Bluetooth today doesn't address major requirements that in my opinion are decisive for mainstream consumer audio applications: audio quality, low latency, and synchronized multichannel. Any of those is a major challenge and it's been very hard to optimize products for two of those factors, without affecting negatively one of the three.

Most likely one of the features that consumers will quickly embrace in the renovated Bluetooth LE Audio specifications is Audio Sharing. If this becomes easily available and robust, it will make Bluetooth even more popular. But as the history of wireless audio technologies tells us, that's not easy to do based on standardized technologies.

For example, Multipoint connectivity, which Apple users have long enjoyed and is essentially described in the Bluetooth specifications, Audio Sharing, or the ability for a user to share its wireless stream with another user, is also something that Apple has made possible for at least 3 years—effectively introduced with AirPods and an iOS update without much fanfare. For some reason it's not available from a macOS device (e.g., an iMac or MacBook). But Apple pulled a few extra tricks to make the whole experience seamless when using an iOS device (iPhone/iPad) as source. Starting with the way it offers an option to "share audio" when the user is about to press Play, and includes the suggestion for the second user to approach its AirPods to the streaming source and guide the user to "pair" the second pair of AirPods (and it also works with Beats). Apple does it with a combination of standard protocols and a peer-to-peer link to achieve a very robust Audio Sharing connection, that currently the Bluetooth LE Audio specification describes but is not yet generally available. When it does, I am certain that wirelessly sharing audio from a source to two users—or more than two users with Auracast broadcast—will become something not only popular, but something that will again reinforce consumers perception of the technology's convenience.

In the sequel to the Market Update in this issue, I am discussing wireless audio for home theater 5.1 and immersive formats, which is still a challenge and it shouldn't be. Also, I will explore what is being done to address the bandwidth limitation of Bluetooth and enable real "lossless," or even uncompressed, high-resolution 24-bit audio.

This issue of *audioXpress* also features two important articles that reveal—each one in a different perspective—how important digital signal processing technologies are in today's audio industry. When I invited submissions for this issue focused on DSP, I was not surprised to see that the articles received all had voice in mind. The two examples included in this issue reveal how voice processing applications powered by artificial intelligence are determining what is possible today and in the future.

Having multiple synchronized 24-bit/96kHz audio channels with the lowest latency being streamed wirelessly should also happen very soon. What will trigger that evolution, we still don't know.

J. Martins
Editor-in-Chief

OUR NETWORK

audioxpress | circuit cellar | LIS Loudspeaker Industry Sourcebook | VOICE COIL

SUPPORTING COMPANIES

Contents

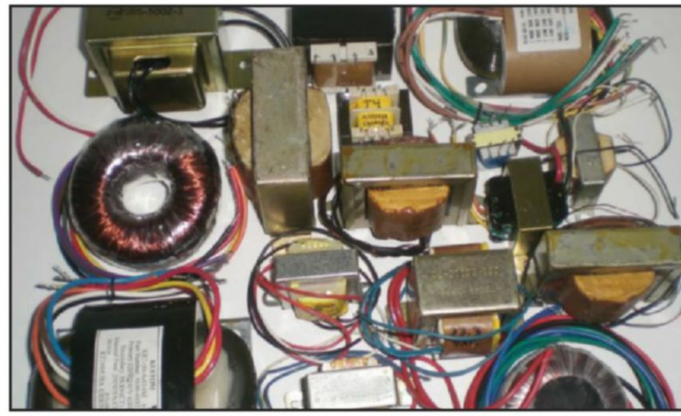

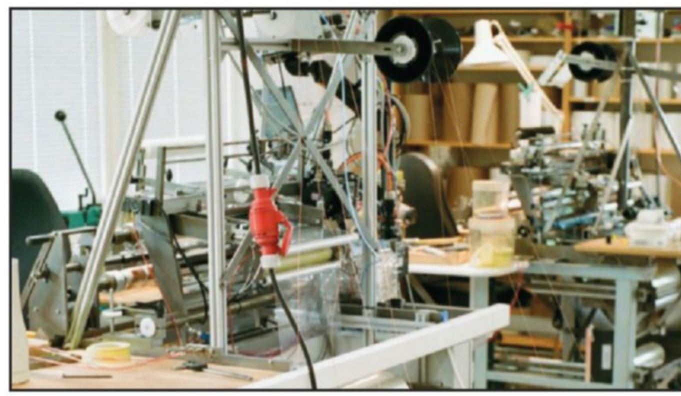

Features

Volume 53 – No. 11

November 2022

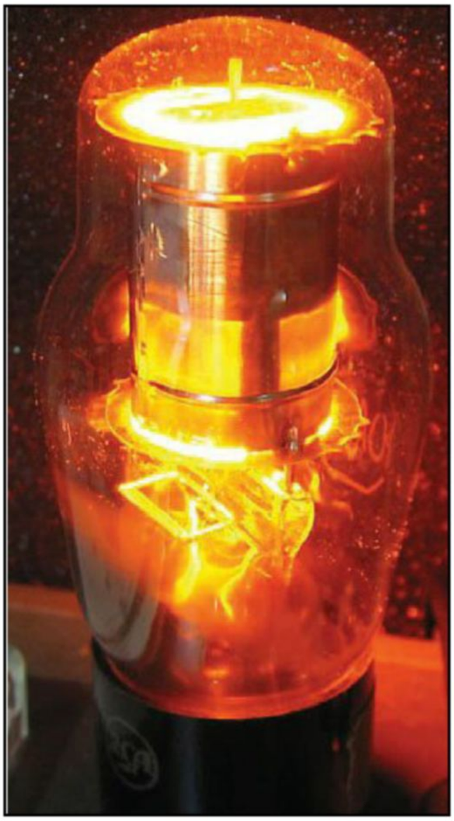

Market Update—Residential Integration

Columns

Departments

Websites

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

@audioxp_editor audioxpresscommunity
linkedin.com/company/audioxpress
audioxpress

Parts Express 2022 Speaker Design Competition

Photo 1: The Parts Express showroom front

Earlier this year Thomas Perazella received a call letting him know that the Speaker Design Competition sponsored by Parts Express was going to be held again, after a two-year hiatus due to the pandemic, and that he was again invited to be one of the three judges. This amazing event, held August 5-6, enables DIY speaker builders from all over the country to bring their prized speakers to Springboro, OH, to be appraised, compared and enjoyed. Some of the amazing speakers presented at this year's design competition are described here.

By
Thomas Perazella

First, a little background. In prior years, the Speaker Design Competition was part of an event called Midwest Audio Fest. In addition to the speaker competition, there was an auto sound competition, a customer used equipment flea market, and a tent sale put on by Parts Express and a number of vendors. Since only the speaker competition was being held this year, before it started, I decided to visit the Parts Express catalog showroom (**Photo 1**). On display are a huge number of items, and if you don't see it on display, it is probably available from their huge warehouse since the retail store is attached to it (**Photo 2**).

Photo 2: Inside the Part Express showroom

Regarding the competition, knowing that there is a tremendous range of potential speaker designs, four categories were set up to make the competition better matched to the characteristics of speakers from the very basic to the most complex. The categories are just slightly different than before, primarily due to the effects of inflation. Three of the four categories require a passive iteration, that is only one amplifier is required for operation and no DSP is allowed.

The Categories

The "entry" level category has a primary requirement that all the drivers for a stereo pair must have a total retail price of no more than $300. Formerly, that limit was $200. This does not take into account all the other parts including crossover components, enclosure materials, finishes, terminals, ports, wiring, and of course labor. When I first started judging the competition I had very low expectations for the results in this category. Well, I was wrong. After hearing some of the entries in this category I was reminded how easily preconceived notions can be shattered. In fact, the last time the competition was held, an entry in this category got the highest number of total points when aggregated from the three judges.

The Over $300 category has the same requirements as the entry level one with the exception that the total driver cost can be anywhere from $300 and up. Classically, this has been one

Photo 3: The 2022 Speaker Competition reception area

of the best categories with high-quality materials and sound designs being employed.

The Dayton Audio category allows any number or price of drivers as long as they are of the Dayton Audio brand, which is the Parts Express house brand. At first this might seem to restrict choices, but the last time I looked at the Parts Express website I found more than 150 drivers ranging from a 21" subwoofer with 21mm Xmax to planar magnetic and AMT tweeters with lots in between.

The Unlimited category is basically what it says. Any drivers and designs can be used. The only limitation is that it must be active and that if the same basic design but without an active implementation is built it cannot be entered into any other category.

In each category, there are three winners for First, Second, and Third place. Winners in each receive a trophy plus a Parts Express certificate for purchases in the amounts of $250, $100, and $50, respectively. During the meet and greet, there were also raffles for some nifty merchandise.

The Venue

Changes were made this year to the venue based on feedback received by Parts Express. The judging area used to be a large room with some material hung around the area. The speakers were toward the front of the room with a table for the judges in a prime listening position and behind them rows of chairs for all the contestants and other observers. Although the audience members were usually well behaved, maintaining quiet during the auditions was difficult at times due to a natural reaction of people to make comments. In addition, the size of the room was not indicative of the "average" room most people would use for their listening.

This year, it was held in a different venue that allowed for four main areas. Upon entry there was a large open area with chairs, tables, and large video displays (**Photo 3**). It was used for initial reception and registration as well as the place where the Meet and Greet was held. It also had a list of all the names of all the speakers that were entered displayed on a large monitor.

A large room with tables to place small speakers and floor space for floor mounted speakers was used to receive the speakers as they arrived and to stage them for judging. At that point they all had an identifying placard with the name of the entry, the number, the category, and the entrant's name.

The actual room used for judging was more representative of the size of a typical home listening room. It had one door but no windows. The ceiling had typical office-type ceiling tiles and material was hung along the walls to tame early reflections without being overly damped. Reference tape marks were placed on the floor to locate the speakers and a table was placed in the prime listening position at which the judges sat. A small number of chairs were placed behind the judges for the entrant and their helpers. In **Photo 4**, the judges are shown from left to right, Tom Perazella, Javad Shadzi, and Jerry McNutt. Information about them can be obtained from the Speaker Design Competition website.

After judging, the speakers were moved to the "jam room," which was a larger room with rows of chairs set up where everyone could listen to the speakers just judged. In that room, any music that the entrant or others desired to hear could be played with a 10-minute total time limit.

The Process

On Friday afternoon, the entrants began to arrive. They registered their entries in the appropriate categories and were given project numbers for each. A person could enter more than one speaker if desired and some did. From there, they brought their speakers into the staging room where the discussions between entrants started.

Later in the afternoon a Meet and Greet session was held where Parts Expressed provided pizza and soft drinks for everyone. It was a great chance for the entrants, the judges, and the folks from Parts Express to meet, ask questions, and share stories.

After the Meet and Greet, the judging began with the Unlimited category being first up. The reason for doing the unlimited first is to allow for the extra setup time needed in this category. Since DSP and multiple amplifiers are often used, setup does become more involved. It is not unusual for someone to need an adapter of some sort to interface with the Parts Express front end. For each of the categories, the things that remain consistent are the placement of the speakers and the distance from the speakers to judges. The tracks to be used in the evaluations are those selected by the judges for that category. A list of the tracks with times and categories is shown in **Figure 1**.

Six criteria were used to judge the speakers. They were Clarity, Craftsmanship, Dynamic Range, Originality/Design, Soundstage/Imaging, and Tonal Balance. Numerical rankings ranged from 1 to 10. Ratings of 1-2 indicate Needs Work, 3-4 Below Average, 5-6 Average, 7-8 Above Average, and 9-10 Excellent. In addition to the places for numeric scores, each tally sheet had a large area for judge's comments. As each audition was completed, the sheets were collected by Parts Express personnel for inclusion into a master scoring spread sheet.

The Results

This year, there were 42 entries. They were divided into the categories as follows: Under $300 – 13, Over $300 – 8, Dayton Audio – 8, and Unlimited – 13. Again, it was amazing to see some of the original thinking and craftsmanship that went into the entries. As far as awards, not everyone could be a winner, but in terms of the experience of being there and seeing what others were doing, everyone benefitted.

Under $300

In the under $300 category, it was really amazing to hear the quality of sound that some of the entrants achieved with a very difficult limitation of all drivers for a stereo pair retailing for no more than $300. Last year, that category (it was under $200 then) had an entry with the highest number of points of any category. It showed two things to me. First there are lots of people who have the knowledge and experience to choose the right drivers and second to implement them in a well performing speaker.

Photo 4: Judges from left: Tom Perazella, Javad Shadzi, and Jerry McNutt

The winner in this category was called "Duetta" from Paul Kittinger. In Italian, the word Duetto, close to duetta, means a musical piece for two instruments or singers. And sing, the Duettas did. They had a shape that resembled an inverted triangle that was 40" tall, 10" wide from top to bottom and 13 ¼" deep at the top. A tapered support on the rear provided the stability needed. Two drivers were used per speaker—a 7" Dayton paper cone mid woofer and a Dayton 1 1/8" fabric dome tweeter. The enclosure was a tapered and mass-loaded transmission line (**Photo 5**).

Second place in the under $300 category was not only a well-performing speaker but a rather unique design that brought laughs from the judges. It was a pair called "Carl & Stuart" made by Nick Santorineos

Artist	Title	Category	Start	End
Jimmy Buffett	Diamond as Big as the Ritz	Dayton Audio	0:25	1:27
Deadmau5	Seeya	Dayton Audio	0:15	1:25
Jacqui Naylor	Ain't No Sunshine	Dayton Audio	2:50	4:00
Bruno Mars	Uptown Funk	Open Unlimited	0:00	1:01
Sub Urban	Cradles	Open Unlimited	0:00	1:12
Toto	I Will Remember	Open Unlimited	2:41	3:41
Charole	Andante Sostenuto	Over $300	7:30	8:29
Nickel Creek	Destination	Over $300	0:00	1:22
Habib Koite, Eric Bibb	Blowin' In The Wind	Over $300	0:42	1:42
Natalie Merchant	The Peppery Man	Under $300	1:30	2:30
James Taylor	Line 'Em Up	Under $300	0:00	1:00
Marion Hill	Differently	Under $300	0:00	1:19

Figure 1: List of sound tracks used for judging

Photo 5: Duetta by Paul Kittinger

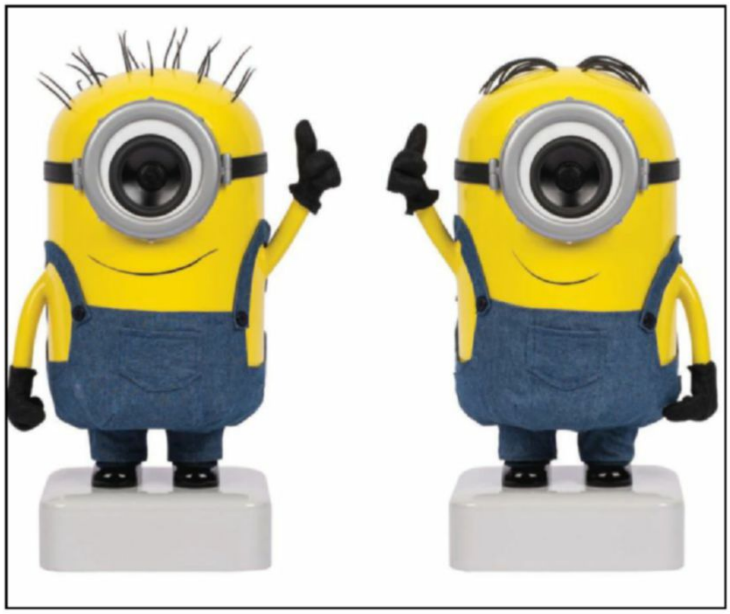

Photo 6: Carl & Stuart by Nick Santorineos

Photo 7: Rear view of Carl showing the woofer port

that were modeled after the Minion characters of *Despicable Me* movie fame. The speaker named Carl was made to look like the Minion Carl with spikey hair, who is described as cheerful and easy to get along with. He also liked to sing along "be-doo, be-doo, be-doo." An appropriate name for a speaker. The speaker named Stuart had the appropriate split hairdo as the movie character. On both speakers, the hair was made by rubbing rubber on a sander to make round "hairs." Stuart is funny, playful and likes to play the guitar and ukulele (**Photo 6**). Both speakers had a coaxial mid/tweeter made by Nick as the "eye" In addition to looking comical, what got the judges laughing was the port arrangement and placement for the bass driver, which was downward firing. When the judges asked Nick where he placed the port, he replied at the lower rear, of course (**Photo 7**). A set of denim pants was made for each speaker with a flap covering the rear port. When music was played, the flap moved back and forth in time with the bass notes. It was a very welcome comic relief in the hectic schedule of the judges. This entry also took the "Fan Favorite" award given by vote of the other contestants.

Third place in this category went to an entry named "LCR mains" by John Hollander. It is an MTM design that got high marks for tonal balance, clarity, and imaging (**Photo 8**).

Over $300

Typically, some of the best entries are in this category, and this year was no exception. Although somewhat conventional in electrical design requiring a passive approach, novel physical designs and a wider choice of premium drivers can result in exceptional performance. In many cases, the entries that I have heard in this category could rival, if not best many of the "high end" commercial speakers.

The first-place winner exemplifies the advanced thinking and construction that is seen in some of these projects. The name of this entry is "Waveguide Omni" by Dan Neubecker. For a long time, there have been conflicting demands to provide the greatest sense of image detail and placement provided by monopole speakers with the soundstage capabilities of omnipole speakers. Think of a small stand-mounted monitor compared to a Walsh driver design used, for example, in Ohm speakers that radiates equally in all directions. Without going into all the details provided by the entrant, suffice it to say various design elements combine to provide a "Di-cardioid"

About the Author

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

Photo 8: LCR Mains by John Hollander

Photo 9: Waveguide Omni by Dan Neubecker

Photo 10: Nighthawks by Adam Malito

Photo 11: Diffractorama by Bill Schwefel

horizontal pattern and a cardioid pattern for the midrange and tweeter through the crossover range of the midrange and tweeter (**Photo 9**). The goal of the latter is to help reduce interaction with the front wall allowing closer placement to that boundary. The resulting sound was exceptional resulting in high marks from all the judges.

In second place was the "Nighthawks" by Adam Malito. They are a three way using a Dayton 8" woofer in a ported chamber, a dipole mounted Tang Band 5" mid, and a Fountek horn-loaded tweeter. The front plate is the front of the woofer enclosure and extends as a single piece for the mid and tweeter. What was most striking about the appearance of the design is the alternating strips of oak and aluminum that make up the front plate (**Photo 10**). It got uniformly high marks in all parameters.

The "Diffractorama" by Bill Schwefel focused on minimizing diffraction problems. A tall rectangular enclosure handles the bass duties, but the most interesting feature was an asymmetrical curved upper mount for the mid and tweeter (**Photo 11**). Smooth curves can assist in minimizing frequency anomalies caused by the pressure changes that can occur when sound waves pass sharp edges.

Dayton Audio

Previously I mentioned how preconceived notions can be blown away by a particular speaker design. The first place winner in this category was a prime example. The "Bantams" by Tom Zarbo are very small two ways with a volume of about a half cubic foot using a 3-½" mid woofer and AMT tweeter (**Photo 12**). I can tell you that they punched way above their size and although the bass would not shake the floor, the overall bass was amazing considering the size of the woofer, passive radiator, and box volume. The response was smooth and the imaging great.

The "Mini Lx521" by John Hollander took second place. The design was inspired by the Linkwitz Lx521 but had an interesting difference for the

Photo 12: Bantams by Tom Zarbo

bass enclosure. To increase the volume, the tall rectangular hollow stand was open to the lower of the two bass drivers. Having the upper baffle and bass enclosure natural wood and the stand blue further enhanced the impression of a classically small speaker on a tall stand (**Photo 13**).

Photo 13: Mini Lx521 by John Hollander

Photo 14: Plumdingers by Keith Etheredge

Photo 15: Baffle-less-ness by Charlie Laub

Rounding out third place were the "Plumdingers" by Keith Etheredge (**Photo 14**). Looking at the color of the enclosure, you can see the derivation of the name. The color was carried through to the base of the stands and complemented not only the black drivers, but the phase plugs of this nicely performing MTM design. Clarity of this entry was very good.

Open Unlimited

A design by Charlie Laub called "Baffle-less-ness" garnered first place in this category. It was a four-way system with an intriguing combination of the tweeter and midrange drivers hanging from a wire frame with no baffles and a bass driver mounted on a half baffle and the frame (**Photo 15**). Low bass was provided by a sealed box sub. Extensive DSP was used for crossover, EQ, and time alignment functions. The baffle less implementation certainly eliminated any boxy sound. Another goal was to not have a heavy bass presence. The 12" driver in the subs is one of the Ultimax series that I have previously tested in the 18" version. That series is capable of prodigious bass and his implementation nicely integrated the bass into the rest of the system for very good balance.

Another unconventional by intriguing design was the second place winner, "T3-0 Omnipresence" from Julian Franke. A three-way design in a very unconventional housing, a metal bridge was used to center two different sized wooden spheres in front of the tweeter and midrange drivers to increase the dispersion. The fit and finish were first class. The light wood finish contrasted well with the black metal bridge and the metal feet (**Photo 16**).

Third place was "Tubular Belle" by Jack Putti. It featured a large cylindrical ported bass cabinet with a downward firing woofer, an upward firing midrange with a waveguide above it and at the top a bare AMT tweeter that had provisions to swivel in the vertical direct for optimum adjustment when at the seated position. Another unusual design that looked great (**Photo 17**).

Other Notable Designs

When you have as many outstanding entries as we had, it was very difficult to pick three winners in each category. Doing so was like choosing a favorite child. The majority of the entries would stack up very well in both design, craftsmanship, and performance to expensive commercial speakers. I would like to show a few other examples of this high level of results. They are detailed in alphabetical sequence of their names.

"Bella Sonus Model 1" from Michael Hadjinian was a tall freestanding design with a curved back


Photo 16: T3-0 Omnipresence by Julian Franke


Photo 17: Tubular Belle by Jack Putti


Photo 18: Bella Sonus Model 1 by Michael Hadjinian

and a B&W-like treatment of the tweeter. Finish on the wood and stand was very good (**Photo 18**).

"Bitches Brew Live Edge Dipole" from Perry Marshall as the name states is a large dipole configuration that was named after a Miles Davis album. It is a three-way system made from book matched slabs of spalted birch with a natural finish (**Photo 19**).

"Over Easy Eggs" from Clay Allison were a three-way design with a spiral low-frequency waveguide made in a mirror imaged pair. The egg shape was similar to what I used for the mid and high section of my reference speakers as an aid to eliminate edge diffraction that can cause frequency response anomalies (**Photo 20**).

"Peerless Uncubed" from Shawn Kipka was a three way that had an intriguing asymmetrical shape and finish (**Photo 21**). Having built asymmetrical enclosures, I realize that the cutting angles and assembly are quite challenging.

T-ALPHA 1 from Julian Franke was another outstanding example of expert woodworking both in the speakers with its many layered cabinet construction that was mirrored in custom stands (**Photo 22**).

"Titans" from Eric Woodring was an interesting combination of a pyramidal-shaped bass cabinet with the driver firing upward. Above that was


Photo 19: Bitches Brew Live Edge Dipole by Perry Marshall


Photo 20: Over Easy Eggs by Clay Allison

Photo 21: Peerless Uncubed by Shawn Kipka

platform that held a dipole mounted midrange and tweeter. Again, the quality of the woodworking and other pieces was great (**Photo 23**).

"The TS265" from Eric Opett was a very high-tech-looking design with a white enclosure contrasting black drivers sitting on top of a custom stand using extruded aluminum posts and white top and bottom pieces. Vertical blue LED strips in the stand added to the look (**Photo 24**). The photo unfortunately does not show the stand lighting.

The Bottom Line

Putting together an event like this is a herculean effort. It is hard to extend too many thanks to the contestants who, as you can see from the photos, spend huge amounts of time and in many cases their own money to create some exquisite speakers. The folks at Parts Express consistently show their dedication to the industry by having a large number of their employees involved in the preparations including promotions, which are both time consuming and expensive. They charge no entry fees and also provide perks such as food and drinks to all the contestants. The judges also do a lot of preparation work deciding what test cuts are appropriate for each speaker category and then listening to all the cuts from each judge before the competition. Personally, I reviewed the cuts several times in the weeks leading up to the competition, culminating in a session the day before I left for the event lasting over three hours reviewing the 12 one-minute cuts to be sure I heard all the details of the pieces.

For a judge, the most difficult part of this competition is choosing a few "winners" from among all the outstanding entries. You know that being critical of someone's speaker is going to be received poorly by some. But you have to be as objective and honest as you can, knowing that as an individual, your ranking of the different speaker characteristics may be at odds with others. That is the reason for having three judges who all have significant experience with speakers. The score sheets, which are given to the entrants, have comments about the characteristics that were felt to be strong or weak.

Overall, the Speaker Design Competition hosted by Parts Express is one of the most valuable events available to the DIY speaker community. If you have never participated, you should consider it. Having the ability to meet in person with a lot of other very dedicated speaker builders, exchange information at all levels, and have your results critiqued is, as Martha Stewart would say, "a good thing."

Photo 22: T-ALPHA by Julian Franke

Photo 23: Titans by Eric Woodring

Photo 24: TS265 by Eric Opett


StreamAmp High-End Audio Solution
A complete solution to build active high-end streaming speakers with outstanding sound, designed with premium components and integrated streaming services.
• StreamSDK based software
• Spotify Connect / TIDAL Connect
• AirPlay 2 / Roon
• Airable Radio & Podcasts
• Wi-Fi & Ethernet connection
• iOS / Android control app
• Hypex NCORE amplifier
• Illusonic audio processing
• Digital audio link with cable & / or WiSA
• Pre-certified & compliant
• Designed & manufactured in Europe
StreamUnlimited is the global system integrator for solutions in:
• Immersive audio
• Music Streaming
• Voice assistant integration
• IoT interoperability
StreamUnlimited
www.streamunlimited.com
Contact us: StreamAmp@StreamUnlimited.com

Market Update — Wireless Audio

Improving Wireless Audio

Part 1 — The Growing Pains

By
J. Martins
(Editor-in-Chief)

Wireless audio evolution still can take place and is in no way constrained by the recognized limitations of current Bluetooth technology—Classic and LE Audio. This article discusses the efforts and the most promising platforms and technologies for wireless audio.

This article is about wireless audio transmission technologies. It is not about streaming audio formats and codecs, even though it's impossible to discuss wireless audio without mentioning them. Basically this is intended as an update to our original series of detailed Wireless Audio articles published in 2014, when those technologies were actually in their infancy. For this report we are focusing exclusively on the core consumer wireless audio technologies.

It might seem that we had everything covered at that time, but basically what we had then, and in most parts we still have now, are building blocks that don't fit very well. The market was naturally offered a mix of standard and open technologies, with lots of patented and proprietary IP that allowed some sort of compatibility—the term often used—but not interoperability.

The history of wireless audio is still very much a work in progress and this report mainly intends to place the multiple options for progress in perspective—essential for wireless audio to offer wired-quality sound without the wires.

After all, in 2014 we had wireless audio systems available on the market for a while, although most used proprietary technology while sharing the available unlicensed 2.4GHz spectrum of the original Industrial, Scientific, and Medical (ISM) radio bands. That part of the spectrum was quickly becoming saturated with the popularity of many home devices, including garage door openers, and baby monitors, while the electromagnetic interference also increased in home environments. The more robust short-range schemes operating in that environment were naturally prevailing, while more sensitive operations such as audio streaming, which usually require robust transmission over sensitive distances (>10m), were still somewhat unreliable. Particularly with multiple synchronized receivers, as Sonos was already doing for multiroom configurations—and the reason why Sonos still promoted LAN-cabled connectivity by that time.

The WLAN Road

The Sonos approach was based on standard Wi-Fi networks but using its proprietary SonosNet protocol that it developed for wired LAN connections, which had the advantage of supporting both direct streaming from Internet sources and files stored on the local network. The system could also stream in multiple source formats, including FLAC (up to 16-bit/44.1kHz) and ALAC, while allowing the mentioned direct access to streaming services over Wi-Fi, even before Spotify Connect.

In 2014, there were still many companies attempting to promote their own approaches to wireless audio streaming, including Imagination Technologies with its Caskeid technology supporting synchronized wireless multiroom audio streaming over Wi-Fi+Bluetooth. Part of that technology would end up in Apple's hands, which hired most of Imagination's employees and still licenses many of its patents to this day.

And there was also Qualcomm, which proposed its own AllPlay platform for wireless whole-home audio over Wi-Fi, then aiming to compete with Apple's AirPlay, not so much with Sonos. Qualcomm would eventually focus on the mobile industry and consequently on Bluetooth technology with the acquisition of CSR at the end of 2014 (more on this ahead).

Other companies were more keen to emulate Sonos' approach, before realizing that this would be an environment with many expensive barriers to surpass due to the fact that many of the fundamental building blocks had long been protected with patents.

One of the very few companies that managed to navigate that environment successfully, even if subsequently forced to settle IP-related disputes with Sonos through a settlement in 2020, was Lenbrook Industries and its Bluesound/BluOS systems. When originally accused by Sonos of patent infringement, Lenbrook highlighted the fact that its “high resolution audio capabilities substantively differentiates Lenbrook’s products from those of Sonos and many other of Sonos’ actual direct competitors.” But that was beyond the point of the robust Sonos IP portfolio, which has meanwhile been successfully defended even against giants such as Google.

Effectively, among the many companies that attempted different approaches to wireless whole-home audio solutions, using purely Wi-Fi or Bluetooth+Wi-Fi as a foundation, Bluesound was the only one that originally approached it with 24-bit digital audio streaming and native support for high-resolution audio files using 802.11n Wi-Fi. Bluesound betted correctly that the evolution in Wi-Fi and WLAN technologies in general would enable consistency of high-quality streaming, even if originally limited. A prediction that is now awarding Bluesound a unique market position.

AirPlay Ecosystem

Of course, in 2014 users of Apple equipment were already used to the power and convenience of AirPlay wireless streaming over Wi-Fi, and Apple had started to license third-party devices. But no matter how good it was, AirPlay was just another proprietary wireless communication protocol, then restricted by Apple to be licensed as a third-party software component technology.

Like SonosNet, AirPlay was designed to run on Wi-Fi or Ethernet. Apple has had a pivotal role in promoting Wi-Fi since 1999 when it introduced support in its original iBook laptops and quickly created an ecosystem of AirPort hardware products, which created a foundation to approach wireless audio confidently. And even though it had started to open that ecosystem and was licensing AirPlay to other manufacturers, Apple never hesitated to go first and stride away from standard specifications, whenever if felt the benefits outweighed the risks. That’s what it did when it introduced support for Wi-Fi Direct allowing devices to connect without a LAN. Basically, Wi-Fi Direct is a peer-to-peer wireless connection, much like Bluetooth, but Apple quickly evolved the approach into a proprietary MultipeerConnectivity feature to create a mesh network between iOS devices. But we digress...

The advantage of the original AirPlay protocol stack and its wired LAN foundation is that it used UDP for streaming audio, based on the Real-Time Streaming Protocol. This allowed two-channel audio streams to be transported using the Apple Lossless codec (ALAC 44.1kHz) encrypted with AES to a receiver with the appropriate key. The trade-off was that the mechanism required packet-buffering, creating a latency that made synchronization of more channels problematic. Apple compensated with many other technologies that allowed for robustness over the primitive Wi-Fi systems of the time.

In 2017, Apple addressed those latency and synchronization limitations with AirPlay 2, expanding on the original AirPlay technology to introduce the ability to stream music to multiple

An estimated 374 million Bluetooth speakers are expected to ship in 2022, with 425 million forecasted annually by 2026, as projected by the Bluetooth SIG. As overall demand and the desire for more flexibility and mobility increases, the adoption of Bluetooth technology in portable or home speakers will also continue to expand.

In 2021, Sonos introduced Sonos Roam, an ultra-portable Bluetooth speaker built to deliver the best possible sound for its size, both at home and outdoors with adaptable features. It can be fully connected to a Sonos system on Wi-Fi at home and automatically switch to Bluetooth when the user is on the go.

Apple propelled the concept of wireless audio streaming and even multiroom with its AirPlay technology over Wi-Fi. Officially discontinued in April of 2018, AirPort Express units are still highly sought after to stream music to connected powered speakers or home audio systems, and its connectivity also supports uses as an Ethernet hub, or connecting a printer, simple to use and configure. Although now discontinued, Apple has provided a firmware update that allows it to be used with AirPlay 2 for streaming audio.

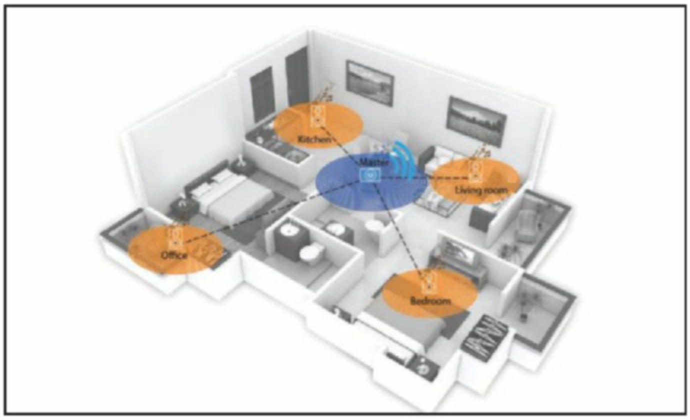

Wireless audio solutions for the home always struggled with the challenges of multiple synchronized channels. Point-to-multipoint applications for multiroom systems over a wireless LAN are achievable but streaming audio requires buffering, which in turn causes latency. After more than 20 years, full home audio wireless systems remain complicated concepts for the average consumer.

A brand of the Lenbrook Group, Bluesound has consistently expanded the range of products and the technology for its wireless multiroom systems. With an ecosystem powered by its own operating system, BluOS, Bluesound is the only system to support 24-bit/192kHz uncompressed High Resolution Audio on every player in the network, and benefits from numerous integrations and an expanded range of products, for both home and professional integration.

Another example of a typical "wireless audio" integrated system, the Bowers & Wilkins Zeppelin continues to be updated and re-imagined as the technology evolves. First introduced in 2007 as the ultimate iPod dock, the B&W Zeppelin was relaunched in 2021 with a redesigned core now supporting updated wireless streaming protocols over Wi-Fi and Bluetooth, including AirPlay 2, Spotify Connect, and aptX Adaptive.

AirPlay 2-compatible speakers simultaneously, creating a true multiroom streaming system that kept the lossless audio qualities (now finally recognized by everyone as "an improvement").

By 2014, when we published a complete series of articles on wireless audio technologies, we addressed the then-emerging topic of the transition from wireless LAN to wireless Personal Audio Networks (PAN), which was already bringing Bluetooth into the limelight. But Wi-Fi was effectively the only technology platform for the audio industry to address home systems as even Apple acknowledged. By then Wi-Fi 802.11n was established, and 802.11ac (now Wi-Fi 5) products started to appear, bringing Multiple-Input and Multiple-Output (MIMO) support and the use of MIMO-OFDM (Orthogonal Frequency Division Multiplexing), and allowing consistent data rates in crowded RF environments, essential for audio streaming.

Wi-Fi 6 and 6E (802.11ax) with dual band 2.4/5GHz and triple-band 2.4/5/6GHz, respectively, have now improved this foundation for wireless audio applications even more. In comparison, in 2014 Bluetooth was in its 4.1 infancy and it wasn't even much of an option for home audio, even though it was already being explored to create control mechanisms and discovery.

Bluetooth and Mobile

The reasons why one technology would be preferred by consumers was not completely apparent, as even Apple failed to predict. In 2014 it was clear that the iPhone (the iPhone 6 was launched that year) was going to disrupt everything, while smartphones sales were quickly becoming massive compared with any class of products. Wireless audio on Bluetooth and Wi-Fi had already surpassed the basic quality thresholds and the introduction of Bluetooth 4.2 that same year was a pivotal moment for wireless audio. Consumers were increasingly adopting wireless audio for the convenience and not worrying so much about reliability. Coupled with the consumer's perception of "quality," essentially associated with the popularity of streaming music services on mobile devices, created a turning point that would propel Bluetooth over any other wireless audio platform, invading even the home, first with portable Bluetooth speakers, followed by smart speakers. Even Apple failed to anticipate this, when it didn't consider Bluetooth audio for its HomePod speaker, leading to its quick demise, no matter its pioneering merits.

After all, because consumers were able to stream music from the Internet to any mobile device on a mobile network, they obviously expected that same "stream" to continue from the same mobile device through a wireless home network directly to speakers, hi-fi systems, or wireless headphones. But wireless audio at home, and specially the wireless link to speakers, was purely Wi-Fi only at that time.

Sonos started the "revolution" in consumer perception. It was globally promoted by Apple's AirPlay, and reached new levels with services such as Spotify Connect, which enabled streaming content directly from a wireless speaker or an AV receiver. But even when using Spotify Connect, consumers still selected the track or playlist on the mobile device, which meant that the association (and indirectly with Bluetooth, used for control) was a lasting one. In every home, whoever set the Wi-Fi and configured the home audio system understood the role of Wi-Fi and the quality and reliability implications. The other members of the family never did (and still don't), and always associated the experience with Bluetooth because they used their personal devices—the smartphone or even tablet—to control it.

But even before Bluetooth became commonly associated with streaming audio, there were also other wireless technologies that users associated with mobile freedom, on headphones for example. One of those technologies was the Kleer wireless transmission, which Sennheiser and Sony had long adopted for wireless headphones, anticipating what would be possible with Bluetooth.

Until 2014, Kleer was a clear wireless audio alternative. SMSC acquired Kleer Semiconductor (originally from Cupertino, CA) in 2010 after the company had already established important relationships with major consumer electronic brands. In 2012, SMSC was acquired by Microchip, which started licensing the technology rebranded as Kleernet. The Kleernet platform enabled a high-quality, uncompressed, low-latency Bluetooth alternative for wireless distribution of digital content to headphones, speakers, and other audio devices. It supports triple-band transmission on 2.4GHz, 5GHz, and 6GHz and allows users to simultaneously stream a signal to four wireless audio receivers. The technology is still available and offers a completely dedicated hardware and software solution that works only with Kleernet-compatible products. Likewise, Microchip offered its JukeBlox Wi-Fi connected audio platform, originally developed by BridgeCo, as another wireless audio solution that was simpler to implement.

But the Bluetooth growing popularity and growing consumer association was unstoppable. We already mentioned Qualcomm's AllPlay solution, originally intended to compete as a higher quality alternative over Wi-Fi. Yet, Qualcomm soon dismissed those efforts with the acquisition of CSR plc (formerly Cambridge Silicon Radio), mainly because of its strong Bluetooth audio technology portfolio, which also included the aptX family of codecs. That part of the story was recently shared in a dedicated article about the "Story of aptX" (*audioXpress*, October 2022). What is less known is that apart from its leading Bluetooth and audio coding and processing technologies, CSR also offered its VibeHub Networked Audio Platform with SyncLock technology to support a wide array of connectivity options, including Bluetooth, Wi-Fi, or Ethernet, with extensible networking capabilities. In 2015, after the CSR acquisition by Qualcomm was finalized, the VibeHub Wi-Fi development solutions for wireless audio, originally promoted as making it far easier for device manufacturers to deliver multiroom networked audio around the home, faded to oblivion, as Qualcomm focused all its efforts on Bluetooth.

The Bluetooth Question

Fast forward a few years, and in December 7, 2016 the Bluetooth Special Interest Group (SIG), the trade association that oversees Bluetooth technology, officially adopted Bluetooth 5. Key updates to Bluetooth 5 included longer range, faster speed, and larger broadcast message capacity, as well as improved interoperability and coexistence with other wireless technologies.

As the Bluetooth SIG stated in 2016, "The increased speed of Bluetooth 5 lays the groundwork for the next generation of Bluetooth audio, and the increased range will deliver reliable connections

Wireless audio applications for the home frequently bridge personal audio technologies and concepts. An example is the Sony SRS-NS7 Wireless Neckband Personal Theater, intended to provide an immersive home theater experience, with sound tailored for personal listening based on Sony's unique 360 Spatial Sound Personalizer technology. To render the Dolby Atmos source content, the new neckband needs to be paired with the WLA-NS7 Bluetooth wireless transmitter that helps minimize audio delay.

According to the Bluetooth SIG, Bluetooth LE Audio was designed to improve wireless audio performance via the use of the LC3 codec, add Bluetooth support for hearing aids, and introduce Auracast broadcast audio and audio sharing, two new Bluetooth capabilities that will enhance the way users engage with others and in social environments.

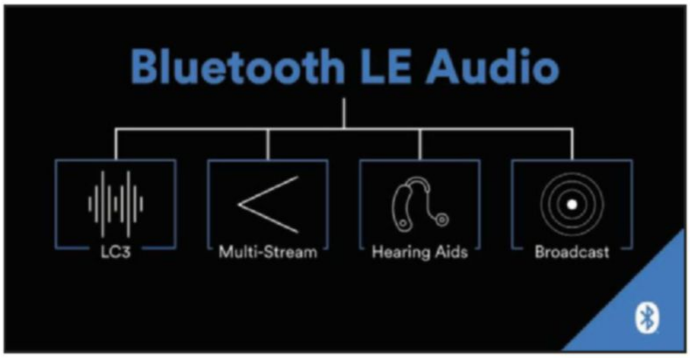

This was the original plan for Bluetooth LE Audio in January 2020, and we now know it was just a plan. Multistream is a big technical advancement. Broadcast is now Auracast.

that make full-home, building, and outdoor use cases a reality. All this while maintaining its low-energy functionality and flexibility for developers to meet the needs of their device or application." That was a correct prediction, as Bluetooth entered and dominated wireless home audio.

Audio Sharing (now Auracast) is probably the main thing that consumers will notice in the next generation of LE Audio wireless audio products, even though the hand-off between users listening to different streams to share the same stream is not yet very clear how it will happen from minimalistic TWS earbuds with very little control options.

Bluetooth Audio Codecs	Maximum Bitrate kbps	Theoretical Latency ms
SBC	345	250
AAC	320	100
aptX	384	150
aptX HD	576	430
aptX LL	352	50
LDAC	990	480
16-bit 48kHz uncompressed	1500	–

Bluetooth audio codecs maximum bitrate and theoretical latency comparison

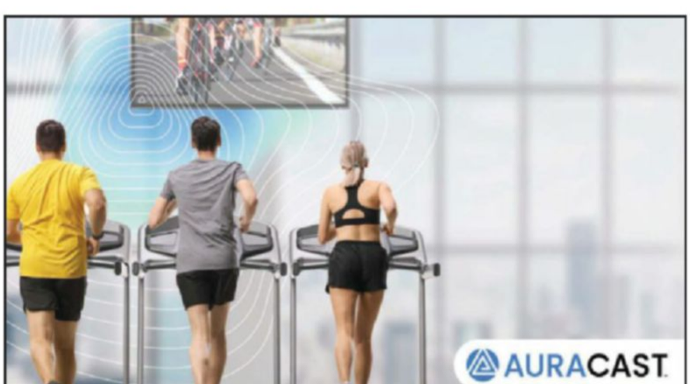

The Bluetooth SIG announcement of Auracast broadcast audio adds an unprecedented new capability provided in LE Audio for sharing audio with nearby listeners. This includes the ability to invite others to simultaneously share a personal audio stream as well as for venues to deliver audio directly to earbuds when watching a muted restaurant/bar or gym television, enjoying a movie at the cinema with high-quality sound coming directly into advanced hearing aids, or receiving an airport public announcement broadcasted privately in the preferred language. As a new brand, Auracast makes it possible to promote applications that consumers don't associate with what is possible today with Bluetooth technology.

Bluetooth 5 was essentially aimed at robustness and improving on the technology's success, unexpected in the case of wireless audio streaming where the default audio codec was always clearly labeled as inefficient, particularly under the bandwidth constraints. In fact, the use of the extremely efficient AAC codec over Bluetooth, available in Apple, Microsoft, and Android devices was the reason why the audio quality over Bluetooth suddenly seemed to be "acceptable," leading many to still assert today that no further improvements would be noticeable.

Less obvious was the fact that the Bluetooth 5 specifications pointed at capabilities that offered a "mobile" wireless LAN platform for the Industrial IoT and M2M markets, filling a few gaps that Wi-Fi was not able to meet. To this day, there is an attempt at expanding Bluetooth use cases that remains highly disputed as many other dedicated industry efforts have pursued the same goals in different domains. More importantly, the limitations in point to multipoint connectivity and bandwidth limitations for audio were not addressed in Bluetooth 5, leading many companies to look for alternatives—even if based on proprietary approaches. Effectively, many companies tried to grab the opportunity to create solutions that performed better in many areas of applications, but apart from introducing new optional codecs, nothing substantial happened to improve the use of Bluetooth for audio transmission.

When in July 2010, CSR acquired Audio Processing Technology (APT) it was because of the fact that aptX offered a scalable and higher-quality audio codec that proved to be particularly suited to Bluetooth wireless audio due to its stronger error resilience, lower latency, and low power consumption, suitable for portable, and battery-driven applications.

LE Audio and LC3

In March 2022, the Bluetooth SIG declared that the LE Audio specification project was "nearing completion," and naturally the consumer electronics and mobile industries were showing increasing signs of impatience. And the root to the problem—apart from the apparent slow progress in the specifications and availability of the technology itself—was mainly to do with a growing anxiety about what was being promised not being in line with market requirements and ambitions.

The situation was not new. In 2014, when the Bluetooth 4.0 low energy (LE) specification and Bluetooth Smart technology were published, the focus was mainly on power-friendly applications with wearables and the Internet of Things (IoT), enabling for better control and faster file transfers in packets, rather than live streaming. And the emerging audio use cases clearly required something else.

And yet, in the next seven years, the intense battlefield for conquering home audio and home theater applications over Wi-Fi was almost forgotten by many companies, given the sheer product volumes happening in the mobile and personal audio spaces, which boosted Bluetooth to become the most widely adopted technology for wireless audio. Bluetooth 5 in 2016 enabled a more robust implementation for wireless audio applications. It was far from solving all the audio industry's concerns, but it was much more robust than what existed and it actually paved the way to huge growth in sales and, more importantly, consumer confidence.

Bluetooth LE Audio, the next generation of Bluetooth Audio was announced in January 2020 at CES and promised to be released with Bluetooth 5.1. Then the global pandemic hit, and much work continued out of sight, but not much progress resulted in the more than two years since that historic announcement. And we just need to compare the scope of what Qualcomm announced during the same period, including the announcement of aptX Adaptive and a lossless CD-quality streaming tier over Bluetooth, to understand that the Bluetooth SIG was not very ambitious and should have moved faster.

What is more noticeable is that the focus for the Bluetooth specification updates was completely centered around an expansion into wearable devices, asset tracking, and network lighting control, in consumer, commercial, and industrial use cases—as determined by the companies that make up the dominant Bluetooth SIG membership. On the audio front, although audio streaming was and is expected to remain the largest Bluetooth solution area for the foreseeable future, the target use cases were centered around hearing aids and true wireless stereo earbuds, converging around the concept of OTC hearing assistance device and consumer hearables.

There's no disputing the massive numbers associated with those segments. By 2026, "annual Bluetooth earbud shipments will climb to 619 million, making up 66 percent of all wireless headsets," predicted the organization.

Finally, in July 2022, the Bluetooth SIG announced the completion of the full set of specifications that defined LE Audio. The release of the complete set of specifications finally enables manufacturers to design and release products supporting the full scope of improvements promised for the next-generation Bluetooth audio, including the new Auracast broadcast audio features.

Bluetooth LE Audio was designed to improve wireless audio performance essentially via the use of a more efficient audio codec (LC3), an alternative to using the Bluetooth Classic radio—also referred to as Bluetooth Basic Rate/Enhanced Data Rate (BR/EDR), a low power radio that streams data over 79 channels in the 2.4GHz unlicensed ISM bands. Supporting point-to-point device communication, Bluetooth Classic will remain in use for the foreseeable future, as it became the standard radio protocol behind portable wireless speakers, headphones, and in-car entertainment systems.

As the names suggest, Classic Audio operates on the Bluetooth Classic radio while LE Audio operates on the Bluetooth Low Energy radio. The Bluetooth Core Specification was enhanced to enable delivery of audio over Bluetooth LE, including the new LE Isochronous Channels feature, enabling the creation of new products and use cases. LE Audio supports multiple streams that are highly synchronized, allowing for source devices to communicate directly to multiple sink devices. This fundamentally solves the separate left/right channel that always complicated true wireless applications, but will also enable streaming to left and right speakers and deliver stereo sound using a standardized approach.

The new Low Complexity Communications Codec (LC3) was developed by the Fraunhofer Institute, also the source for the Advanced Audio Coding (AAC) technology and more recently the more advanced open standard LC3plus codec that combines support for high-resolution audio with low delay. The LC3 audio codec was chosen for its robust low-power implementation possibilities—extremely suitable for hearing-aid applications, although not yet evolving toward the possibilities enabled by the LC3plus codec or similar technologies already available.

Bluetooth LE Audio development solutions are now starting to become available. Nordic Semiconductor has used the latest trade shows globally to showcase a series of new development solutions targeting Bluetooth LE Audio.

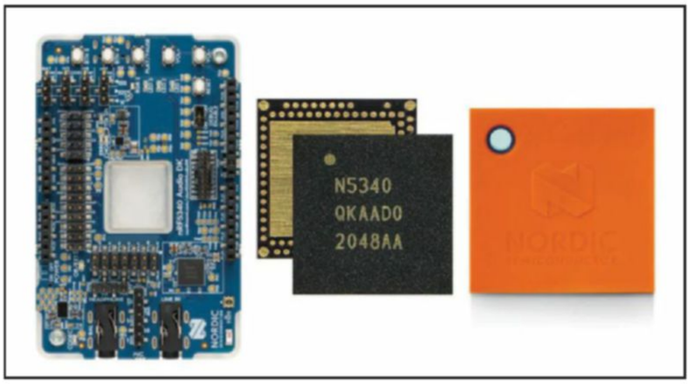

The new Nordic nRF5340 Audio Development Kit, a design platform for rapid development of Bluetooth LE Audio products, based on Nordic's nRF5340 SoC, already supports Auracast for both Broadcast Audio and Audio Sharing.

Packetcraft, the wireless software licensor, was one of the first companies to announce that its Bluetooth LE Audio software achieved qualification from the Bluetooth SIG enabling developers to quickly refresh products with support for audio sharing, and the recently announced Auracast broadcast audio features.

Auracast? No. This is a promotional image from Quiet Events, a Silent Disco headphone rental and event planning company that has been operating since 2012. The entertainment concept is simple. Party-goers get special wireless headphones that can change between 3 Live DJs by flipping a switch, this also changes the LED colors of the headphones. Headphones can be rented for a party at home or a huge gathering of thousands. It works well because it's built over a very robust RF implementation not within the reach for Bluetooth.

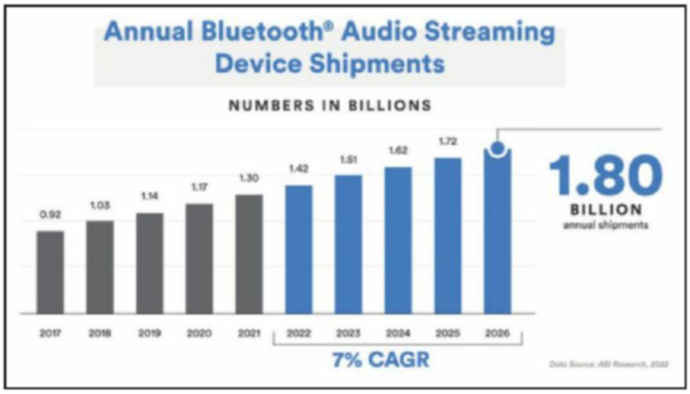

The new LE Audio specifications are supposed to support the development of "next generation Bluetooth wireless streaming applications." The new Low Complexity Communication Codec (LC3) with higher quality, lower power wireless streaming is certainly a great evolution compared with existing Classic Bluetooth audio solutions. But it is only a small part of a solution for a much bigger problem that just keeps getting bigger.

Sennheiser—or better said, Sennheiser Hearing, the consumer division of Sonova—has leveraged the latest Bluetooth LE Audio technology and LC3 codec to update its TV Listening product portfolio. The Sennheiser RS 120-W wireless on-ear headphones became the first in the market to support the updated Bluetooth specifications and offer no less than three different sound modes, working with advanced LE audio transmission to enhance one of the oldest concepts in the brand's portfolio.

Hope remains that the Bluetooth SIG could expand the specification scope in the future to accommodate higher quality audio use cases, but for now, all bets are being concentrated not on quality concerns but on the low-power efficiency for true wireless stereo, and Auracast use cases.

It's more or less obvious that 95% of consumers will embrace the more robust LE Audio implementations and the extended battery life allowed by their new-generation wireless designs using LE Audio and LC3, without even questioning the quality (there's still the question of spatial audio applications, but that remains a topic for a separate article).

Yes, LC3 offers much better quality compared to SBC, but the majority of people who truly care about audio quality are already using wireless earbuds and headphones supporting AAC, aptX Lossless, aptX HD, LDAC, or LDHC. And yet, those remain less than 5%. In general, the other 95% of consumers will not notice any difference, and they are not going to check the bitrate to compare.

And most importantly, no company wants to explain the benefits of LC3 to consumers by stating that the SBC codec and what they had so far was very bad, and that the future with LC3 is basically about "less water pressure in the shower." But that's what LC3 means for any user that cares about audio quality.

The main differentiation for new Bluetooth LE Audio products will come from things that weren't previously possible, and that starts with audio sharing, which is something new, and very likely will also imply that Bluetooth products (True Wireless earbuds in particular) behave much more reliably when it comes to pairing to sources, and toggle between sources.

Implementation will vary from manufacturer to manufacturer initially, since the concept of "switching" between Bluetooth sources is not yet clear. Being able to "find" another Bluetooth audio streaming source that another user wants to share with us—when most likely we're busy listening to our own stream—is not obvious, but is a welcome feature.

The other differentiation comes from Auracast broadcast audio in public spaces. It's up to the manufacturers to find the right interface for users to be able to "tune in" to different Bluetooth streams, and that starts unavoidably with an app interface, which in turn runs on a smartphone that in turn supports Wi-Fi, which is a much more robust and efficient way to distribute broadcast audio channels. To switch on screen-less Bluetooth devices, passcodes, NFC, QR codes, channel scanners, all those things (might) work in some contexts but not always. Earbuds alone with touch sensors, voice commands, and even physical buttons will not cut it. Of course, there's always the possibility to combine some sort of pairing using some of the tools in the Bluetooth set—but the usability and user interaction remains unclear for now.

Bluetooth Auracast

The "new" Bluetooth Auracast broadcast audio capability announced separately and prior to the publication of the LE Audio specifications seemed to be a pure marketing reboot of existing promises for a specification that was delayed for two and a half years.

Previously called Broadcast Audio and Audio Sharing, the new Auracast technology enables an audio transmitter (e.g., a smartphone, laptop, television, or public address system) to broadcast audio to an unlimited number of nearby Bluetooth audio receivers, including speakers, earbuds, or hearing devices.

Basically, this is the same functionality that was available for many years using the SKAA Wireless Audio proprietary technology developed by Canadian firm Eleven Engineering and today available in numerous products. Those systems can receive audio via Wi-Fi or Bluetooth and then transmit using SKAA to multiple receivers. The technology is now seeing a resurgence in DJ and party-speaker applications where dedicated receivers are not an issue. SKAA was always very poorly marketed, otherwise it could have completely dominated the space, at least for professional audio applications (see the side text).

Other earlier attempts to implement the "party" mode functionality over Bluetooth have been demonstrated by companies such as Tempow, meanwhile acquired by Google. Their first product—the Tempow Audio Profile (TAP)—was an updated Bluetooth protocol allowing any Bluetooth chip to stream audio to multiple Bluetooth audio outputs simultaneously. And of note, it was a 100% software solution and worked with any brand of Bluetooth speaker on any chip.

Before being acquired by Google, Tempow focused on pushing the limits of Bluetooth technology and was ahead of the curve of the Bluetooth SIG efforts and meeting real market requirements. Even when LE Audio was announced, in 2020, Tempow released a new Bluetooth Software Stack for LE Audio development and launched a partnership program for chipset manufacturers to use Tempow's implementation since it offered two things that manufacturers where looking to have. The perspective that Google simply "offered" the solution to be part of the Bluetooth specification after acquiring Tempow—hence explaining the requirement for a separate branding—is highly likely.

One way or another, the specifications that define Auracast broadcast audio are now a part of the Bluetooth LE Audio specification suite. The Auracast implementation has two different levels of sophistication. The original Audio Sharing feature depends mostly on the smartphone, earbuds/headphones and speaker manufacturers to enable it in existing chipsets—it's mostly software, as Tempow demonstrated. The Broadcast Audio implementation will require more hardware solutions to meet all different application scenarios—and there are multiple.

Auracast is being promoted for the ability to "Share Our Audio," allowing users to invite others to share their audio experience—on a personal level, and "Unmute Our World," for pure broadcast audio applications. Examples long envisioned and already available over Wi-Fi will now be implemented over Bluetooth (e.g., the ability to enjoy television sound in public spaces). Silent televisions in public venues (e.g., airports, gymnasiums, restaurants, and waiting rooms) will be able to broadcast audio that any visitor with Auracast-enabled Bluetooth earbuds or hearing aids will be able to hear.

Another logical avenue of the technology is the "Hear Our Best" application, where Auracast broadcast audio will support hearing assist when visiting a public venue such as a transit center, cinema, conference center, or house of worship. Visitors to those spaces will be able to receive audio broadcasts from the public address system directly into their Auracast-enabled Bluetooth earbuds or hearing device.

In a clear sign that Bluetooth is not meeting the audio industry requirements, every day manufacturers introduce new solutions that work around quality and latency limitations. An example is this dedicated Bluetooth audio transmitter recently launched by Creative Technology with support for Qualcomm's latest aptX Adaptive high-resolution audio technology and supporting a dedicated low latency mode on 2.4GHz for gaming consoles and applications.

There will be many challenges to this strategy and application scenarios. Experience in this domain tells us that the challenge is always in the implementation of the source signals.

For reference, we just have to look at the actual experiences from Listen Technologies or Sennheiser with its MobileConnect solution. Both companies have extensive experience and they had the technical solution fully implemented. Smartphones all have Wi-Fi, an app allows for low latency and ideal user interface, and radio coverage over Wi-Fi is superior. Neither of the two companies has achieved significant success with their solutions (if we ask Listen, it probably considers its business very successful, maybe just not for the metrics used by a Samsung or Google). Before selling that part of that business to EPOS and then Sonova, Sennheiser was well-aware that true wireless earbuds would be used in both business and education applications, as well as assistive listening for users with hearing disabilities.

Already in 2022, Qualcomm announced new options for brands and developers to offer the latest technologies in wireless audio. Qualcomm's latest Snapdragon Sound S5 and S3 Sound Platforms offer aptX Lossless CD-quality streaming over Bluetooth, stereo recording, ultra-low latency sound with in-game chat, and adaptive active noise cancellation. Both platforms are optimized for dual-mode and support LE Audio.

Many companies have never accepted the bandwidth limitations of Bluetooth and clearly believe there is a market opportunity there. HED Technologies, a company with a focus on audio and based in Los Angeles, Geneva, Taipei and Dublin, launched Unity Full-Fidelity wireless headphones that are able to receive lossless 24-bit/192kHz (FLAC) audio over Wi-Fi. Retailing for $1,799.

Looking Ahead

Market requirements for more bandwidth, lower latency, and a more robust and seamless user experience with multidevice connectivity remain critical points that the Bluetooth specifications need to address—required by the whole audio and gaming industries, which together account for the vast majority of Bluetooth devices shipping today.

Not conforming to the current state in wireless audio technology, over Wi-Fi and much less over Bluetooth, large technology companies such as Apple or Qualcomm are clearly looking at setting their own pace and strategies for wireless technology. Apple is already exploring alternative technologies, including Ultra-WideBand (UWB)—www.firaconsortium.org—to do most of the things that the Bluetooth is aiming to do in areas such as access control, location-based services, and device-to-device (Peer-to-Peer) connectivity and tracking. And Qualcomm is pushing its Snapdragon platforms to encompass a complete strategy that includes all wireless protocols and promotes aptX Adaptive and optional profiles as a solution to solve the problems that the more discerning consumers recognize today with Bluetooth—even if still conditioned to working with the all the identified Bluetooth technology limitations.

Following the original Snapdragon Sound announcement in 2021, Qualcomm has expanded its strategy and announced the availability of aptX Lossless as another major tool in the platform arsenal for differentiation in new true wireless stereo earbuds, headphones and speakers. Qualcomm's two new ultra-low power wireless audio platforms, the new Qualcomm S5 Sound Platform (QCC517x) and Qualcomm S3 Sound Platform (QCC307x) are optimized for Bluetooth Classic and the latest LE Audio technology, already including support for audio sharing and broadcasting. And with the expanded aptX suite and aptX Adaptive, manufacturers can add aptX Lossless for CD-quality audio, and still be able to offer 24-bit/96kHz high-resolution Bluetooth audio quality with aptX HD.

The Qualcomm S5/S3 Sound Platform SoCs also support stereo recording and gaming mode capabilities that are based on the Bluetooth LE Audio specification. For gaming, they enable 25% lower latency audio (compared to previous generations) at only 68ms, and with voice back channel.

In different ways, both companies are looking for a solution that improves on the fundamental requirement of "more bandwidth." And while both continue to promote industry standards and are part of consortiums and common industry efforts, both recognize that "being ahead" of standards requires the need to control both the source and the receiver (sink) side of the equation—and both are fine with that, and even see it as a strategic benefit to build a long-lasting market position. Specifically for Apple, which offers both source (iPhones, Macs) and receiver devices (AirPods, Beats) all options are on the table, as long as its products meet user expectations. And we know how large those expectations are for the future generations of AirPods, just as an example. Wireless audio streaming remains at the core of those expectations.

In the next article, we will explore the latest approaches for wireless audio transmission, including the evolution of Wi-Fi, dedicated solutions including WiSA, DTS Play-Fi, SKAA, KleerNet (yes, they are still around), and the promising prospects on the UWB front.

The Bluetooth Multipoint Challenge

Multipoint or multidevice connectivity over Bluetooth remained an afterthought on Bluetooth for far too long. Bluetooth Multipoint allows users to pair their headphones or earbuds with more than one source, and instantly switching between them when needed. This feature suddenly started to be advertised by virtually all companies when announcing new Bluetooth 5.2 and Bluetooth 5.3 products. The Bluetooth SIG calls this "multi-source Bluetooth," because basically it involves having the same sink device being able to receive sound from two "source" devices. Never simultaneously, since current Bluetooth technology doesn't allow multistream.

Around for some time but never properly implemented, Bluetooth multipoint or multi-source was once advertised as a feature in office headsets that needed to switch calls quickly between a mobile phone and a PC. But effectively it never worked exactly as advertised. And the main reason has to do with the problem of switching between two different source devices that might be using different Bluetooth profiles and codecs. It is possible to pair to two different source devices and switch between them, as long as devices are streaming an AAC codec, for example. This worked well for Sony, Microsoft, and naturally Apple, which prioritize using AAC.

But even when it works once, it rarely works again as intended, because the switching is always cumbersome—usually it takes so much time that the user decides to "do something" about it, and stops the process. As an advanced feature in office headsets, this was explored mainly as way for users to be able to manage and switch between calls, and not high-quality music streams—

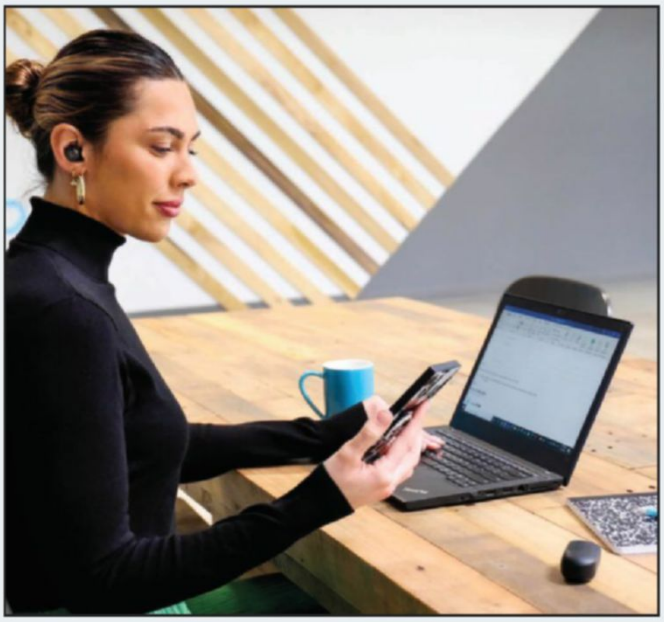

Bluetooth multipoint allows users to connect their earbuds, headset, or headphones (sink devices) to maintain simultaneous connections to at least two source devices (e.g., a laptop and a smartphone).

which require more time to be detected and for the streaming to be switched. Again, reliability was always an issue, and this was before true wireless—where the two independent left/right earbuds almost always use proprietary schemes to enable the correct channel separation and synchronization.

For all those reasons, very few true wireless earbuds support multipoint and even Apple has been hesitant to promote the concept, even if it does support multi-pairing and multipoint source switching. In fact, that feature is supported at the operating system level in Apple devices. But Apple usually implements auto-switching when it detects the user has stopped listening to Apple Music on the iPhone and started watching a movie on a MacBook, as an example. That way, it avoids the possibility of user error and at least it is able to prioritize calls. In its documentation for developers, both for Mac OS and the External Accessory MFi program, Apple clearly advises: "Sound should not automatically switch from one device to another if you're in a conversation, like a phone call, a FaceTime call, or a video conference."

Trying to do the same with a Windows laptop and an Android phone is a nightmare for audio manufacturers. The Bluetooth 5 specification theoretically allows multi-source pairing, but manufacturers are afraid to tell consumers that it "only works with Samsung" or that they need to only use the latest-generation devices with the latest OS, Bluetooth profile, and codec... Also, sometimes it does work the first time, and then one day it doesn't—the pure definition of unreliability.

Bluetooth LE Audio, when available, will potentially help to make this work much better since it supports multi-stream and isochronous audio streams with multiple devices. But this is the opposite example of Audio Sharing and Broadcast Audio (Auracast) mode—one source, many sink devices. Multipoint is about switching between multiples sources to the same sink device. And the missing key component is the definition of "switching," which is associated with human interface.

Lundahl Transformers

How High Performance Transformers Are Made

By
Scott Dorsey

Is there anything better than traveling through Scandinavia during summer weather and visiting leading manufacturers for the audio industry? Not for Scott Dorsey, who recently fulfilled his dream of visiting the Lundahl Transformers factory, located in the small town of Norrtälje, just outside Stockholm, Sweden.

Lundahl Transformers was founded by Lars Lundahl in 1958 to make small transformers for the growing electronics industry. Making commodity transformers for television sets and the defense industry helped the company grow.

In 1970, the company began making audio transformers for the Swedish Broadcasting Corp., which brought them to the attention of a number of manufacturers in the British market and soon they were in the audio market full-time. In 1985, with the resurgence of interest in tube audio, they began making higher power transformers for tube audio amplifiers.

Photo 1: Winding six coils at a time on a stick

Lars' son Per joined the company in 1990 and today runs the operation, and continues to make transformers of the highest grade. He says he is "lucky enough to find a niche where we can charge as much as it costs to make a quality product," and that is a rare thing in today's world.

What makes them most interesting from our perspective is that they manufacture C-core transformers, which is a European configuration that is seldom used by American manufacturers and is constructed in a very different way than the common EI-core transformers. There are very few chances to actually see C-core transformers being made.

Lundahl basically has a single production system optimized for this design, and so they create new designs based on required electrical specifications, which can be easily built on their existing line.

Why C-Core?

The C-core method means:

1. Stick winding is possible for consistency and higher production rates (the next section will explain this).
2. Insulation between windings is very effective and insulation between layers of the same winding is possible.

3. Use of a core that can be inserted into the windings—no need to wind around the core itself.

Making the Windings

A custom-made winding machine wraps a fixed number of turns over a wrapped Mylar sheet and then additional inter-layer sheets can be added in the process. This gives you a stick with a number of individual winding combinations on it, and the stick is cut apart to produce the winding assemblies that will later be mated with the cores.

You can see in **Photo 1** how six different coils are being wrapped at the same time around a winding bar with a Mylar sheet on it. Additional sheets will be added after the windings are finished and the whole stick of six coils removed from the machine.

This "stick winding" technique creates very clean and even winds since there is no axial force from the edges of the bobbin to make for a poor winding lay on the edges. It also allows a layer of insulation material to be put between each and every layer of copper wire, which reduces interwinding capacitance and increases breakdown voltage, making for more consistency between units, better high frequency performance, and more reliability.

The stick winding method was developed when wire varnish was not as good as it is today, when high voltage breakdown at high frequencies was difficult to obtain. This is much less critical today for power transformers but still a great advantage for tube amp transformers that might experience very high voltage spikes during faults.

The individual sticks are not cemented or glued in any way; they are held together only by the tension of the windings and the Mylar against one another. This gives good long-term stability since there is no cement to degrade.

Most of the Lundahl designs are dual-coil transformers, as you can see in **Photo 2**. This gives better symmetry and reduces magnetic coupling to the outside world.

Often the winding assemblies are made in large quantities and then stored for future use. The same winding assemblies might be used in several different transformers; for example, a single-ended output transformer and a different push-pull model might use the same winding assemblies with different core gapping. The winding assemblies are usually referred to by the model of transformer for which they are intended (**Photo 3**).

Making the Core

A pre-slit strip of grain oriented silicon iron steel is wrapped by an automated machine around a rectangular form to create a continuous rectangular core. Since the wrapping process introduces stress into the material, the cores are annealed after wrapping.

In the annealing process the cores are kept in desired form in a corset arrangement. After cooling down, the material is stress free and has accepted the new, more rectangular form (**Photo 4** and **Photo 5**). The cores are then cut in half on a saw

Photo 2: Double coils are assembled onto mounting plates ready for the cores.

Photo 3: Stock coils are kept on shelves waiting assembly.

Photo 4: Cores after wrapping but before annealing

Photo 5: Squared-off cores after annealing

into two C-shaped sections and the edges carefully polished so that they fit together precisely. You cannot completely eliminate the gap caused by the surfaces never touching perfectly, but you can get it very small with careful lapping (**Photo 6**). The two sections are kept together so every core shipped is made from one individual wrapped core.

Assembly

The two halves of the core are fitted into the winding assemblies. If there is a gap required, plastic spacers of a very precise thickness are inserted between the two halves (**Photo 7**). A band is put around the assembly in order to keep the core perfectly stable for the life of the transformer, and then the whole assembly is impregnated.

After the assembly is constructed, the transformers are put into a tank of varnish and a moderate vacuum applied to remove all the air and make sure the varnish gets into every corner (**Photo 8**). This gives better electrical stability but since the varnish does not get absorbed into the Mylar it does not give better mechanical stability the way it does with paper insulation.

Lundahl does still use paper insulation on a couple of models of power transformers for that reason; the mechanical stability reduces acoustic hum from the transformer windings and high frequency, high voltage breakdown is not as much of an issue as it would be for output transformers.

Lundahl stamps and forms its own mu-metal shielding boxes, into which the assemblies are inserted (**Photo 9**). They don't stretch or expand the metal and then they re-anneal the shields the same way they do the cores.

Testing

Once assembled, every device is hipot tested (a high potential, high voltage test, also known as a Dielectric Withstand Test). The devices are then tested for winding impedance with other windings left open (**Photo 10**). Distortion is tested and a final check is made on the turns ratio just to make sure

Photo 6: Core after annealing and cutting

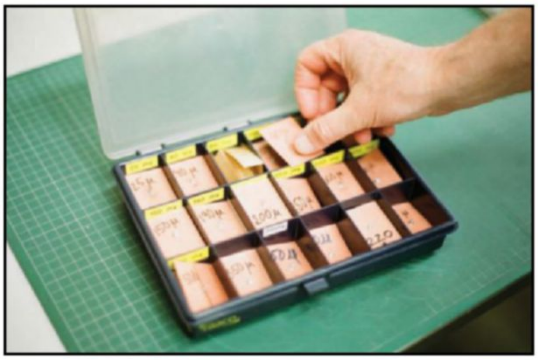

Photo 7: Spacers for core gaps

Photo 8: Impregnating finished transformers

Photo 9: Mu metal box cut and formed on site

Photo 10: Extra-large winding machines for large coils

Photo 11: Production testing

something wasn't mislabeled along the line. The testing process is automated and test records are kept on file of every device (**Photo 11**). **Photo 12** shows the completed product.

Photo 12: Tube audio manufacturers have wanted Lundahl to continually push performance levels, and to further study the intricacies of transformer design to meet their very special demands. The first manufacturer to use Lundahl tube amplifier transformers was Shindo Labs, a company in Japan from where several prestigious transformer brands originated.

Other Types

In addition to C-core transformers, which are primarily used for pro audio line output transformers and for tube amplifier transformers, Lundahl make transformers with conventional mu metal lamination cores and with strip wound amorphous cobalt cores. These transformer types are primarily for small signal input applications. Lundahl has a proprietary process of forming the amorphous cores in place inside the windings. I think this is unique. It is definitely different than the way most companies make amorphous core transformers by wrapping windings around a toroidal core, and it allows physically smaller transformers to be made which would be difficult to wrap with standard methods.

Historically, Lundahl has made all manner of other small signal transformers but with the changes in the market it no longer does.

About the Author

Scott Dorsey has a degree in electrical engineering, during the pursuit of which he worked in the broadcast and recording industries. After several years working at a major studio, he took a job with a defense contractor. This left him time to do live concert recording for acoustical music and to design and build audio devices for personal use and on contract to several audio manufacturers and importers. Scott is a regular contributor to several audio magazines. He has been publishing equipment reviews and DIY projects since the mid-1980s. He is probably best known in the general audio community for his retrofit electronics designs in inexpensive Oktava, AKG, and Feilo microphones.

When Lundahl started out in business, every table radio had an audio output transformer, and television sets were filled with wideband magnetics not just for audio output but for the horizontal and vertical sweep circuits. There was a huge demand and large volume production of small audio transformers, but there was also a lot of competition. That market is gone, and what is left are two rather different groups of people in the pro audio and audiophile markets, both looking for much smaller quantities of transformers and wanting transformers of much higher quality. Lundahl, like Jensen, Sowter, and a small handful of other companies, focuses on these markets today. It's a good place to be.

Timeline of Lundahl Transformers

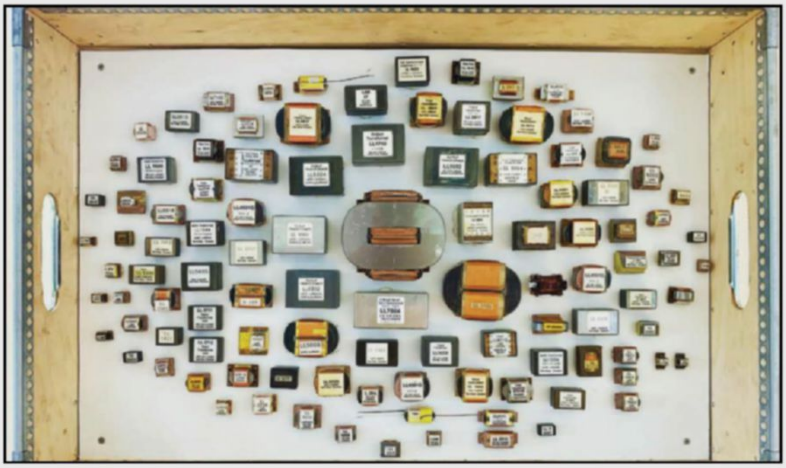

1958
Lundahl Transformers is founded. After his exam from the Royal Institute of Technology in Stockholm, Lars Lundahl started working for AGA, a big gas corporation who at that time had ambitions to start manufacturing electronic products, both for the professional and consumer markets. During his years at AGA, Lars worked with magnetic amplifiers for control applications, and was introduced to the newly invented transistor. He also discovered that it was hard or almost impossible to find good transformers, so he decided to start his own company to manufacture transformers.

The first winding machine was installed in the basement of the family house, soon to be followed by more production machines as well as a second hand lathe and ditto milling machine.

1960
Following its ambition to become a consumer product manufacturer, AGA had started manufacturing radio and TV units. One (at least) of the TV models was equipped with a Lundahl audio isolation transformer.

1963
Due to the early success of the company, new premises were needed. After visiting most towns in the Stockholm region, an empty building-site was found on the outskirts of Norrtälje, a small town located 70km northeast of Stockholm.

1965
Lars Lundahl discovered the advantages of the stick winding technique. The normal way to wind a transformer was (and still is) to use a bobbin. But in a bobbin there are axial forces from the sides of the bobbin, causing wires to cross in an uncontrolled way). With the stick winding technique there are no axial forces, and in addition you can put insulation material between each layer of copper wire. In spite of the added internal insulation (which reduced internal capacitance) the density of a stick wound winding is higher than the density of a bobbin-wound winding.

1970
Lundahl develops a set of transformers (SR501, SR502, SR503 etc.) for the Swedish Broadcasting Corp. (SRT). A British engineer visits SRT and is introduced to Lundahl transformers. He brings some transformer to show to other engineers in the UK and the brand starts to be used by a number of major UK sound equipment manufacturers such as Trilogy, Soundcraft, Calrec, Focusrite, and others.

1975
Already manufacturing its own C-cores using grain oriented silicon-iron sheet metal, Lundahl also starts to manufacture its own mu metal housing to isolate the transformer magnetically from external noise.

Per and Lars Lundahl

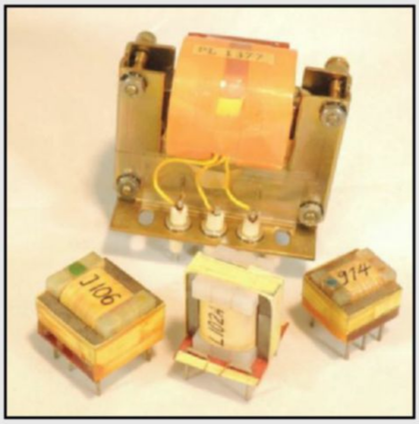

The first transformers manufactured by Lundahl

1976
With a workshop in house, and with full control over winding size, C-cores and mu metal housings, it was possible for Lundahl Transformer to tailor the shape of transformers to the available space. In some cases a low profile was demanded, in other cases only a small board space, but plenty of height, was available. Lundahl is able to design and manufacture transformers for any requirements.

1980
Lundahl starts using amorphous metal and develops a technique to manufacture strip wound amorphous core transformers, highly appreciated in audiophile circles. In such transformers the amorphous strip is wound inside existing coils, as opposite to toroid transformers where the copper winding is wound inside an uncut core.


In 2020, Rikard Wallin was appointed CEO of Lundahl Transformers, leaving Per Lundahl free to focus mostly on product and production development.

1985
Lundahl starts making transformers for tube amplifiers. It was a natural move, because Lars Lundahl had a background as a tube amplifier hobbyist.

1990
The Lundahl factory building is expanded. Per Lundahl joins the family business and the company starts exhibiting at Audio Engineering Society conventions.

2000
With Per Lundahl as managing director, the company's international recognition expands and the company decides to expand the factory once more, to meet increasing demand.

2007
To serve the audiophile market always in search for the best possible sound, Lundahl develops a group of moving-coil step-up transformers using oxygen free, post annealed ultra-pure copper wire from Cardas Audio in the United States. In spite of the high price, the new product range was so popular that Lundahl expanded the use of Cardas copper wire in other, newly developed transformer types.

2009
For the extreme high-end purist Lundahl starts offering transformers with silver wire windings, as well as a limited range of interstage and output transformers with silver wire.

In the same year, the company also starts supplying housings for Lundahl Transformers in response to demand from tube amplifier DIY customers who prefer to place the transformers on top of their amplifiers.

2022
Today, Lundahl offers a complete range of products for audio, including tube amplifier output transformers, Moving Coil step-up transformers, line input and line output, interstage, headphone and line output transformers, transformers for ground isolation, for splitting and for balanced-to-unbalanced conversion, mains transformers and chokes.

Updated Smartphone SPL Apps

By
Richard Honeycutt

Recently, a colleague asked Richard Honeycutt about smartphone apps for measuring sound pressure level (SPL). His first reaction was to refer him to the Sound Control articles published in *audioXpress* back in 2015. But then he realized that information could be outdated. Some research verified his suspicions. Turns out more recent studies have been done, some new apps are available, and older apps are no longer offered. This article provides an updated look at what is currently available.

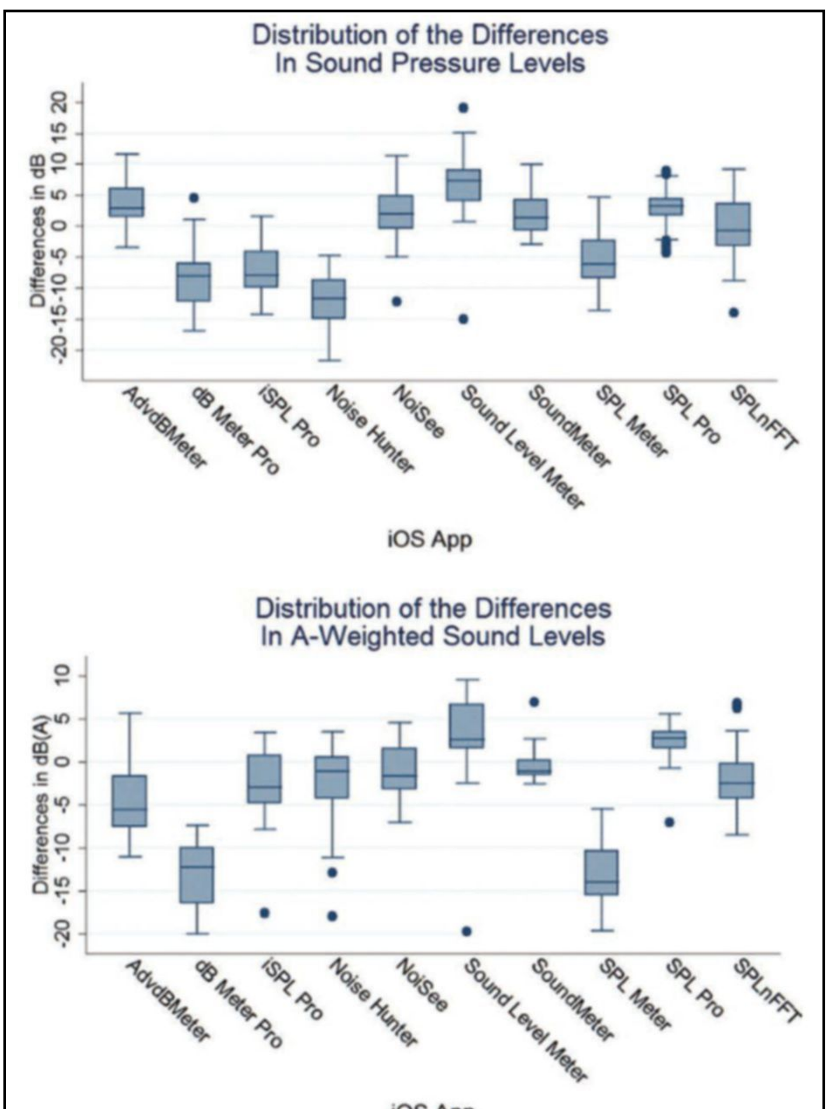

Figure 1: Kardous and Shaw's paper showed wide performance variations among smartphone SPL apps.

In April 2014, Chucri Kardous and Peter Shaw of National Institute for Occupational Safety and Health (NIOSH) published a paper in the *Journal of the Acoustical Society of America*'s "Express Letters," [1], in which they concluded that:

"Certain Measurement apps for Apple smartphones and tablets may be considered accurate and reliable to be used to assess occupational noise exposures."

"Android and Windows developers do not offer apps that meet the functionality needed for occupational noise assessments."

"Field measurement results may vary greatly due to the effect of temperature, humidity, long-term use, object interference, and overall stability of the microphone and electronics in these devices."

The same app used on different generations of iOS devices yielded somewhat different results.

The accuracy of these measurements was determined in comparison with a newly calibrated Larsen-Davis model 831, Type1 sound level meter (SLM) using a model 2559 random-incidence microphone. Deviations varied from 0.07dB (for unweighted measurements) to well over 10dB for the worst case. Due to the superior performance and repeatability of iPhone apps, the study primarily evaluated these. Only four Android apps "partially met" their selection criteria. **Figure 1** shows a graph of the results of their tests. While a few of the apps' measurements came close to those of the Larsen-Davis SLM, interestingly, the apps that performed best on unweighted measurements did not always show up best on A-weighted measurements. Not all the apps they tested are still available.

Cell phone mics have evolved over the years: From the Motorola DynaTAC concept phone built in 1973 (**Photo 1**), using a carbon mic, through early designs having electret capacitor microphones (ECMs), to the introduction of microelectronic microelectromechanical system (MEMS) mics used in the iPhone 4 and later, portable phone mics have increased in quality and decreased in price. Even so, Kardous and Shaw were aware that microphone performance could be a limiting factor in the accuracy of smartphone-based SPL meters.

Therefore, in 2016 they published a follow-up paper reporting the results of comparing different SPL apps, using external calibrated microphones, thus removing the microphone as a variable [2].

Limiting their study to the four iOS apps found in the previous study to exhibit mean differences of ±2 dB of a reference sound level measurement system, they obtained the results shown in **Figure 2**. Although their results showed measurements within + 1dB of the reference, they cannot be generalized to say that jut any SPL app on just any iPhone will produce accurate results if used with a good external mic. But the study does verify the value of using external calibrated microphones in improving the accuracy of smartphone SPL measurements. The external mics they used were an i436 mic by MicW (**Photo 2**) and the iMM6 by Dayton Audio (**Photo 3**)—both are ECMs that can be calibrated. Their specs are shown in **Table 1**.

It should be noted that the prices given were determined in 2016, and are different now.

Three other measurement mics designed for use with iOS devices should be mentioned. Studio Six Digital introduced the iTestMic, which has been replaced by the iTestMic2 (**Photo 4**) and AudioControl manufactures an iOS-compatible measurement mic: the SA-4140i-SPL, which looks very similar to the Studio Six iTestMic2 and has two gain ranges: 57dBA to 137dBA and 60dBA to 140dBA.

Choices

There are many reasons why you might want to measure SPL. You may just be curious how loud a sound is. A good friend of mine often becomes vary annoyed about the noise level produced by high-air-velocity electric hand dryers. One of my own pet peeves is restaurants in which one has to shout to communicate with one's table-mates. In both cases, curiosity could inspire us to measure the SPL. A more important reason for measuring SPL is hearing protection. In the former case, a "ballpark measurement using either flat or "A" weighting would be sufficient. A "slow" response makes the SPL easier to measure in the normal

Photo 1: Motorola's Martin Cooper is shown holding a 1973 prototype of the Dyna TAC cell phone.

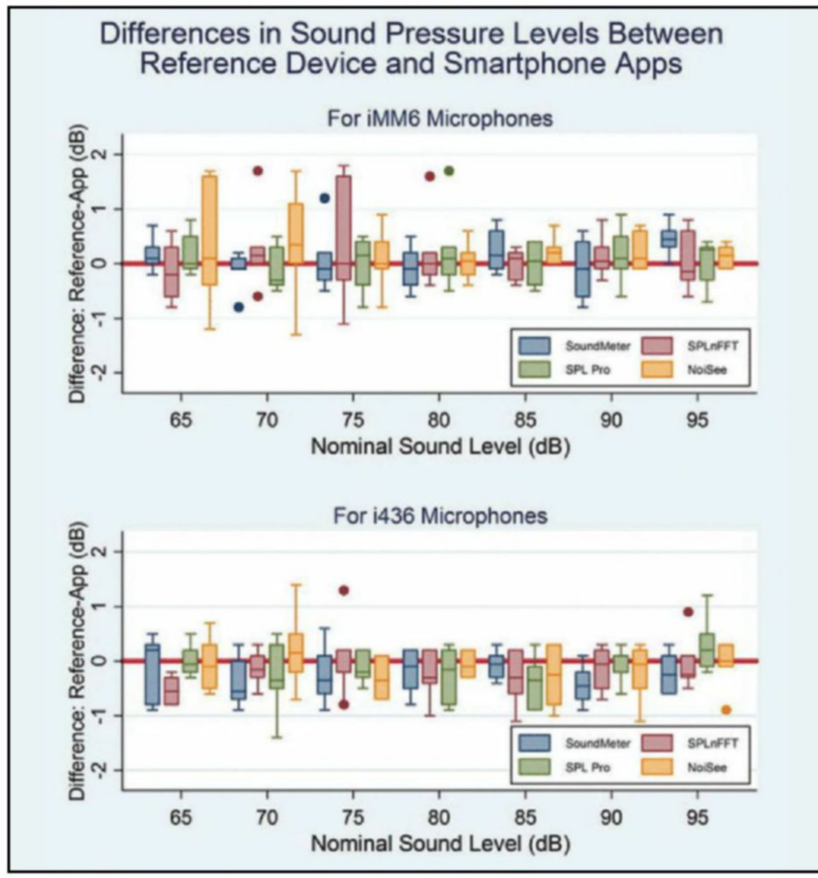

Figure 2: Kardous and Shaw tested the four best-performing apps using identical external mics.

Photo 2: The i436 test mic by MicW can improve measurement accuracy of iPhone SPL measurements. Unfortunately since Apple removed the jack from the iPhone this can only be used in older models.

Photo 3: The Dayton Audio 1MM6 is another alternative calibrated external test mic.

Microphone	Cost	Capsule Size	Sensitivity	Frequency Response	S/N Ratio	Maximum SPL
i436	$150	7mm	6.3mV/Pa	20Hz to 20kHz	>62dB	128dB
iMM-6	$15	6mm	10mV/Pa	18Hz to 20kHz	70dB	127dB

Table 1: These are the specifications of external mics used by Kardous and Shaw.

Photo 4: The Studio Six Digital iTestMic2 is terminated with a Lightning Connector for easy connection to an iPhone or iPad.

Photo 5: The Studio Six SPL Meter includes choices of "A" or "C" (flat) weighting and "SLOW" or "FAST" response time.

About the Author

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

situation in which the sound level is continuously varying. Maximum/minimum indications are also handy in these cases (**Photo 5**).

For hearing protection, the duration of noise exposure is critical, leading to the need for a noise dosimeter that automatically calculates the combined effects of SPL and exposure time (**Photo 6**). And sometimes it is helpful to be able to graph the noise level versus time (**Photo 7**).

Some venues and some communities have noise regulations, and sound-system operators need to be able to monitor the SPL continuously to ensure compliance, or law-enforcement personnel need to be able to document illegally high noise levels. "Traffic-light" displays can ease this task (**Photo 8**).

SLMs employed to document compliance or noncompliance with noise ordinances must comply with national and international standards, such as American National Standards Institute (ANSI) S1.4-1983 (R2007), Specifications for Sound Level Meters (ANSI, 1983), and International Electrotechnical Commission (IEC) 61672-1 (IEC, 2013). Both standards specify that SLMs must pass certain acoustical and electrical tests with indicated tolerance limits and measurement uncertainties,

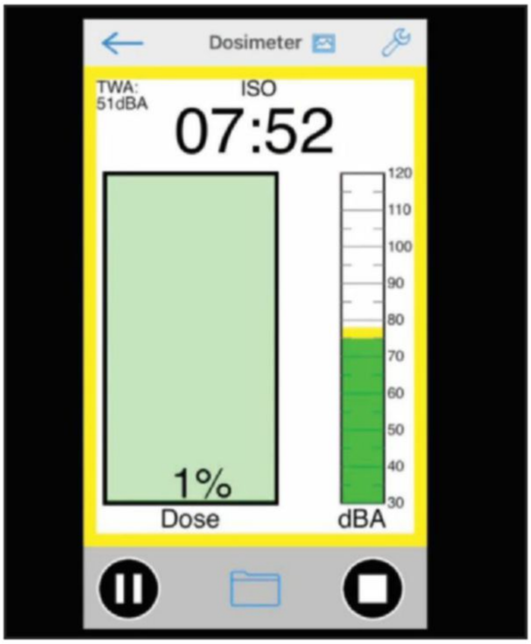

Photo 6: The Studio Six noise dosimeter shows the combined effects of SPL and exposure duration.

specified in decibels, over a wide frequency range (typically from 10Hz to 20kHz).

Such tests must certify level linearity, directionality, time and frequency-weighting responses, tone bursts, radio frequency interference, and atmospheric and environmental conditions. The standards also specify that these tests shall be made on the complete instrument, including the microphone and preamplifier. Because of the variations among apps (e.g., mics, smartphone models, and others), only one smartphone-based app has met—or is likely to meet the requirements of IEC or ANSI standards.

The NIOSH Sound Level Meter used with a suitable external mic meets the Type 2 requirements of IEC 61672:3. (The older standard IEC 60651 referred to the grade as "Type," whereas the new standard IEC 61672 refers to it as the "Class.") Despite ANSI and IEC type specifications' applying only to complete SLMs, some manufacturers of measurement mics describe their mics as "ANSI Type such-and-such" or "IEC class such-and-such" mics. By this, they usually mean to indicate that a particular mic meets the tolerance specs of the

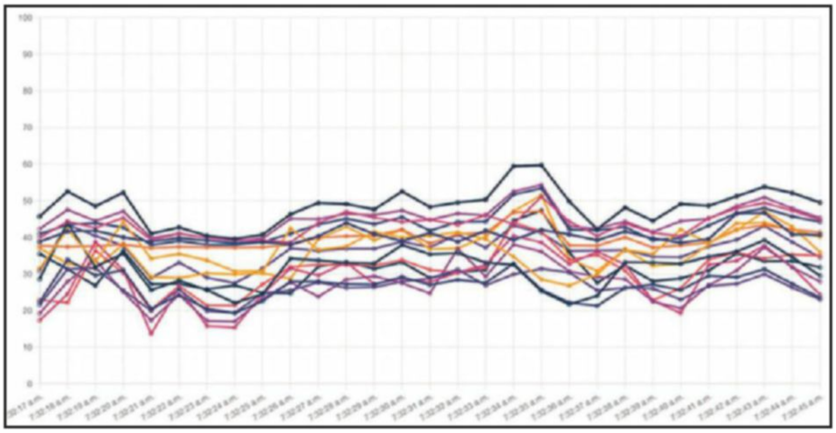

Photo 7: The Studio Six SPL Graph app can log the levels for multiple octave bands.

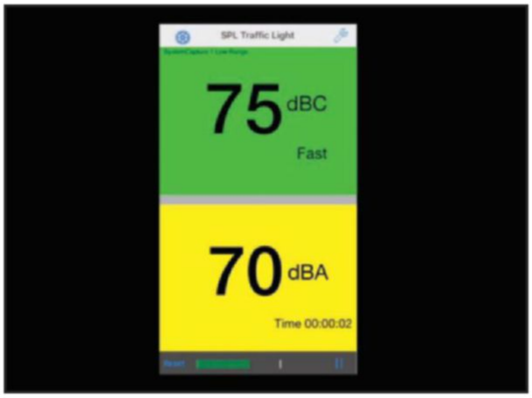

Photo 8: The Studio Six dual traffic light helps keep sound levels within stated limits.

listed ANSI type or IEC class. Examples for the ANSI types are shown in **Table 2**, for the frequency range 20Hz to 20,000Hz.

Class	SLM Accuracy	Calibrator Accuracy
Type 0	+0.4dB	+0.15dB
Type 1	+0.7dB	+0.3dB
Type 2	+1dB	+0.5dB
Type 3	+1.5dB	What calibrator?
Not spec'd	Who knows?	Oh, come now

Table 2: Sound level meters must meet these specs in order to be classified as a specific ANSI type.

Photo 9: The Studio Six Octave and 1/3 Octave Band Logging function displays the frequency content of noise.

Of course, using a mic that meets ANSI Type 1 standards does not mean that your smartphone can perform measurements with ANSI Type 1 accuracy. In community noise and industrial noise investigations, and in mitigation of HVAC noise, the spectrum of the noise provides essential information (**Photo 9**).

Pro User Suggestions

Online conversations with several highly respected audio professionals have revealed several recommended apps for pro audio use. Perhaps the one receiving the most positive comments is the NIOSH Sound Level Meter (**Photo 10**) This free app is tested and validated (accuracy ±2dBA) according to standards in a reverberant chamber at the NIOSH acoustics lab. In addition to the normal SLM function, it also functions as a dosimeter, providing averages such as LAeq and TWA, Max and Peak Levels, Noise Dose, and Projected Dose according to NIOSH and OSHA standards, and all three major weighting networks (A, C, and Z). When used with an appropriate external calibrator, the NIOSH SLM provides calibration capability for either internal or external mics.

The Sound Meter from Melon Soft (**Photo 11**) received positive comments from one of the audio pros. Figure 1 and Figure 2 show the performance of Sound Meter measured by Kardous and Shaw.

Decibel X Pro is another app favored by some audio pros. Since this app was not tested by Kardous and Shaw, I tested the app for accuracy, using as a standard a recently calibrated NTi AL1 with a Sabine SQ-1001 calibrated microphone.

References

[1] C. A. Kardous and P. B. Shaw, "Evaluation of smartphone sound measurement applications," *The Journal of the Acoustical Society of America*, 135, EL186, 2014, https://doi.org/10.1121/1.4865269.

[2] C. A. Kardous and P. B. Shaw: "Evaluation of smartphone sound measurement applications (apps) using external microphones—A follow-up study," *The Journal of the Acoustical Society of America*, 140, EL327, 2016, https://doi.org/10.1121/1.4964639.

Sources

AudioTools - dB, Sound & Audio
Apple App Store
https://apps.apple.com/us/app/audiotools-db-sound-audio/id325307477

SA-4140iSPL High-SPL iOS Test and Measurement Microphone
Audio Control
www.audiocontrol.com/pro-audio/ios-microphones/sa-4140i-SPL

iMM-6 iDevice Calibrated Measurement Microphone
Dayton Audio | www.daytonaudio.com/product/1117/imm-6-idevice-calibrated-measurement-microphone

NIOSH Sound Level Meter
EA Lab |
https://apps.apple.com/us/app/niosh-sound-level-meter/id1096545820

Sound Meter Decibel for PC
Melon Soft | https://desktoptwo.com/app/app.melon.sound_meter

i436 Measurement Microphone
MicW | www.micwaudio.com/product.php?id=3

Decibel X: Pro dBA Sound Meter
SkyPaw Co., Ltd. | https://skypaw.com/decibelx. html

iTestMic2 Test & Measurement Microphone
Studio Six Digital | www.studiosixdigital.com/audio-hardware/itestmic2

Because of the strong dependence of measurement accuracy upon the microphone, I used the same Studio Six iTestMic2 and my iPhone 6 for the Decibel X test. Its measurement of an 114dB 400Hz tone exactly matched the NTi/Sabine reference.

SLM apps in the AudioTools package have been mentioned several times in this column. The basic AudioTools set of apps includes the SPL Meter with an analog display. Five more specialized SLM apps are available as in-app downloads. ax

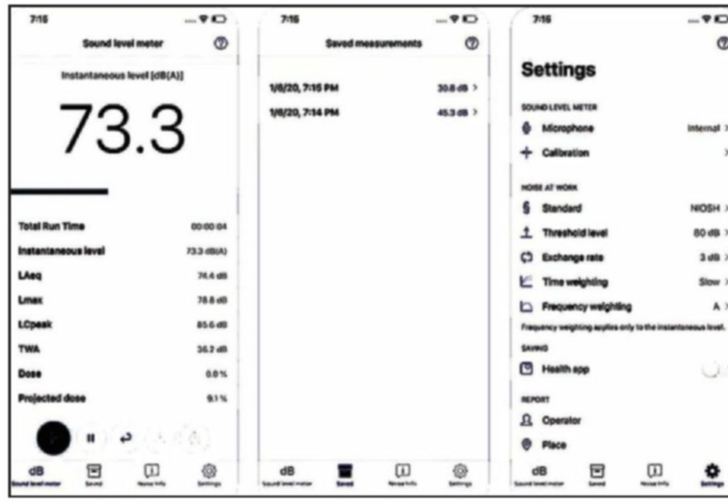

Photo 10 (shown above): The NIOSH Sound Level Meter provides multiple useful functions.

Photo 11: The Sound Meter displays the noise waveform as well as the level.

Development and Design

Solutions for Building Voice User Interfaces at the Network Edge

Voice interfaces require AI and machine learning processors directly on-device. The level of computation for voice processing using edge-AI optimized processors also enables exciting possibilities for highly sophisticated audio processing in applications from low-power earbuds to automotive interfaces.

By
Matthew Mitschang
(DSP Concepts)

As an example of voice applications on the edge, Sensory created new custom trained speech recognition models that understand the unique linguistic patterns associated with children's speech and unlocks a higher level of accuracy and privacy for apps, toys, kid's wearables, and education technologies, all key market segments for voice interface applications.

Edge computing, where the processing of data occurs at the "edge" of a local network rather than through communication with a cloud provider, has prompted a shift away from cloud computing in recent years. Information security and privacy, reduced latency, and an increase in intelligent applications are primary reasons driving the adoption of edge devices for infrastructure and enterprise, as well as in consumer market verticals.

Devices featuring a voice user interface are on the rise as well, according to a report by Juniper Research that outlines expectations that the global market for voice assistants will reach a value of $99 billion USD by the year 2025, representing a CAGR of more than 85% from 2021 to 2026. With applications spanning from smartphones to automotive, smart speakers and wearables, these processors are increasingly powering the more than 350 million wearable and hearable devices shipped in 2022, according to ABI Research.

In general, voice user interfaces (VUI) appeal to those with convenience in mind, as well as those who are reticent to use shared, touch-based controls due to hygiene concerns. However, wider adoption of these products is hampered by concerns about privacy, accuracy, and the ability of devices to recognize various accents or dialects. Other obstacles that can be encountered with voice-enabled products are sub-optimal acoustic environments and extra latency incurred by the interpretation and processing of commands. Consumers are more likely to adopt smart devices if its functionality has an improved ease of use and can provide a seamless experience, rather than a reaction to user input that has a low response accuracy rate or a cumbersome feel.

To capture the most of their market segments, voice-enabled products must have reliable performance and be adaptable to different languages or geographical regions. The development and production costs, along with the often-difficult integration of voice technology, can be a problem for product makers. They must grapple with the development time needed to create and deploy competitive voice recognition functions while also adapting their designs for different languages and regional accents.

Use Cases

For edge-based voice recognition, there is no reliance on cloud services for speech-to-text and command recognition. While a cloud-based alternative does lend itself to certain use cases (e.g., performing a web search via voice), processing commands at the network edge provides a much lower latency, and its offline nature is private by design. Where the cloud may offer access to more information and more potential processing power, edge-based voice user interfaces are embedded directly on the device and are able to respond without extra delay, resulting in a natural-feeling interaction with the device.

There are a number of use cases that can benefit from an edge-based VUI. Control panels for security systems and other always-listening devices may employ on-device voice commands so that there is no reliance on cloud services. This allows the panel to perform critical functions, such as arming or disarming, without being affected by network interruptions or service outages. Any air gapped system, such as closed-circuit cameras, can also employ a VUI without sacrificing device security, with voice commands that are processed locally. With the increased responsiveness of a VUI that directly processes commands, devices such as lighting control panels can operate via voice in way that matches the immediacy of toggling a physical switch and yet can be controlled from anywhere in the room.

Assistive devices can greatly benefit from the use of edge-based VUI, especially in cases where the user relies on voice control functionality for tasks that they would otherwise be unable to complete due to physical limitations such as visual impairment or dexterity issues. Reliability is key and using a device with local voice control means that its assistive functionality is unaffected by any network instability or third-party service interruptions.

Edge-based voice control requires a capable audio front end (AFE) comprising the input stage of the VUI. With microphone arrays and signal processing to reduce interfering noise, the AFE must deliver a clean and intelligible signal to the speech recognition engine. Voice-controlled products, whether the VUI employs local processing alone or relies on cloud services, require an AFE that can be scaled to match the physical constraints, available processing power, and use cases of the device.

A well-designed AFE often includes features such as quiescent sound detection, enabling the device to remain aware while in a low-power state; multi-microphone arrays with beam forming, to eliminate interfering noise by narrowing its focus to the direction of spoken commands; noise reduction, to eliminate stationary noises such as appliance hum or fan noise; and automatic echo cancellation, which removes echoes caused by acoustic coupling between the built-in speaker and microphones.

Such an AFE is responsive to voice commands from a user positioned in the far field (several meters away from the device) even in problematic acoustic environments. Once a voice signal is received and processed by the AFE, it is passed on to the VUI itself.

Audio Weaver

DSP Concepts has developed Audio Weaver, the development platform for the Audio of Things, which offers a comprehensive set of embedded audio processing technologies with proven and easy-to-use tools to design, test, and deploy a full range of sound and voice features to various products. As a hardware-independent platform, designs created in Audio Weaver can be developed with or without target hardware, and then deployed when ready to a target MCU, SOC, and DSP, without the need for redesign. With algorithms developed in-house and by third parties, Audio Weaver is a powerful and flexible solution that can streamline the entire development workflow. Audio Weaver supports a wide range of instruction sets, and the binaries have been optimized for use with most major silicon vendors.

The functionality of Audio Weaver helps product makers approach the future by facilitating rapid innovation and mitigating risk. Designs are created by placing the signal processing building blocks known as modules on a virtual canvas, connecting them with virtual wires, and adjusting module properties to tune the design. Designs can then be auditioned from within AWE Designer using the PC's sound card. Multiple team members can each approach the creation and tuning of different portions of the design concurrently, developing features in parallel and later combining them into a final design. With this collaboration and the ability to quickly and seamlessly test iterations and new designs, the entire process is streamlined.

Edge IP

A variety of terms may describe the capabilities of the speech recognition engine, depending on the specifics of its functionality. Wake words (sometimes

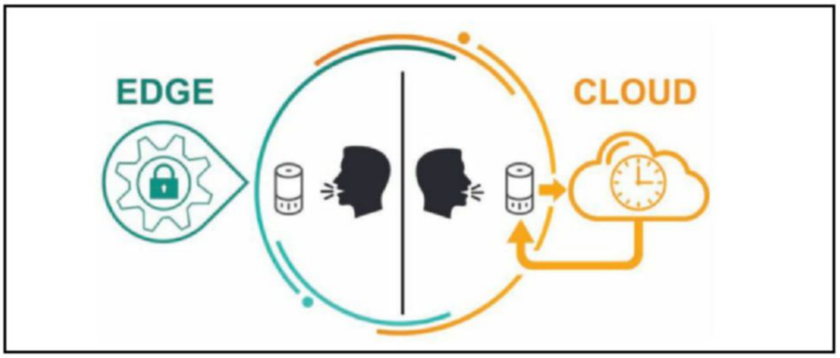

Voice-enabled products must have reliable performance and be adaptable to different languages or geographical regions.

MCU's and Application Processors	Audio Weaver Libraries
ARM® Cortex®-M4, -M7, -M33	Optimized
ARM Cortex-A	Optimized
Cadence® HiFi 2 and HiFi 3 (Fixed-point only)	Optimized
Cadence HiFi 4	Optimized
Cadence HiFi 5	Optimized
ADI SHARC®, SHARC+®	Optimized
Texas Instruments C66x	Optimized
CEVA-X2™	Optimized
Qualcomm® Hexagon™	Optimized

Audio Weaver supports a wide range of instruction sets, and the binaries have been optimized for use with most major silicon vendors.

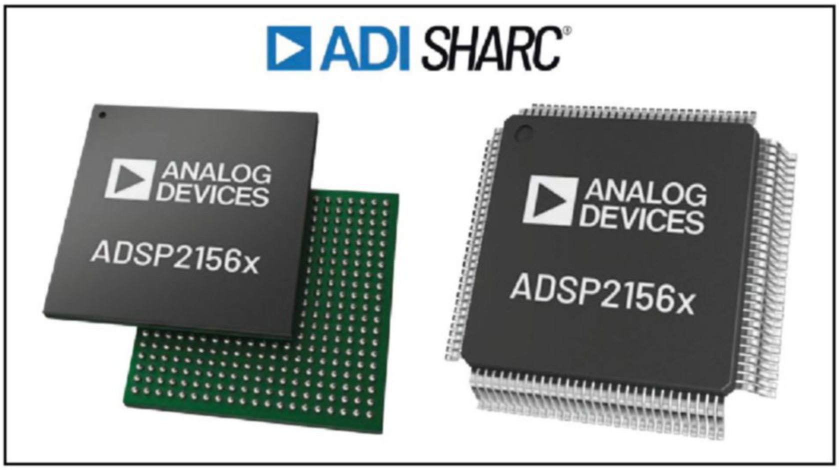

The ADSP-2156x series of processors from Analog Devices features a SHARC+ DSP core, upgraded from its predecessor SHARC core to natively support double-precision floating-point operations. The ADSP-2159x series adds a second SHARC+ DSP core to double the performance. Ideal for consumer, pro audio, and automotive audio applications, DSP Concepts worked in collaboration with Analog Devices to optimize the performance of Audio Weaver Core runtime libraries to take full advantage of the SHARC+ core and to leverage the powerful SHARC+ DSP FIR/IIR hardware accelerators.

called "trigger words") describe the keyword or phrase used to prompt the VUI into a listening state. Once activated, the VUI may employ phrase-spotting processing to isolate relevant commands, speech-to-intent processing, which can deconstruct syntax to identify variable or synonymous phrasing, or even natural language processing that utilizes embedded AI to distill commands from fluid speech.

One of the challenges of embedding these functions on the device itself is choosing a flexible solution with high accuracy, which makes efficient use of the available processing power. Companies such as DSP Concepts, Sensory, and Fluent.ai have approached the edge market with a variety of technologies that help product makers design their own products to include sophisticated edge-based VUI.

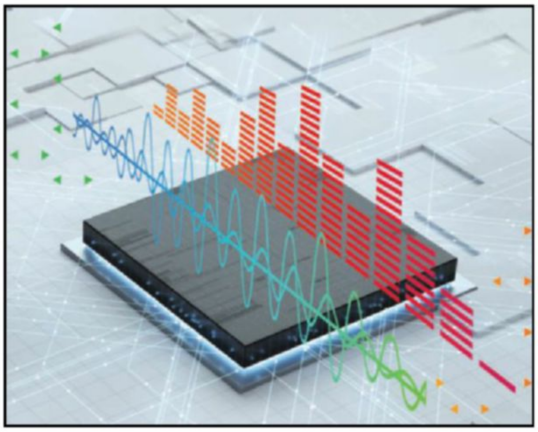

Blending a neural network (NN) with digital signal processing, the Cadence Tensilica HiFi 5 DSP easily executes the two complementary aspects of speech recognition: audio pre-processing and speech recognition/keyword detection. This is a platform that is able to power edge-AI and rich voice interactions for appliances, smart home or automotive dashboard control.

DSP Concepts offers up the Audio Weaver embedded platform that incorporates the TalkTo AFE. With a selection of flexible microphone arrays and a host of advanced signal processing, TalkTo allows the easy addition of far-field voice control to product designs created in Audio Weaver.

TrulyHandsfree and TrulyNatural are technologies from Sensory that span dozens of languages and offer wake-word detection, phrase spotting libraries, plus a natural-language-understanding engine that is ideal for large-vocabulary continuous speech recognition.

Fluent.ai's Automatic Intent Recognition (Air) is a direct speech-to-intent spoken language understanding system. With edge-based universal language support, Air utilizes on-device command sets with highly accurate and concurrent understanding of multiple languages and accents, with no reliance on the cloud.

With these tools, it becomes much easier and economical to develop an edge-based VUI that performs well, can utilize expansive and customize command sets, and has built in support for various languages.

Embedded Models

Embedding the voice recognition engines in the devices themselves requires an efficient use of resources—especially with low-power devices such as wearables and hearables. The edge-based IP available in Audio Weaver is efficient and scalable based on implementation and the number of embedded commands.

Fluent.ai's Air operates by directly extracting intent from speech, directly distilling the intended action from input speech, and can be trained for use with any accent or language. This approach eliminates the need to utilize a speech-to-text element before performing natural language processing.

Fluent.ai provides benchmarks for the Arm Cortex-M4 platform showing that the Air system with custom, low-power keyword spotting and on-device voice control with up to 1,000 commands requires the use of only 99KB RAM and 992KB Flash memory and can be executed with merely 100MHz processing headroom. This is made possible by the use of a voice-to-intent neural network, which operates in overlapping intervals due to the real-time nature of voice control utterances. Air applies weighting and processing to each frame then passes an output vector to be combined sequentially. This results in a short decoding time once Air has determined that a valid intent has been received, since the latency depends in part on the length of the most recent interval rather than the length of the entire speech command.

Sensory TrulyHandsfree can utilize two methods of wake word training when embedding the IP. With a predefined wake word, these models have the option of being fixed, with a pre-trained and speaker-independent model right out of the box for accurate performance based on the target demographic and use case; or enrolled, where the model is trained by collecting recordings of the end user uttering the set wake word.

In the listening window immediately after detection of a valid wake word, the command set model listens for valid commands from sets that can be customized by the product maker using the VoiceHub tool or the Sensory linguistics process, as needed for product requirements and use case. Sensory also makes models available for security purposes such as Speaker Verification (SV) and Speaker Identification (SID) for secured wake words that can identify a specific user based on individual voices.

Sensory's benchmarks of its TrulyHandsfree IP running on an Arm Cortex-M4 platform vary depending on the employed wake word model, with a single fixed wake word model running at a nominal 21.6 MIPS (23.1 peak MIPS) and occupy only 123KB flash memory and less than 15KB RAM. When embedded, single enrolled wake word models run at a nominal 22.1 MIPS (53.1 peak MIPS) and occupying only 112KB flash memory and 12KB RAM. Even with other models such as Low Power Sound Detector coexisting on the same device, TrulyHandsfree occupies a footprint as large as 2MB or as small as 20KB depending on the chip and use case, while still performing well in the key metrics of false acceptance and false rejection.

Conclusion

Due to their operation either offline or at the periphery of the network, edge-based voice-enabled devices built with these technologies can operate as contained systems that do not require transmission to and reception from an external online service. This method of local operation significantly decreases processing delay in recognizing and processing commands, allowing for a seamless user experience that is private by design. With powerful audio processing and the third-party voice recognition solutions available on multiple platforms, product makers can now easily incorporate this technology into a range of devices that span multiple market segments.

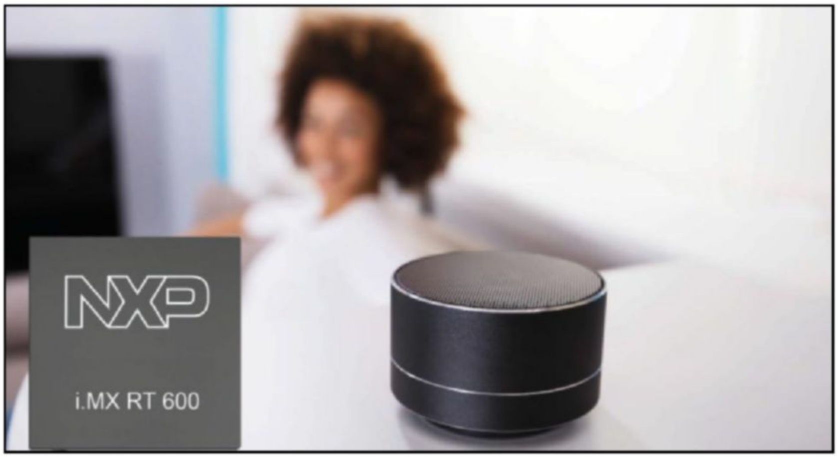

NXP i.MX RT600 crossover MCUs are optimized for 32-bit audio playback and voice user interface applications combining a high-performance Cadence Tensilica HiFi 4 audio DSP core and Arm Cortex-M33 core. NXP partnered with Alango Technologies, DSP Concepts, and Sensory to provide high performance voice pre-processing and recognition software, plus professional audio libraries and tools.

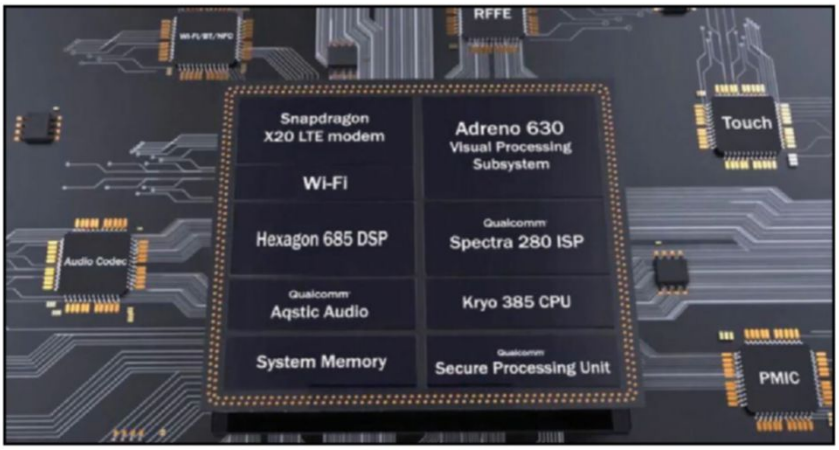

The Snapdragon 845 system-on-chip is an advanced solution combining powerful CPU cores with the Hexagon 685 DSP tailor made for artificial intelligence and machine learning. The Hexagon DSP uses a more efficient instruction set architecture (ISA) with zero-latency thread scheduling, meaning it delivers excellent performance for audio processing tasks.

About the Author

Matthew Mitschang is a Technical Content Specialist at DSP Concepts. Focusing largely on audio-related technical publications, Matthew has a deep appreciation for documentation and all things related to sound. He has spent the past decade writing certification requirements, user manuals, and other fascinating content for THX, Meyer Sound Laboratories, and Avid Technology.

Resources

ABI Research, "Mobile Accessories and Wearables Market Share and Forecasts," www.abiresearch.com/market-research/product/7780556-mobile-accessories-and-wearables-market-sh

Juniper Research, "Hearables: Emerging Opportunities, Competitor Leaderboard & Market Forecasts 2022-2026," www.juniperresearch.com/researchstore/devices-technology/hearables-market-research-report

Juniper Research, "Voice Assistants: Monetisation Strategies, Competitive Landscape & Market Forecasts 2021-2026," www.juniperresearch.com/researchstore/devices-technology/voice-assistants-market-research-report

How to Choose the Right Signal Processing Technique

By
Dave Betts
(AudioTelligence)

Noise is all around us—at work and at home—making it difficult to pick out and clearly hear one voice amid the cacophony. Electronic devices have the same issue—the audio signals picked up by microphones are often contaminated with interference, noise, and reverberation. Signal processing techniques, such as beamforming and blind source separation, can come to the rescue. But which should you choose for which applications—and why?

"Sorry – can't hear you—someone at the next table is talking really loudly."
"Hang on, there's a plane going over..."
"Just a minute, I'm going to a quieter room..."

Sound familiar? The number of phone and video calls we make is increasing all the time, so we are all finding ourselves saying things like this more and more often. Thanks to technology and to widespread network coverage, we can now make calls from almost anywhere, using our phone, laptop, or tablet. But that often means we are doing so in noisy cafés, on trains and buses, or perhaps while walking down the street. And, if we work in an open-plan office, we may have to contend with the background chatter from our colleagues' calls and meetings. Our world has grown noisier, and although the human brain is very clever and can pick out the voice we want to hear even in a cacophony of competing voices and noise, it is still difficult for many of us to hear the voice we want to listen to as clearly as we would like.

Electronic devices, which also need to pick up audio signals clearly, have the same issue: In an increasingly noisy world, the signals reaching their microphones are a mix of the relevant voice, background noise, room reverberations, and other interference. This means the quality and intelligibility of the speech that the devices are designed to capture can be badly affected, leading to poor performance. Intelligible speech is crucial for a huge range of modern technology—not just the phones and computers we use for calls and VoIP, but also for conferencing, transcription, car infotainment, home assistants, and of course, hearing assistance.

Signal processing techniques (e.g., beamforming and blind source separation) can help—but they have different benefits and drawbacks. So which technique is best for which application?

Beamforming Uses

Audio beamforming is a technique that has been available for a long time, and it is one of the most versatile multi-microphone methods for emphasizing a particular source in an acoustic scene. Over the years, many different types of beamformers have been developed and they can be divided into two types, depending on how they work: data-independent or adaptive. One of the simplest forms of data-independent beamformers is a delay-and-sum beamformer, where the microphone signals are delayed to compensate for the different path lengths between a target source and the different microphones. This means that when the signals are summed, the target source coming from a certain direction will experience coherent combining and it is expected that signals arriving from other directions will suffer to some extent from destructive combining.

However, in many audio consumer applications, these types of beamformers will be of little benefit. Why? There are a couple of reasons.

The first has to do with the size of the array compared with the frequencies of normal speech. For delay-and-sum beamforming to do a good job, it needs the wavelength of the signal to be comparable with the size of the microphone array. The wavelengths in audio range from millimeters to meters and, from the physics of antenna theory, we know that the optimal microphone spacing is one quarter of the wavelength of the sound you wish to receive. The low frequencies of speech have a very large wavelength—measured in meters. This is why top-of-the-range beamforming conference microphone arrays are typically 1m in diameter and have hundreds of microphones, which allow them to cover the wide dynamic range of wavelengths. They can work very well but, of course, are very expensive to produce and are suitable for the business conferencing market only.

When it comes to devices designed for consumers, they usually have only a few microphones in a small array. In these use cases, the delay-and-sum beamformer really struggles as it is contending with the large wavelengths of speech arriving at a small microphone array. A beamformer the size of a normal hearing aid, for example, cannot give any directional discrimination at low frequencies—and at high frequencies it is limited in its directivity to a front/back level of discrimination.

Another problem relates to the way sound behaves. It does not move in straight lines: A given source has multiple different paths to the microphones, each with differing amounts of reflection and diffraction. This means that simple delay-and-sum beamformers are not very effective at extracting a source of interest from an acoustic scene. But they are very easy to implement and do give a small amount of benefit, so they were often used in older devices.

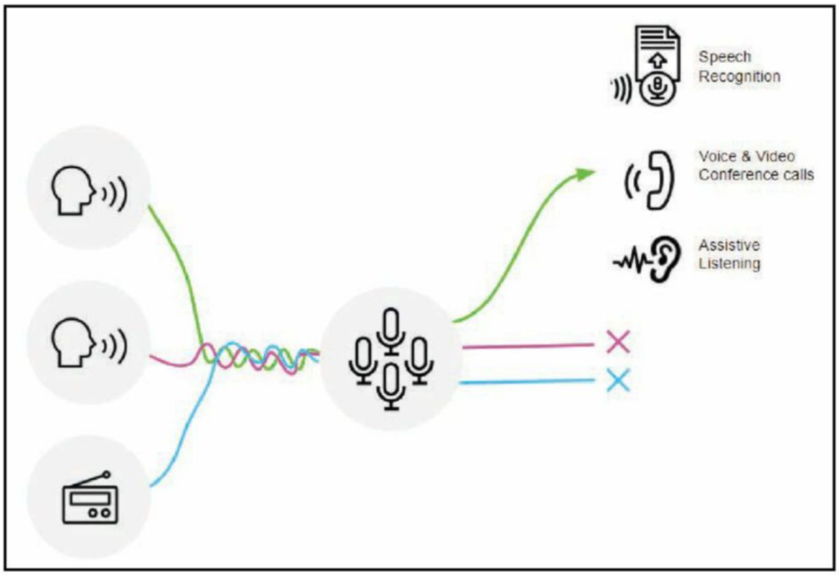

Figure 1: AudioTelligence's AISO technology uses Blind Source Separation (BSS) to improve clarity and intelligibility of speech.

Over the years, many more advanced beamforming techniques have been developed and are now available. One of the most well-known adaptive beamformers is the minimum variance distortionless response (MVDR) beamformer. This tries to pass the signal arriving from the target direction in a distortionless way, while attempting to minimize the power at the output of the beamformer. This will have the effect of trying to preserve the target source while attenuating the noise and interference.

This technique can work well in ideal laboratory conditions but, in the real world, microphone mismatch and reverberation can lead to inaccuracy in modeling the effect of source location relative to the array. The result is that these beamformers often perform poorly because they will start cancelling parts of the target source. Of course, a voice activity detector can be added to address the target cancellation problem, and the adaptation of the beamformer can be turned off when the target source is active. This can perform as desired when dealing with just one target source but, if there are multiple competing speakers, this technique has limited effectiveness.

And, again, MVDR beamforming—just like delay-and-sum beamforming and most other types of beamforming—requires calibrated microphones as well as knowledge of the microphone array geometry and the target source direction. Some beamformers are very sensitive to the accuracy of this information and may reject the target source because it does not come from the indicated direction.

More modern devices often use another beamforming technique called adaptive sidelobe cancellation, which tries to null out the sources that are not from the direction of interest. These are state-of-the-art in modern hearing aids and allow the user to concentrate on sources directly in front of them. But the significant drawback is that you must be looking at whatever you are listening to, and that may be awkward if your visual attention is needed elsewhere—for example, when you are paying for your drink but talking to someone next to you in the line.

The BSS Route

Luckily, there is another way of improving speech intelligibility in noise—blind source separation (BSS). BSS is a family of techniques which has been the subject of scholarly articles and laboratory experiments for some time—but here at AudioTelligence we have developed BSS algorithms, which are successfully running on products in the real world. You can hear the difference between BSS and beamforming here: https://bit.ly/3QQ4wip

How does BSS work? There is more than one way of performing BSS. Time-frequency masking estimates the time-frequency envelope of each source and then attenuates the time-frequency points that are dominated by interference and noise. At AudioTelligence, we prefer to use another method—linear multichannel filters. We separate the acoustic scene into its constituent parts by using statistical models of how sources generally behave. BSS calculates a multi-channel filter whose output best fits these statistical models. In doing so, it intrinsically extracts all the sources in the scene, not just one.

We chose to use this method (**Figure 1**) because it can handle microphone mismatch and will deal well with reverberation and multiple competing speakers. Importantly, it does not need any prior knowledge of the sources, the microphone array, or the acoustic scene, since all these variables are absorbed into the design of the multi-channel filter. Changing a microphone, or a calibration error arising, simply changes the optimal multi-channel filter.

Because BSS works from the audio data rather than the microphone geometry, it is a very robust approach that is insensitive to calibration issues and can generally achieve much higher separation of sources in real-world situations than any beamformer. And, because it separates all the sources irrespective of direction, it can follow a multi-way conversation automatically. This is particularly helpful for hearing assistance applications where the user wishes to follow a conversation without having to interact with the device manually. BSS can also be very effective when used in VoIP calling, home smart devices and in-car infotainment applications.

Sounds easy? Not at all. Developing a BSS solution that works in real products, not just in the laboratory, is very hard—it has taken us years of research and experimentation to create, develop, and test our algorithms. The journey started many years ago, before the advent of the voice recognition revolution. I was working at CEDAR, a company specializing in audio restoration and dialog noise suppression.

We were looking for a long-term research project and identified audio BSS as a key technology to investigate. We predicted that the increase in compute performance and power within a 5- to 10-year horizon would allow the algorithms necessary for BSS to be implemented successfully, although it was not feasible with the state of compute performance at that time. We also identified several key problems that we felt we could solve using some of our proprietary know-how. In particular, we thought we could solve the sub-10ms algorithmic latency required by real-time hearing assistance applications, as well as improving both CPU and acoustic performance.

That decision led to 10 years of research and the creation, development, and testing of proposed solutions. We finally had one we were happy with, but CEDAR considered that turning our algorithms into a product would take the company in a new strategic direction and that it would be better for this to be done in a new company. And so AudioTelligence was born as an independent company—a colleague and I moved to AudioTelligence, and, together with the team we recruited, we have taken the solution from algorithms to actual products.

But BSS isn't without its own problems. For most BSS algorithms, the number of sources that can be separated depends on the number of microphones in the array. And, because it works from the data, BSS needs a consistent frame of reference, which currently limits the technique to devices which have a stationary microphone array—for example, a tabletop hearing device, a microphone array for fixed conferencing systems, or video calling from a phone or tablet, which is being held steady in your hands or on a table.

When there's background babble, BSS will generally separate the most dominant sources in the mix, which may include the annoyingly loud person on the next table. So, to work effectively, BSS needs to be combined with an ancillary algorithm for determining which of the sources are the sources of interest—we combine it with our conversational dynamics algorithm which dynamically follows those sources in a conversation.

BSS on its own separates sources very well, but does not reduce the background noise by more than about 9dB. To obtain really good performance, it has to be paired with a noise reduction technique. Many solutions for noise reduction use AI—for example, it's used by Zoom and other conferencing systems—and it works by analyzing the signal in the time-frequency domain and then trying to identify which components are due to the signal and which are due to noise. This can work well with just a single microphone.

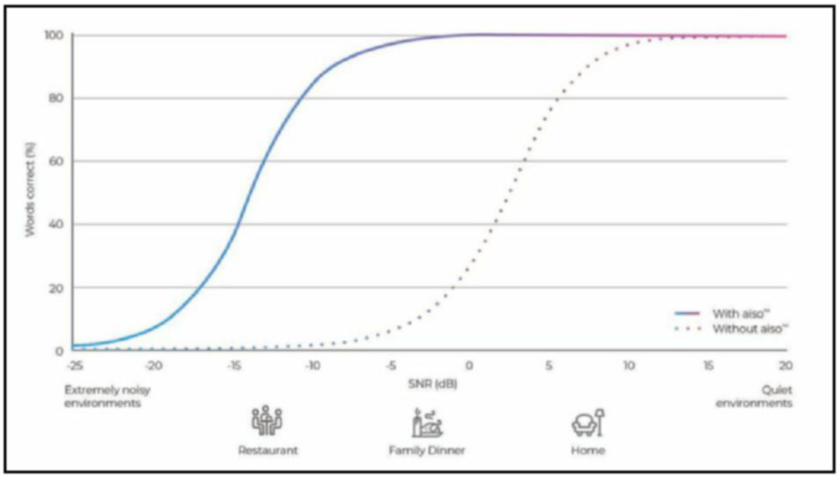

Figure 2: AudioTelligence's AISO technology speech understanding performance

But the big problem with this technique is that it extracts the signal by dynamically gating the time-frequency content, which can lead to unpleasant artefacts in poor signal-to-noise ratios (SNRs), and it can introduce considerable latency.

Enter AISO

At AudioTelligence, we wanted to develop a noise reduction technique to combine with our BSS, which worked well with multi-microphone arrays and did not introduce latency or artifacts. The result is our low-latency noise suppression algorithm which, combined with BSS in our AISO software solution, gives up to 26dB of noise suppression and makes our products suitable for real-time use. Hearing devices, in particular, need ultra-low latency to keep lip sync—it is extremely off-putting for users if the sound they hear lags behind the mouth movements of the person to whom they are talking. This is a difficult problem—but one which we've managed to solve. Despite the complex calculations performed, the underlying algorithms in AISO combined have a latency of just 5ms—which is essential for real-world use—and they produce a more natural sound with fewer distortions than AI solutions.

But when it comes to distinguishing speech in noise, is there an objective measure of performance? One which is commonly used is to measure the percentage of words correctly understood in different SNRs. At -5dB SNR (a typical family dinner), a top-of-the-range hearing aid typically improves speech understanding to 50%—and a top-of-the-range assistive listening device to 80%. **Figure 2** shows the performance of the AudioTelligence AISO algorithms, illustrating an improvement of speech understanding from 5% to 98% at -5dB SNR.

So, BSS can perform really well in increasing the intelligibility of speech in noise. But there is one issue we still have not discussed: echo. I think we have all occasionally experienced the huge irritation of hearing echoes from speaker phones and other devices, when the phone's microphone picks up your voice from the loudspeaker. This echo can occur in other situations as well. The solution for the last 70 years has been to use an adaptive filter to try to cancel out the echo. Many different algorithms have been used, but they all work by trying to predict the echo from the signal driving the loudspeaker. The better they make that prediction, the better they can cancel the echo.

Today, thanks to the recent increases in compute power, more sophisticated signal processing techniques—such as recursive least squares and minimum mean square error—can be used. These algorithms are all forms of linear prediction, and they all struggle with nonlinear distortion. For modern acoustic echo cancellers it is these nonlinear distortions that limit the performance.

At AudioTelligence we needed to solve the echo distortion problem to ensure that our BSS could be used effectively where there is echo. The important property we use is that the distortion occurs mostly in the loudspeaker. This means that, to our BSS, it looks just like any other acoustic source. So, BSS will automatically remove the nonlinear distortion as an unwanted source. This doesn't require any accurate modeling of the nonlinear distortion—we get this effect free as part of BSS.

This will work with most acoustic echo cancellation (AEC) algorithms. Our AEC is particularly powerful as it deals with multiple references such as 5.1 surround sound as well as acoustic path lengths of up to 125ms. It is also designed to preserve the relationships necessary for BSS to work effectively. This ensures that we remove as much of the linearly predictable part of the echo as possible while leaving the residual intact for BSS to perform its magic. The BSS then sweeps up any residual distortion.

Signal processing has come a long way since I started working in audio many years ago. There are multiple techniques from which to choose, all of which are becoming more sophisticated and complex. Selecting the right technique for your application requires consideration not just of the performance you need, but the situation in which you need the application to work, and the physical constraints of the product you have in mind.

About the Author

Dave Betts is the Chief Science Officer at AudioTelligence. He has been solving complex problems in the field of audio for more than 30 years, and his experience ranges from audio restoration and audio forensics to designing innovative audio algorithms used in a huge number of blockbuster films. At AudioTelligence, Dave leads a team of researchers delivering innovative commercial audio solutions for the consumer electronics, assistive hearing, and automotive markets.

Making a DC Power Supply with Audible Benefits

Part 2 — Real Improvements

By
Frans de Wit

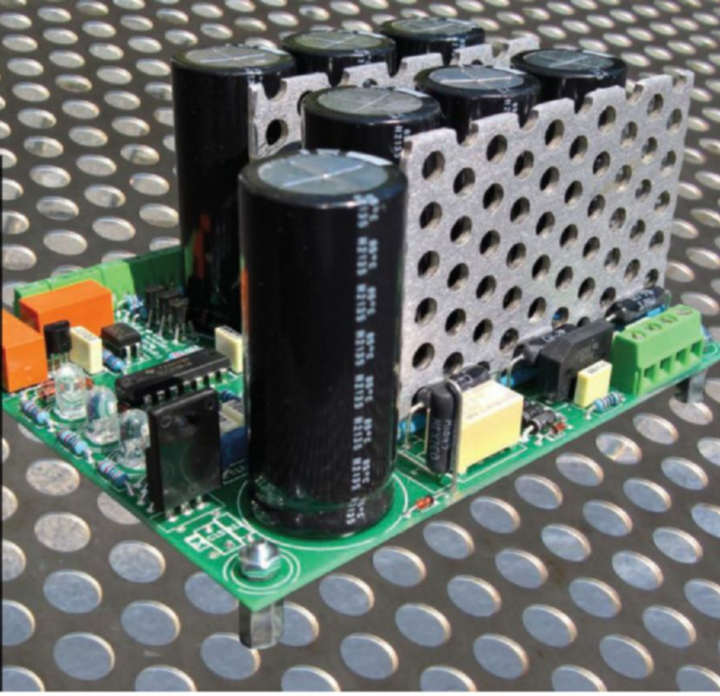

In this two-part article, Frans de Wit considered all the factors and design options for a high-quality power-supply especially targeted toward preamplifiers, record-players, and other low power appliances. In this continuation, he performs simulations that demonstrate the real improvements he achieved.

In this part of the article, we will discuss real improvements and not just the use of a bigger capacitor. So let's improve on the previous simulated circuit. To do that we first need to have one final look at the old circuit (as previously shown in Part 1 of this article), so we know what we are doing. To do this, we need to stabilize the running circuit in such a way that it is possible to "catch" a valid FFT. To enable that we need to wait until the output voltage is stable (e.g., the capacitors are fully loaded [kind of]).

This stabilizing is done by starting the measurement with "soft-start" period of 500ms. After this waiting time, waiting before measuring of 1500ms is added and finally a collect (collecting data, in other words measuring data) of 200ms is executed.

First about the graphs shown in **Figure 5a**, the circuit is unchanged. When we start looking at improvements, it would help if we first looked at where we started. For this we need a few new measurements and more detail. The additional detail is created by a more in-depth measurement. Changes include:

Before) .tran 0 {tSoft+tWait+tCollect} {tSoft+tWait} 1u uic.param tSoft 0ms, .param tWait 0ms and .param tCollect 100ms
After) .tran 0 {tSoft+tWait+tCollect} {tSoft+tWait} 1u.param tSoft 500ms, .param tWait 1500ms and .param tCollect 200ms

Look at the ".tran" statement. It instructs the simulator of the timing and detail of the simulation to perform, in the first line '1u uic' sets the detail level to 1µs steps, giving detail up into the 1MHz range. Leaving out the 'uic' part of the statement instructs the simulator to no longer "Skip initial operating point solution" (e.g., do not try to stabilize the DC behavior of the circuit) and take measurements immediately after starting the simulation (e.g., calculate the DC operating point of the circuit before starting). This will enable measuring the initial (startup) state of the circuit.

The presented '.param' statements are previously 'tSoft 0ms' and now 'tsoft 500ms' creates a soft-start period of 500ms and helps to stabilize the circuit for a more precise measurement. Next, the 'tWait 0ms' is changed to 'tWait 1500ms' adding more time for the circuit to stabilize. And finally 'tCollect 200ms' changes the measurement period from 100ms to 200ms, giving us more measurement data for a higher resolution result.

As far as the graphs shown in Figure 5a, the top pane (box/window) is as before, primary current, secondary current and output wattage. The output voltage detail is being used for the FFT. And, the output current is shown as it flows through the load.

Looking at the added graphs for Figure 5a, we can conclude the following. In the left blue graph we see 670mVPP triangular waves on top of the 66.7V DC output voltage. The next curve, gray, shows us that the loading pulses of the output capacitor are over 4A. The right graph (a linear FFT result) shows that the noise at 100Hz is 136mV and at 1.05Khz it is near 6mV, which doesn't look too bad.

However, taking those start-up pulsed currents that were shown in the first two simulations, together with the continuously pulsing loading currents make this an undesired, even unlikeable circuit. This has to be improved (e.g., using just a large capacitor is not the solution for the PSU problem). P.S. "Not too bad" is not the correct phrasing. For our high-end mains supply, we have to correct this, it is really bad and for our purposes, unacceptable.

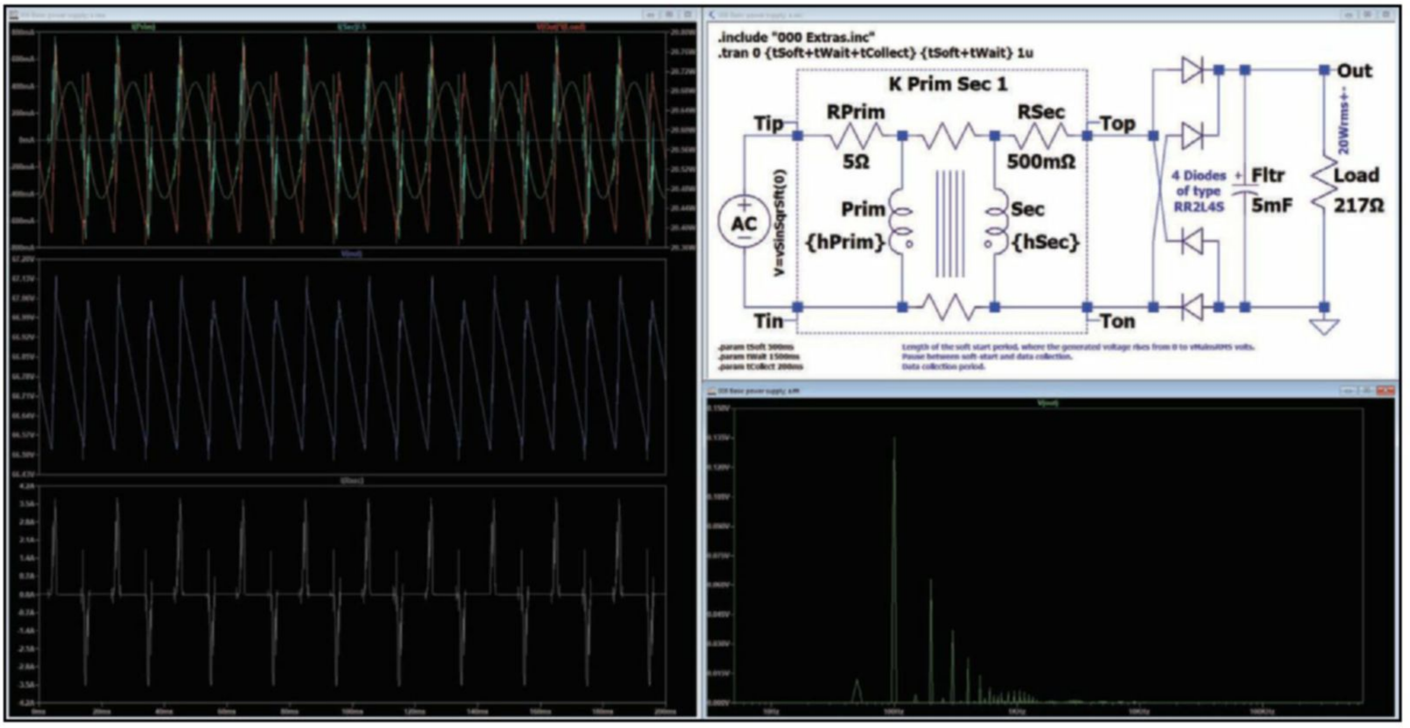

Figure 5a: The circuit is unchanged, but additional detail is show. The top pane (box/window) is as before, primary current, secondary current and output wattage. The output voltage detail is being used for the FFT. And, the output current is shown as it flows through the load

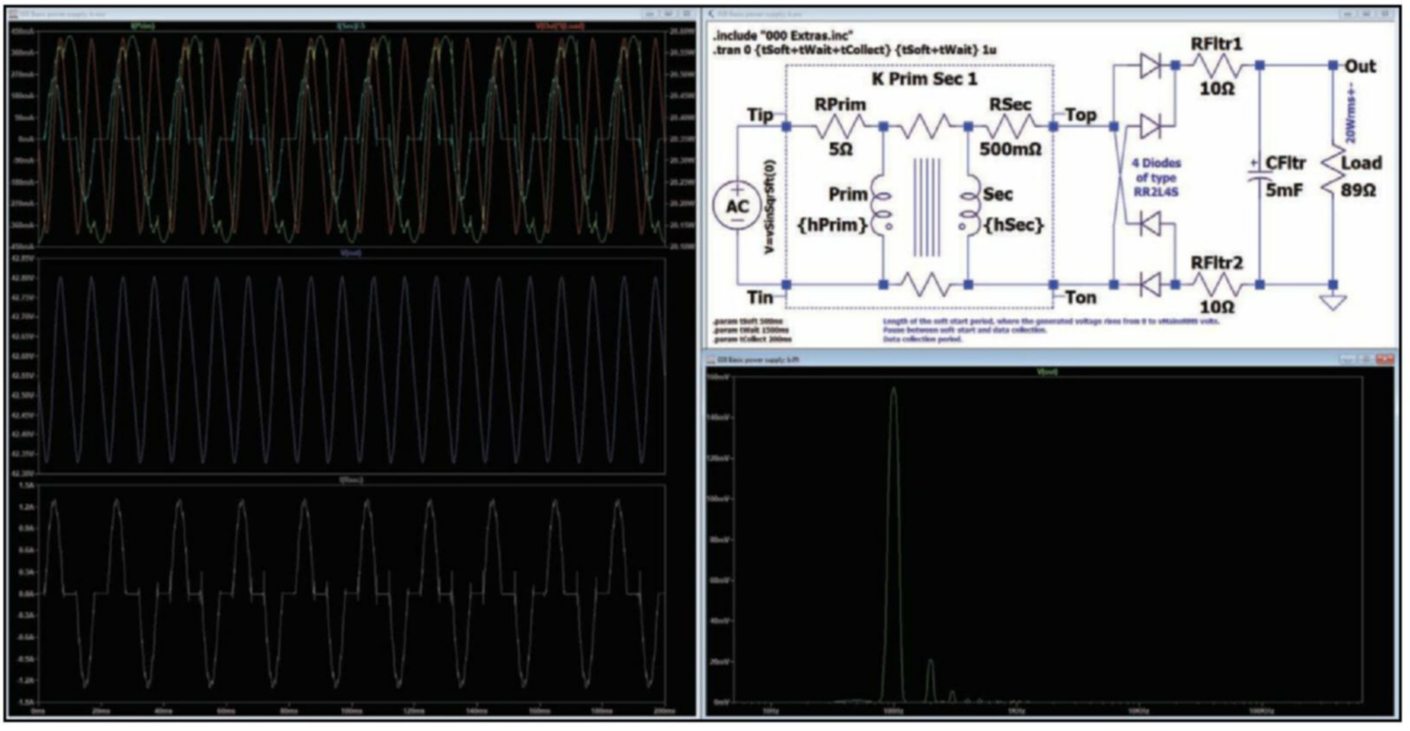

Figure 5b: In this variation, we chose some real filtering by adding RFltr1 and RFltr2, and keeping the 5mF capacitor.

Solving the Problems

Next, we chose some real filtering by adding RFltr1 and RFltr2, and keeping the 5mF capacitor. We will need to add some serious storage and, later we talk about that when making this into a real buildable power supply, while at the same time, improving the power factor (e.g., removing those high current pulses), but more on that later. Figure 5b shows what it looks like.

First, a short note about the large 5mF capacitor. In case of

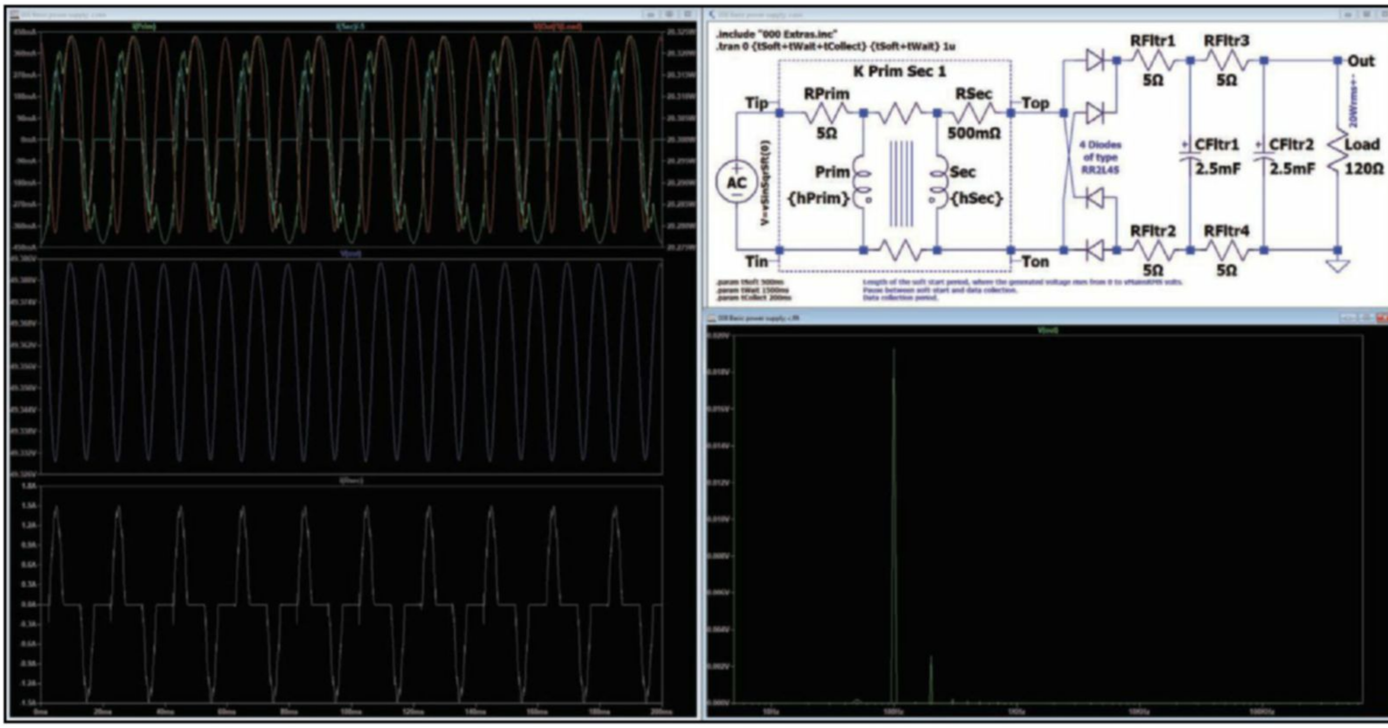

Figure 5c: The solution is not that obvious, but works like this: The resistors are split into four 5Ω resistors, and the capacitor is split into two 2.5mF capacitors, the same total values, 20Ω and 5mF) but a completely different result.

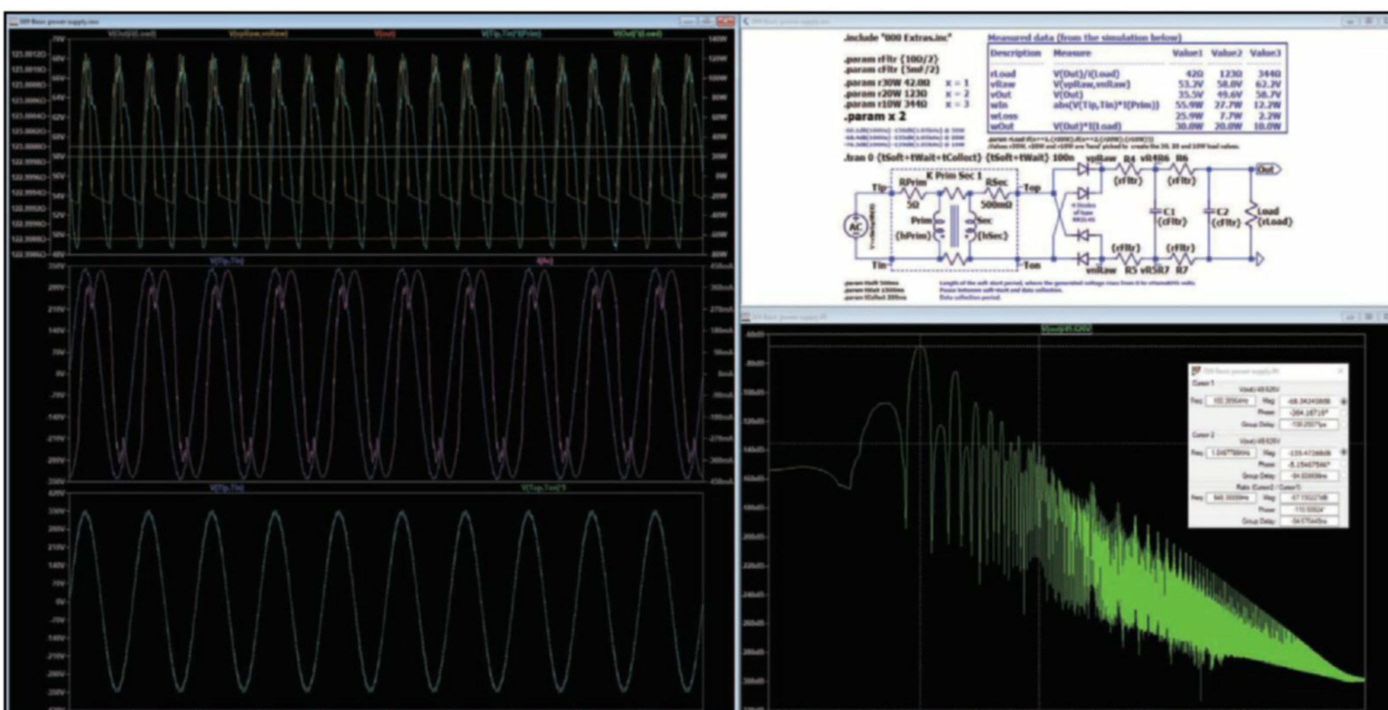

Figure 6: Here is another option, which is in fact the same circuit as the one previously shown with a few more measurements taken at a higher resolution, making a logarithmic FFT more accurate.

mains drop out, it is desirable to have enough power in storage to be able to do a controlled shutdown. In this case, it allows us to mute the outputs (of the now assumed high-end audio application) before losing all power. The half-time that we gain (e.g., the time that the RC filter calculates to) is $2\pi RC$ ($2\times3.14\times(10\Omega\times2)\times0.005F=630mS$ for a 10Ω). The circuit that is drawing energy from this power supply will have an impedance much higher than 10Ω so the available run-down time will be many times larger than that. I would expect up to 3 seconds of useful power from these capacitors (yes there will be more). P.S. The 10Ω×2 is there while the filter resistors are in series.

The 10Ω×2 resistors are needed to create a useful power storage time, but they come at a cost. The resistors dissipate power and generate heat, however, the power loss and the heat are not wanted, especially in a time where we need to conserve energy. Later we will see that this power loss is controllable and not a problem in the final design. In fact, it will consume very little power.

Still, I need to talk the power supply shown in Figure 5b and the effect of those two resistors. It creates a filter and by that you may think that the output noise will go down. But what happened, the measured data shown in Figure 5a was noise at 100Hz 136mV and at 1.05kHz near 6mV, Now it is noise at 100Hz 155mV and at 1.05kHz, it is 750uV (almost 200 times less). So, although the low frequency noise has gone up slightly, the high frequency noise (1kHz and up) is lower into invisibility (in the graph). This is progress, but we need better.

Adding One More Improvement

To get at a dissipation of 20W, we needed to change the load resistance to 89Ω. If we can generate that same wattage in a higher resistor, the circuit will be more economical (and environmentally friendly). The solution is not that obvious, but it works like this: The resistors are split into four 5Ω resistors, and the capacitor is split into two 2.5mF capacitors, the same total values, 20Ω and 5mF) but a completely different result as shown in **Figure 5c**.

This is kind of an extreme change. The measured data has now changed from noise at 100Hz 155mV and at 1.05kHz it is 750μV to noise at 100Hz 19.4mV and at 1.05kHz it is 13μV. That is at 100Hz 7.7 times and at 1.05Khz at 58 times improvement, later we will see how what that looks like in decibels.

There Is Still More

Let me introduce you to another option (**Figure 6**), which is in fact the same circuit as the one previously shown with a few more measurements taken at a higher resolution, making a logarithmic FFT more accurate. First let me show the differences:

```
.tran 0 {tSoft+tWait+tCollect} {tSoft+tWait} 1µ
.tran 0 {tSoft+tWait+tCollect} {tSoft+tWait} 100n
```

Changing the resolution from 1µs(1MHz) to 100ns (10MHz) and making the simulation many times slower (I now have time for more coffee). The measurements in the upper left pane are shown in **Table 2**.

Description	Measure	Value1	Value2	Value3
rLoad	V(Out)/i(Load)	42Ω	123Ω	344Ω
vRaw	V(vpRaw,vnRaw)	53.2V	58.0V	62.2V
vOut	V(Out)	35.5V	49.6V	58.7V
wIn	abs(V(Tip,Tin)*I(Prim))	55.9W	27.7W	12.2W
wLoss	[wIn]-[wOut]=[wLoss]	25.9W	7.7W	2.2W
wOut	V(Out)*I(Load)	30.0W	20.0W	10.0W

Table 2: The measurements for the circuit shown in Figure 6 are noted in the upper left pane and detailed here.

The other measurements are the same, however, the measurements noted in Table 2 are what's important. And then there is the FFT. As can be seen, the graph in Figure 6 is created for 'V(out)/49.626V' (that is the output voltage divided by the output voltage, normalizing the FFT to 1, which makes the results more easily understandable). The 100Hz FFT point is at -68dB. That is a significant number, making for a more then reasonable power supply. The 1.05kHz point is at -135dB making me even doubt the real buildability of a PSU at this level, a power supply at -136dB) is a lot. But this is encouraging, don't you think? It gets better, the

About the Author

Frans de Wit was born and raised in the Netherlands. He went to high school to become an electrical engineer, and found his first job in electronics retail, selling components over the counter and advising customers on component selection and application. Here he was spotted by a headhunter and invited to work for the Dutch importer for Motorola, Texas Instruments, Hewlett Packard, and many other great brands of the time. There he started his career as a Desk-Sales Engineer. During this time, he helped customers building trains, organs, milk-floats, and particle accelerators (yes, the one at CERN) among other things. Next, he co-founded an IT company and until 2016 he worked as a programmer and algorithms specialist building operating systems, compilers, and a large DAM system. In 2016 he achieved a long-time ambition and founded an audio research and development company named Signature Origin.

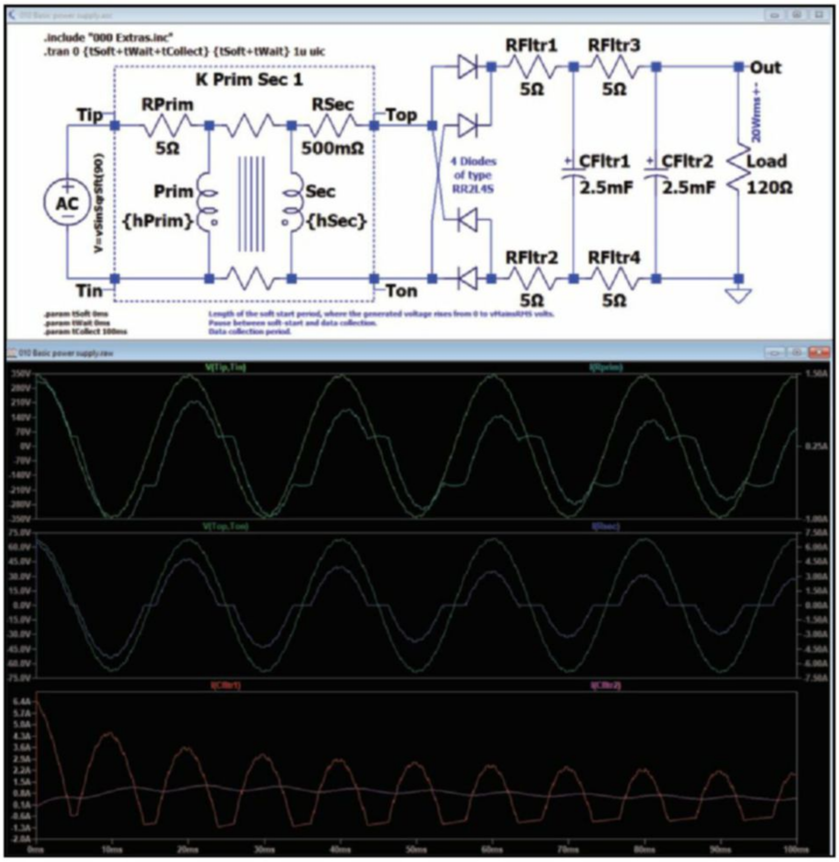

Figure 7: The circuit shown here is a variation of the circuit shown in Figure 5c, with some the measurement conditions as shown in the circuit shown in Figure 4b (shown in Part 1 of this article).

following measurements were also taken using simulated power supply shown in Figure 6:

-60.1dB(100Hz) -130dB(1.05kHz) @ 30W
-68.4dB(100Hz) -135dB(1.05kHz) @ 20W
-76.5dB(100Hz) -139dB(1.05kHz) @ 10W

As we can see, at higher frequencies, the noise goes down into the background noise left over from the big bang, near 300dB, which will not be reached in practice and I think that when using parts of high quality and a design topology that supports high quality, we may be able to (let's be optimistic) reach 150dB. Let me know if you can do better.

If we look at all this data, we can see that a 30W power supply will dissipate lots of power, a 20W power supply is quite okay at 7.7W, and a 10W power supply only dissipates 2.2W. This would be low for an active linear regulator, so this thing is not doing bad.

Notice the value of the rLoad resistor. It is now at 344Ω, about three times higher than the previous value. This leads to a better filter action and lower power dissipation. As can also be seen by looking at the value of wLoss, only 2.2W is lost.

And, before I forget, one more note about the decibel figures, almost -70dB at 100Hz is excellent and the -135dB at 1kHz is even better. Wouldn't it be nice to improve the 100Hz performance to the same level as the 1kHz, and above? Naturally, improvements are always wanted.

But only if it is a real option. Making such an improvement needs larger resistors (RLtr1 to 4) and/or larger capacitors (CFltr1 and 2). Both have their

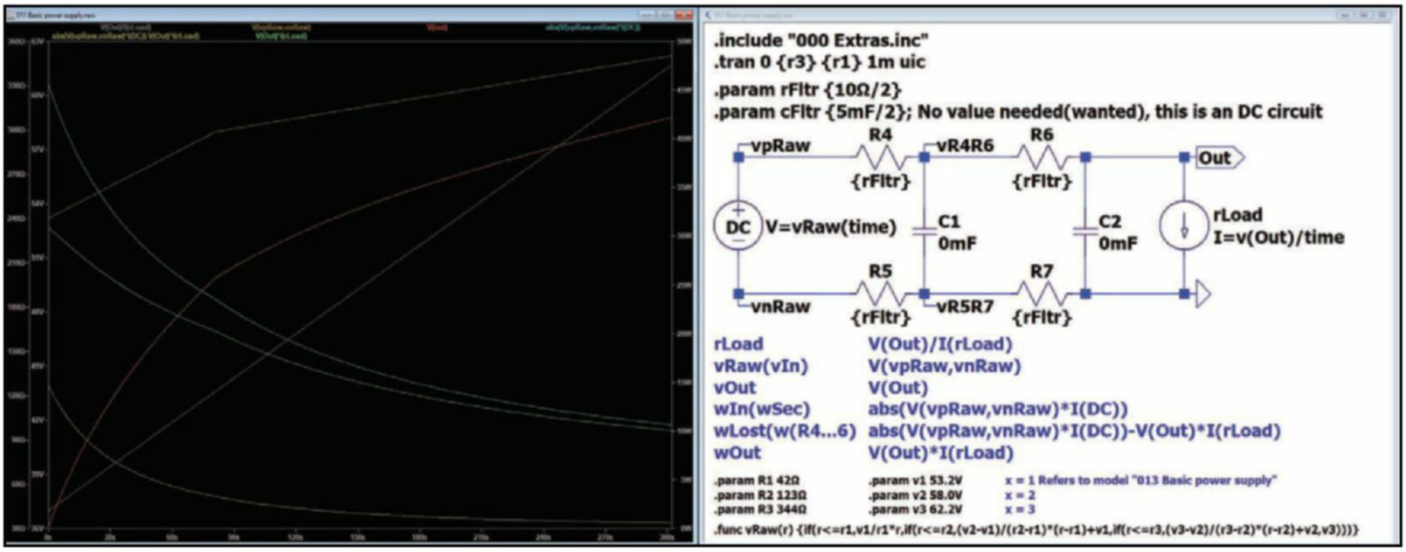

Figure 8: This graph nicely shows the steepness of the dissipation/loss curve (in yellow). It starts off high at 15W and quickly moves down into more acceptable values. The simulation can be used as a reference tool for planning other variants of this circuit. Since it only relies on three data points, its precision is just acceptable, but seeing the curves gives us more insight into the conduct of the circuit under different conditions.

undesirable problems, larger switch-on impulses, a worse power factor, higher dissipation (in case of larger resistors), the circuit gets larger (larger capacitors), and it becomes more expensive.

And yes, we would like to have a better 100Hz filtering value but not at any cost. Like always engineering must weigh the cost and benefits.

On the cost side, are the considerations as just stated—the benefit will be better 100Hz regulation. What we need to consider is that the frequencies in this range are non-transmitting (mostly harmless) and at -70db will do no damage (in this application). Frequencies of 1kHz (and above) will transmit through the circuit but at the level of -135dB should also be considered harmless. Even though it is like that when building the PSU and the high-end application of it, we will take measures to prevent propagation of these signals into the application (and beyond).

Back to the Beginning

Let's refer back to the circuit shown in Figure 5c, with some the measurement conditions as shown in the circuit shown in Figure 4b (the circuit with the worst power-on searching, capacitor loading pulses, and other bad behaviors). The transient command is as before '.tran 0 {tSoft+tWait+tCollect} {tSoft+tWait} 1u uic' and the phase, on the AC power source is set to 90 degrees, fingers-crossed and partying supplies at the ready, let's see the results, which are shown in **Figure 7**.

That is not bad at all, the phase behavior is much improved. The mains voltage and mains current are now about 2.7ms or 48 degrees where it used to be almost 5ms and about 90 degrees in simulation (Figure 5a), so this is a good result. Also we see that the run-in current is now below 1.5A and that is a major improvement, no more fuses will be blown. The next pane shows the primary voltage and current, nicely in phase and well-behaved current peaking, below 3A and at the filter capacitors it gets even better.

Speaking of the capacitors, there we see (third pane red line) a run-in current of 6.4A stabilizing to just above 2A after 100ms. This will improve in time, stabilizing at about 2 seconds. The second line in this pane shows the currents in the second capacitor. There is very low current movement to be seen there, less than 350mA. Remember, it is the current fluctuations that generate magnetic fields, which in turn will cause disturbances/distortion in the application circuit. The general rule is what's not here will not harm, especially the -135dB and more for frequencies above 1kHz, which makes us a proud/happy owner of this circuit and its application.

In Conclusion

I wanted to share an interesting set of graphs, which are shown in **Figure 8**. The data of the circuit shown in Figure 6 has been assigned the parameters R1…3 and v1…3 in addition the transient command is changed to '.tran 0 {r3} {r1} 1m uic' effectively changing the horizontal axes into a facsimile of the resistor value. The displays scale shows 0…300s and actually represents the values R1(42Ω) to R3(344Ω) and the simulator always normalizes the "time" axis to 0s. Thus, there is an offset of 42s(R1).

This graph nicely shows the steepness of the dissipation/loss curve (in yellow). It starts off high at 15W and quickly moves down into more acceptable values. The simulation can be used as a reference tool for planning other variants of this circuit. Since it only relies on three data points, its precision is just acceptable, but seeing the curves gives us more insight into the conduct of the circuit under different conditions.

Although many more analysis can be performed, virtual measurements can be taken, and more paper can be filled, I need to stop somewhere. But do not despair, in the sequel of this article series, I will describe how to turn this data and the benefits we've found into a real circuit. This next article will also include a parts list and circuit board design as well as some images of my actual build.

The Missing Load

My colleague pointed out to me, that something was missing in this section of the article. The person

References

[1] "Luminiferous aether," *Wikipedia*, https://en.wikipedia.org/wiki/Luminiferous_aether#cite_note-1

Resources

"Coil and Transformer Calculator," www.dicks-website.eu/coilcalculator

"LTspice," Analog Devices, www.analog.com/en/design-center/design-tools-and-calculators/ltspice-simulator.html

"Mutual Inductance and Basic Operation, Chapter 9 - Transformers," *Lessons in Electronic Circuits*, Volume 2, All About Circuits, www.allaboutcircuits.com/textbook/alternating-current/chpt-9/mutual-inductance-and-basic-operation

A. Roderick, "Basic Inductance Principles in Transformers," EE Power, June 2021, https://eepower.com/technical-articles/transformer-operating-principles/#

"Series: DD12 IEC Appliance Inlet C14 with Filter, Fuseholder 1- or 2-pole, Line Switch 2-pole," Schurter Electronic Components, datasheet, www.schurter.com/en/datasheet/DD12

"TTSAS0040 - SUPREME AUDIO grade transformer TSAS40VA - voltage to 50V," Toroidy Transformatory Lachowski, datasheet, https://sklep.toroidy.pl/en_US/p/TTSAS0040-SUPREME-AUDIO-grade-transformer-TSAS40VA-voltage-to-50V/379

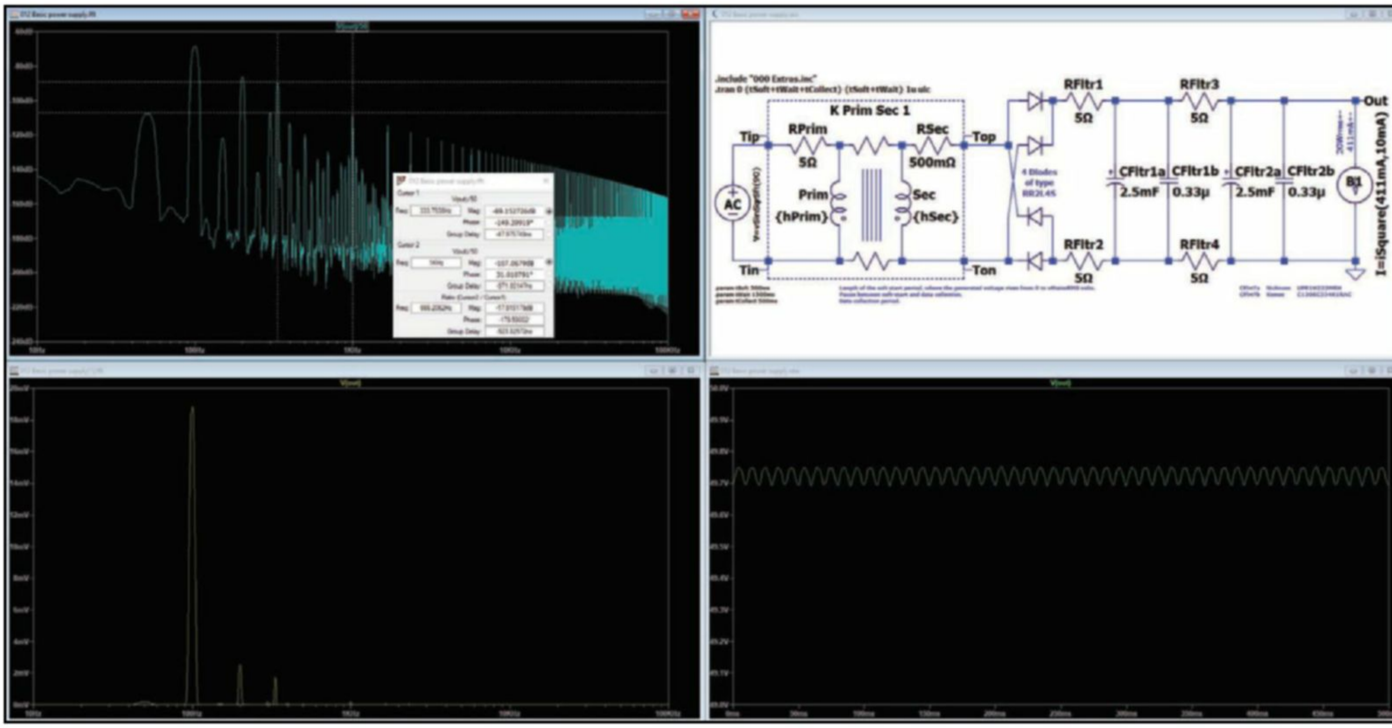

Figure 9: Here is the missing load. CFltr1a and CFltr2a were changed to "real" capacitors (at least as much as possible in a simulated environment) and I also added CFltr1b and CFltr2b, as would be done in an actual build.

asked, "Why did you design this thing (power supply) and how does it perform (quality wise)?" After some deliberation, I had to agree that these are valid questions, so here is the answer: "Not bad at all" or more visually, I can show you in **Figure 9**.

First, I need to explain what was changed. CFltr1a and CFltr2a were changed to "real" capacitors (at least as much as possible in a simulated environment). For this version, I chose the Nichicon UPR1H222MRHs. I also added CFltr1b and CFltr2b, as would be done in an actual build. They are intended to be effective at high frequencies (where the electrolytic capacitors will exhibit a higher impedance). For CFltr1b and CFltr2b, I selected the Kemet C1206C334K1RACs.

I also added B1—this is a current sink that will deliver 411mA (20W). One part of this is 10mA of an AC square wave to represent the AC load (for a mono preamplifier), or for a phono-stage less then 10mA is to be expected (but we exaggerate a bit).

For the the measurements, the right bottom (green) is the residual, mostly 100Hz, noise. We do not worry too much about this. It is only 50mV, that is 1/1000 of the output voltage and (due to its mostly low frequency component) non radiating/transmitting. For one thing, the simulator plots do have the ability to inflate these measurements by (when allowed) spreading them over the full height of the display pane.

The left bottom (yellow) shows a linear FFT, giving us a good idea of the scale of things. At 100Hz, we see 20mV, at 333Hz (the noise at the load) 1.75mV and at 1kHz (the input noise), it is 200µV. When these signals are shown together with the 50V output voltage, they are completely invisible on my 4K video screen.

The usual FFT (blue) shows the FFT normalized to 1V, which equals 0dB. This makes it easily comparable with other FFT readings in decibels. For this reason, the output voltage must be divided by 50. The measurement of 100Hz below -60dB is very good (especially when taking into account this is a passive power supply). And, the noise injected from the load is 333Hz at -90dB. At this level, it may even be lower than when using some "standard" active power supplies. There are no worries whatsoever with the measurement 1kHz at -107dB. We are getting in the territory of the "better" active regulators while we still are in a (now almost) non radiating/transmitting frequency range.

From 10kHz at -130dB on, it is beyond the reach of most, if not all, linear active power supplies. The passive supply, at the least in simulation, shows the "superpower" to go down in noise level, for ever and ever. At 100kHz, it is -156dB and beyond that … who knows maybe -200dB or more will show up. At these higher frequencies, it will be the build-quality that sets what happens there, later we will investigate more about this.

While writing this, I decided to leave it to the numbers, graphs, and short descriptions and let the readers judge for themselves. The next article will detail the benefits of my proposed PSU.

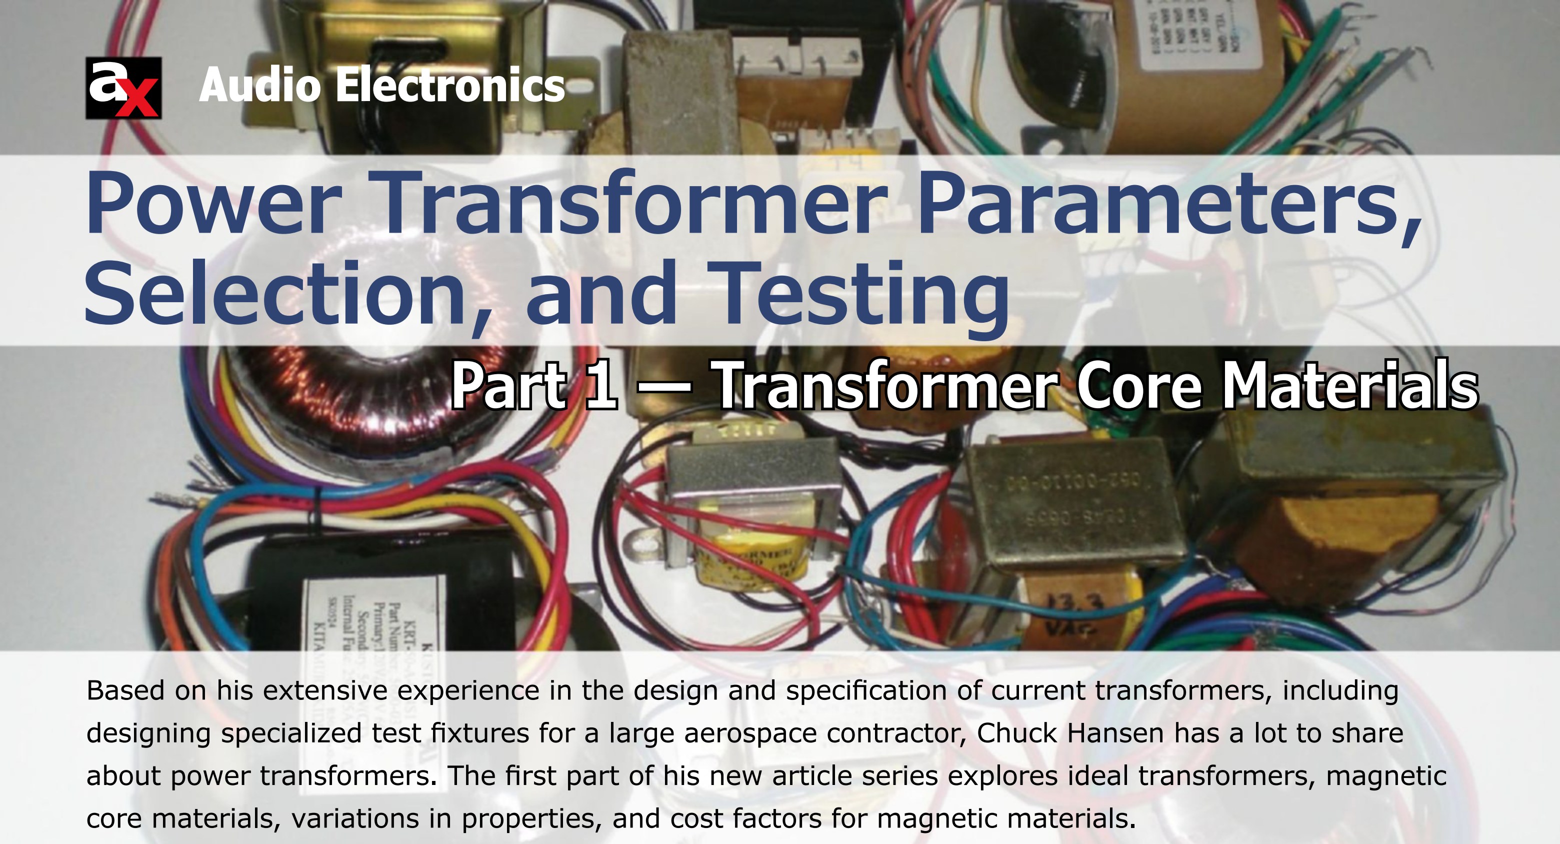

Power Transformer Parameters, Selection, and Testing

Part 1 — Transformer Core Materials

Based on his extensive experience in the design and specification of current transformers, including designing specialized test fixtures for a large aerospace contractor, Chuck Hansen has a lot to share about power transformers. The first part of his new article series explores ideal transformers, magnetic core materials, variations in properties, and cost factors for magnetic materials.

By
Chuck Hansen

Preface

Over the years, several of my esteemed peers including Pete Millet, Gerhard Haas, Ed Simon, Menno van der Veen, and Frans de Wit have written several excellent articles for audioXpress *on topics pertaining to transformers (see Resources). I recommend reading these excellent articles as a preface to this article on power transformers, where I discuss standard 50Hz/60Hz transformers you can buy online, as well electric power utility and aerospace transformer designs.*

I work in the aerospace sector where we design AC and DC electric power systems. In addition to analog circuits, I design current transformers and low-to-medium power step-down and isolation transformers in accordance with the pertinent MIL-Specs. I also specify MIL-Spec and commercial transformer purchases for our Lab test equipment, and design specialized test and calibration fixtures for the Lab.

There are a number of MIL-Specs referenced in this article, and all of them are available free for your reading pleasure on http://everyspec.com.

Power Transformers

As the very experienced engineer who mentored me in my first job used to jokingly say, "Transformers are just bell wire wrapped around a pair of horseshoe nails." Well, we all know there is a bit more to it than that.

A transformer is a static device used for transforming electric power from one circuit to another without changing the AC frequency, based on the principle of mutual induction. Current is applied to one winding wrapped around a high permeability "soft" magnetic material, which causes a voltage to appear at the other winding without any electrical connection between the wires. Transformers can transform the input voltage to the output, either higher (step-up) or lower (step-down).

Commercial chassis-mount power transformers (also mains transformers in Europe) used by audio and electronics hobbyists are designed and manufactured to standards published by the Underwriters Laboratories Inc. (UL, cUL), the International Electrotechnical Commission (IEC), the American National Standards Institute (ANSI), the Institute of Electrical and Electronic Engineers (IEEE), the Canadian Standards Association International (CSA), the Conformitè Europëenne (CE) and the Verband Deutscher Elektrotechniker (VDE) certification requirements.

Many also meet the Restriction of Hazardous Substances (RoHS), which restricts the use of lead in solder, and also applies to metal electroplating, anodizing, chromating, and other finishes on electrical and electronic equipment (EEE) components, heatsinks, and connectors.

Power transformers that are designed to connect to the high-voltage electric power grid are custom designed for the voltage and power level required by the utilization entity (a business, shopping center, planned community, etc.) and then furnished by the utility. Similarly, the pole-top transformers that feed residential areas are also specified and furnished by the utilities. Military transformers are designed to MIL-PRF-27F.

The design of the kind of transformer we are interested in requires control over a large number of physical and electromagnetic variables. Some of the limiting design factors include:

- Cost
- Volt-ampere (VA) ratings, typically from 10VA to 1000VA
- Size and weight limit
- Temperature limit
- Other environmental requirements
- Voltage regulation requirement, no load to rated full load

Power Transformer Core Materials

An ideal transformer core material would confine all of the magnetic flux within the core. The flux would perfectly link all of the windings. The core losses would be zero, the efficiency would be 100%, and the permeability, the ability to support the formation of magnetic fields in a core material, would be infinite. The coercivity would be zero, so the magnetic field required to reduce the AC magnetization flux to zero, would be zero. Ideal windings would have zero resistance. Voltages would transform exactly to the turns ratio of the windings. Current would transform exactly to the inverse of the turns ratio.

Of course, there is no such thing as an ideal transformer. We have to design with all the non-ideal parameter in mind. Selection of the type of core and core material is a tradeoff of both economic and performance factors.

Back when electrical engineers designed DC starter-generators, AC transformers, and brushless AC alternators (still called generators by tradition) with slide rules instead of calculators or computers, frequency was given in cycles per second (CPS), and flux density (B) was given in kilolines per square inch. Many experienced generator designers still comfortably think in those terms. The other older units for the flux density, maxwells/sq in and kilogauss (kgauss), are becoming obsolete. Now magnetic flux density is measured in newton-meters per ampere (Nm/A), also called tesla (T). One tesla is equal to 10 kgauss. We still use kgauss, pounds for core weight, and inches for core dimensions because that is how the core catalogs list them.

The unit of the magnetic flux is the tesla meter squared ($T \times m^2$), also called the weber and symbolized Wb). The older units for the magnetic flux, the maxwell (equivalent to 10^{-8} Wb), is also seldom seen today.

Magnetic Field Strength, H, is measured in Oersteds. The older units for field strength, amperes/inch and amperes/meter, are also seldom seen today.

These newest units are in metric centimeter-gram-second (CGS), while the older units are in US customary and British imperial foot-pound-second (FPS).

Magnetic Cores

It is important to note that iron (Fe), cobalt (Co) and nickel (Ni) are the magnetic elements in all the core materials we will talk about. They are periodic table elements 26 through 28, respectively. Smaller amounts of elements such as boron (B), carbon (C), chromium (Cr), copper (Cu), manganese (Mn), molybdenum (Mo), niobium (Nb), phosphorus (P), silicon (Si), sulfur (S), tantalum (Ta) and vanadium (V) are added to optimize the mechanical and metallurgical properties of the core, such as higher resistivity, less brittleness and more strength.

Core materials used in transformers and other electromagnetic devices are categorized as soft magnetic materials, as opposed to hard magnetic materials that are used for permanent magnets. They can rapidly switch their magnetic polarization under a relatively small applied magnetic field. They are generally characterized by an intrinsic coercivity of less than 1000 A/m. Soft magnetic materials are also used in magnetic sensors and EMI shielding.

The majority of transformers, motors, generators, solenoids, mechanical meters and

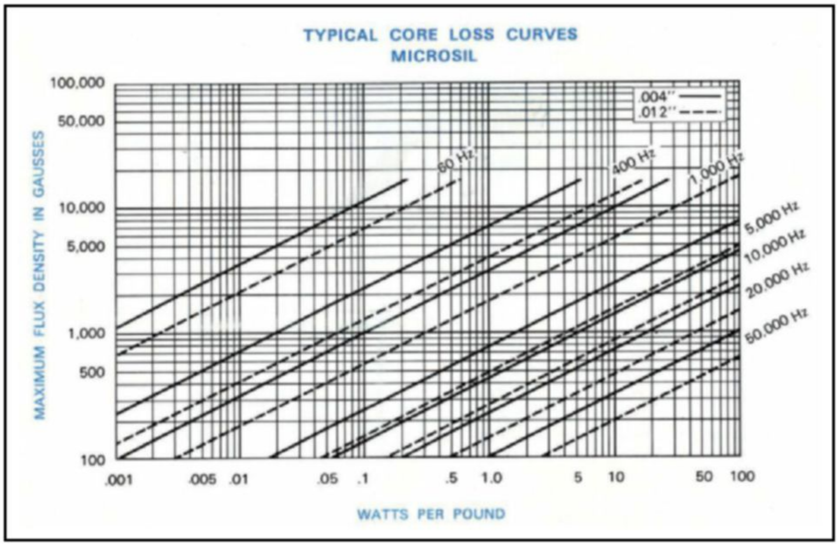

Figure 1: This graph shows core loss vs. frequency for 3% silicon steel, 4-mil and 12-mil laminations. (Image Source: Magnetic Metals, *Tape Wound Design Manual C1-171-20M*)

actuators use 3% silicon-iron by weight (SiFe) as their magnetic material, although it is called silicon-steel in the industry. In addition to the 3% Si, it can contain up to 0.5% each of Al and Mn, also identified as Fe96-Si3-Al0.5-Mn0.5.

The addition of up to 6.5% Si (by weight) yields better electrical and magnetic properties. However, above 3.2% Si, the steel becomes brittle and difficult to work with. For traditional steel manufacturing methods, 3% Si is still the most widely used soft magnetic material.

The silicon also increases the resistivity of the iron by about 5 times compared to low carbon steel. This, and laminating the core into thin strips, limits the eddy current losses (more on losses later). Silicon also lowers the remanence of the steel, which decreases its tendency to become permanently magnetized, and reduces hysteresis losses.

Commercial 50Hz/60Hz transformers use two types of SiFe, depending on the manufacturer's tradeoff between cost, size and performance. The more costly 3% SiFe cold-rolled grain-oriented (CRGO) 14 mil cores (Magnesil, Microsil, Silectron) are designated M16. The grain orientation combines low magnetostriction with enhanced permeability along the length of the laminations and allows operation at a higher flux density. The initial permeability is 1500, the maximum flux density is 15 to 18 kgauss, and the core losses are 0.65 W/lb. Other industry designations for it are grain-oriented electrical steel (GOES) and grain-oriented silicon-steel (GOSS).

Motors and generators use cold-rolled non-grain-oriented (CRNGO or just NGO) since their spinning rotors have a constantly changing flux path direction through the stator core. The permeability of CRGO steel drops by half when the grain angle is just 10° off axis. Economy commercial transformers also use CRNGO, designated type M19. It has a lower initial permeability of 400, a lower maximum flux density of 12 to 15 kGauss, and a higher core loss of 1.0W/lb.

The required lamination thickness for GOSS cores varies inversely with the ac frequency, so thinner laminations are used for higher frequencies (**Figure 1**). Lamination thickness is given in either millimeters (mm), or mils (0.001 inches).

Silicon also reduces the core magnetostriction (the changes in dimensions due to the magnetic field variations). The downside of adding silicon is that, while it reduces core loss, it also reduces permeability, requiring a higher current to produce same flux density. This also increase the copper loss in the windings.

Silicon causes the saturation flux density to drop by about 6% as compared with low carbon steel, but silicon steel also has the lowest variation in peak flux density (Bmax) of all the electrical steels, over a temperature range of 150°C to +375°C (SiFe CRGO curve in **Figure 2**).

The saturation flux density decreases as the temperature approaches the Curie point, the temperature where the permeability goes to unity and the alloy becomes non-magnetic. For CRGO 3% SiFe the Curie temperature is 750°C. For Supermendur it is 940°C. Cobalt increases the Curie temperature while nickel causes it to decrease. For MetGlas amorphous alloys it is lower still. Temporary excursions above the Curie point (which happens during annealing) do not cause any permanent degradation of Bmax.

CRGO silicon steel comes in five standard American Iron and Steel Institute (AISI)/Society of Automotive Engineers (SAE) lamination thicknesses; M-2, M-3, M-4, M-5, and M-6; from 7mil to 14mil, respectively, with much lower core losses than

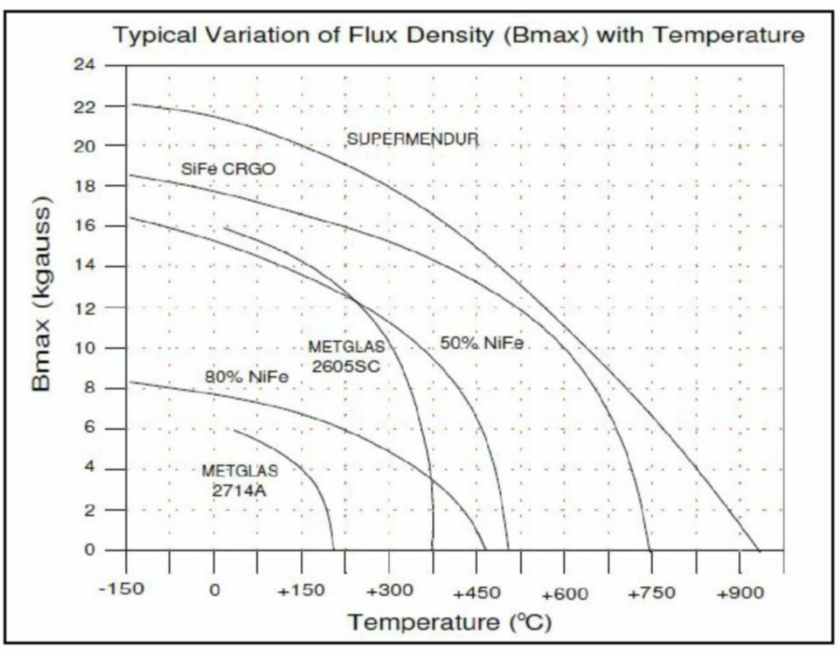

Figure 2: This graph shows the typical variation of flux density, Bm, with temperature. (Image Source: Magnetics Inc., *Design Manual for Tape Wound Cores TWC-300U 9X-1990*)

About the Author

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

CRNGO types. M-4 and M-6 are recommended for instrumentation and current transformers for commercial use. Thinner AISI gauges are available down to 1mil for use in aerospace applications at 400Hz and higher.

CRNGO silicon steel comes in eight standard AISI grades by maximum core loss; M-15, M-19, M-22, M-27, M-36, M-43, M-45, and M-47 from lowest to highest, respectively. In aerospace design we use 14-mil M19 for all our rotating machines that use CRNGO laminations.

Silicon steel also has the lowest cost per unit volume of all the soft magnetic transformer materials (**Figure 3**). Pure iron has a lower cost, but its magnetic properties are not all that useful even at the lowest 50Hz to 60Hz line frequencies.

We used to have an AC Hysteresis Test System that measured the dynamic hysteresis loop of both CRGO and CRNGO silicon steel, from 50Hz to 1kHz. It also measured the static magnetic parameters; coercivity Hc, loss angle δ, total loss Ps, amplitude permeability μa, and specific remanence Br. This was important because we were using 17.6kG-type AZ silicon steel for our large power transformers, rather than the usual 15kG AH steel we used for general purpose transformers. While it cost more, it allowed for the lighter weight magnetics that high-performance aerospace applications require for low size, weight, and power, with as low a cost as possible (SWaP-c) made by qualified suppliers that can meet all the program objectives and specifications.

Silicon-Steel Manufacturing Process

The manufacturing process for silicon-steel starts with cast molten SiFe slabs, which are then hot-rolled to about 2mm (~80mils). Next the rolled slabs are pickled with sulfuric acid (H_2SO_4) to remove impurities, inorganic contaminants, rust or scale; then cold-rolled to 0.75mm (~20mils). This is followed by annealing at 750°C to 900°C, which increases the ductility and aligns the grain orientation in a preferred direction.

After cooling, the slabs are cold-rolled to the final gauges, 0.55mm to 0.04mm (14mils to 1mil), used for generators, motors, transformers, inductors and solenoids. This cannot be accomplished to the required precision by hot-rolling. This is followed by an 830°C to 900°C annealing that decarburizes the surface and re-crystallizes the grain orientation. Then one last annealing is performed at 850°C to 1,110°C.

Finally, the steel sheets are electrically insulated on both sides with a chemical treatment (e.g., CARLITE ASTM Type C-5 or C-10). For large utility power transformers with high voltage per turn, CARLITE C-3 insulation is recommended. The finished steel is wound into large coils for shipment to customers.

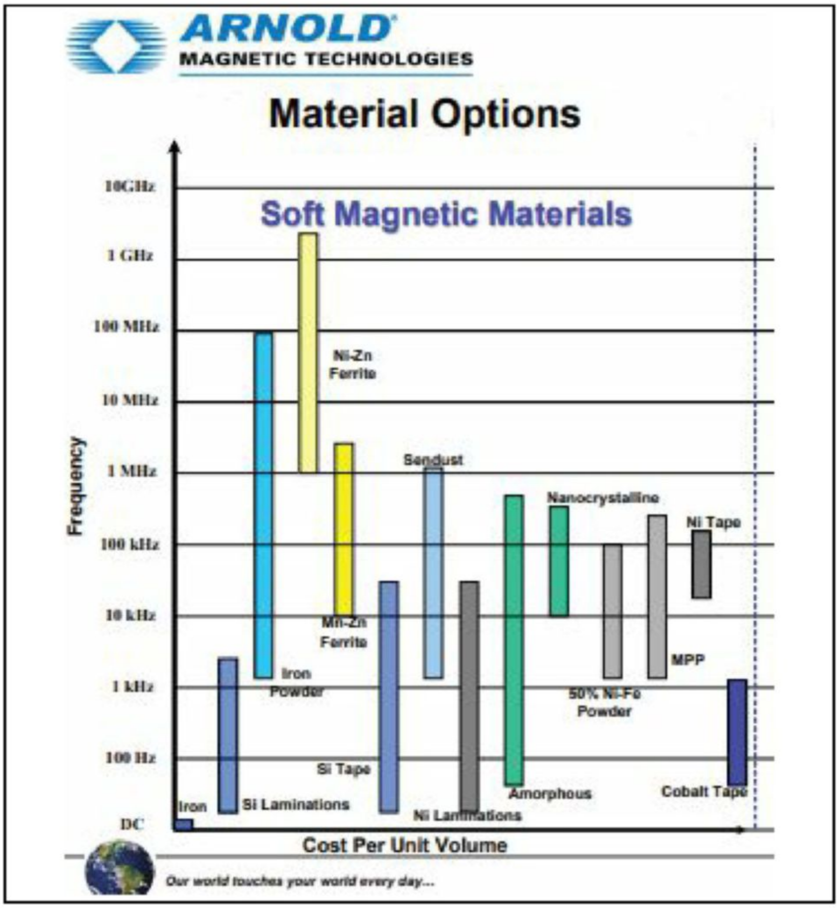

Figure 3: Here is a comparison of soft magnetic materials as a function of cost and usable frequency range. (Image Source: Arnold Magnetic Technologies).

Next month, we focus on transformer cores and their construction methods, from the first soft-iron core transformer design in 1878 to the 1980's development of MetGlas amorphous metal and nanocrystalline transformer core alloys.

Editor's Note: All audioXpress *articles from 2001 to present can be found on the aX Cache, a USB drive available from www.cc-webshop.com.*

Resources

F. de Wit, "Thinking about DC Power Supplies," *audioXpress*, December 2021 (pp. 46-51), and January 2022 (pp. 44-51).

G. Haas, "Transformers for Tube Amplifiers," *audioXpress*, May 2018 (pp. 32-40), June 2018 (pp. 54-59), and July 2018 (pp. 48-58).

P. Millet, "Power Transformers for Audio Equipment," *audioXpress*, June 2001 (pp. 14-19), republished on the audioXpress website on May 5, 2015.

E. Simon, "Powering Your Circuits," *audioXpress*, February 2020 through May 2020.

M. van der Veen, "Vanderveen Trans-SE10 Secrets," *audioXpress*, January 2021 (pp. 58-67).

Voltage Regulator Tubes in Audio Amplifiers

By
Richard Honeycutt

This article details the history of voltage regular tubes and discusses some of the merits of their use in audio amplifiers. Not common today, they certainly added extra visual excitement, but there was a reason why they were used.

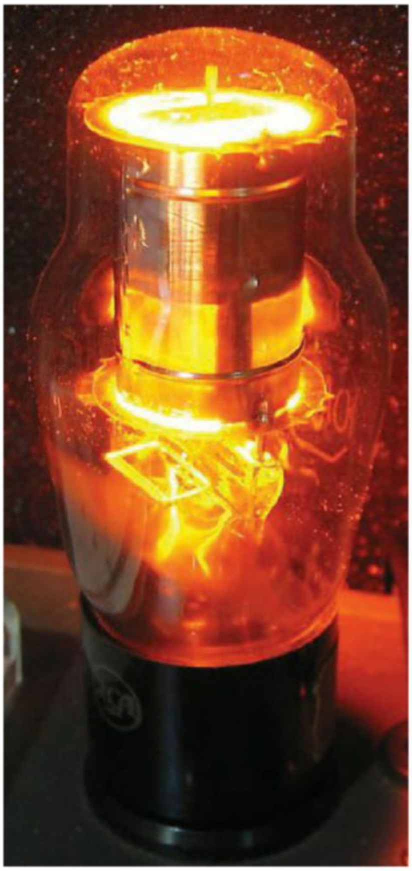

Photo 1: In operation a voltage regulator tube can be visually exciting.

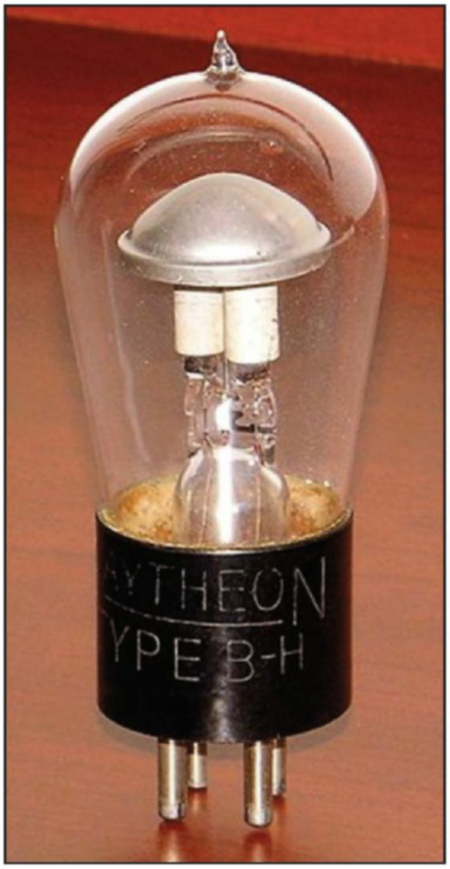

Photo 2: The 1920's-vintage "Raytheon Tube" was the the first VR ube manufactured.

Many hollow-state equipment enthusiasts have mentioned enjoying the "warm glow" of tube equipment, as opposed to the "cold" appearance of solid-state gear. Perhaps the most visually exciting vacuum tube is the voltage regulator (VR) tube (**Photo 1**). These are not really vacuum tubes, as they are filled with a gas such as argon, or a mixture of gases. The "cold cathode" is not heated.

The VR tube was invented by Lawrence K. Marshall, Vannevar Bush, and Charles G. Smith, and was used in a battery eliminator that could be plugged into an AC mains receptacle to allow battery-operated radios to be used without expensive batteries of limited lifetimes (**Photo 2**). As you can see from the battery eliminator circuit shown in **Figure 1**, this tube would probably be more accurately called a gas rectifier than a VR tube, although its operation did depend upon ionized helium gas. As you can deduce from the schematic, significant filtering was needed because of the inherently high level of RF noise produce by the ionization. These battery eliminators also had to be well shielded. Dubbed the "Raytheon tube" (light from the gods), it was manufactured by the American Appliance Co., whose name would later be changed to the Raytheon Co. in 1959.

The devices mainly used as VR tubes today are half-wave diodes rather than full-wave ones such as the Raytheon Tube. They function similarly to a Zener diode, but there are four differences:

- A VR tube's operation depends upon gas ionization rather than primary breakdown of a PN semiconductor junction.
- In order for the VR tube to "strike" (begin ionization), the unregulated supply voltage must be 15-20% higher than the tube's nominal output voltage.
- When conducting, the VR tube exhibits a negative resistance: increased anode current results in a decrease in anode-to-cathode voltage.
- According to the RCA datasheet for the OA3, the VR tube can exhibit a sort of memory effect: For example, the regulation of a tube operated for a protracted period at 5mA and then changed to 35mA may be somewhat different from the value obtained after a long period of operation at 35mA.

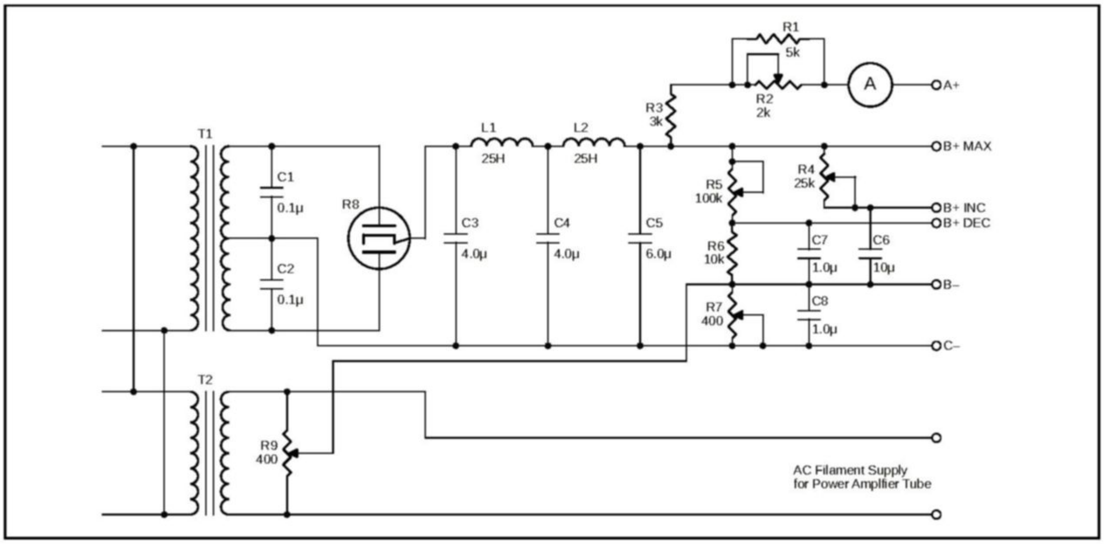

Figure 1: The Raytheon tube battery eliminator provided three different (adjustable) B+ voltages, a low-current "A" filament voltage, and a "C-" control-grid bias voltage. Maximum current output was typically 80mA to 90mA.

In use, a VR tube is connected in series with a resistor, creating a shunt regulator. If the supply voltage is higher than the breakdown voltage of the gas, the gas will ionize and the tube will conduct exactly enough current to bring its plate-to-cathode voltage (=supply voltage minus drop across series resistor) to precisely the breakdown voltage of the VR tube. In this way, the voltage across the VR tube is held to a constant value.

Figure 2 shows a shunt regulator using a VR tube. The current in the 3.7kΩ series resistor is the sum of the output current and the VR tube current. Thus when the regulator's load current varies, the VR tube's current varies in an inverse fashion. Due to the small negative resistance of the VR tube, a lower load current will result in a lower output voltage. This is opposite the effect observed with a Zener diode, which exhibits a small positive resistance.

Although I have not been able to find the approximate value of negative resistance of any of the common VR tubes, one would expect that the negative resistance in the regulators shunt arm would help compensate for the positive internal resistance of the power transformer, which is effectively in the series arm. The most common VR tubes are the octal-based ones listed in **Table 1**. They can regulate from 5mA to 40mA.

VR Tubes in Audio Amplifiers

Very few audio amplifiers use regulated power supplies, since the regulation provided by a well-filtered full-wave power supply is generally considered adequate; the Class-A preamp, phase splitter, and driver circuits draw a constant current, so there is no current variation to cause B+ voltage to vary; the AC power mains voltage is normally pretty well regulated; and the voltage-dropping and decoupling capacitors between the high B+ and the lower preamp B+ generally do a good job of maintaining a pretty constant B+ for the early amp stages.

However, some classic amplifiers do incorporate VR tubes. The Altec a333a schematic shown in **Figure 3** used a VR tube to reduce the B+ voltage for the output tubes' screen grids. The B+ to the center tap of the output transformer is 410V, resulting in the 6L6G output tubes operating at a plate voltage of 395V. The OA3 reduces this to 320V for the screen grids.

Figure 4 shows an LTSpice model for an Altec a333a amplifier. Since I could not find an LTSpice model for an OA3 or an OC2, I simulated the VR tube by using two Zener diodes—a 68V one and a 6.8V one—in series to approximate the 75V breakdown voltage of the VR tube. (I could not find an LTSpice model for a 75V Zener.) With the 10mA current level, the small difference between the VR tube's negative resistance and the Zeners' positive resistances have little effect in the simulation. The a33a used a 6SJ7

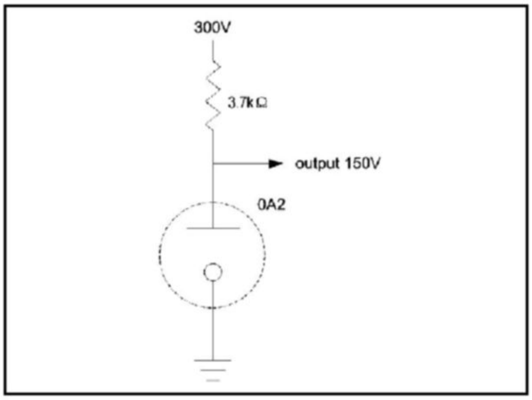

Figure 2: The OA2 serves as the shunt element in this shunt voltage regulator.

Tube Type	Voltage
0A3	75V
0B3	90V
0C3	105V
0D3	150V

Table 1: These are four common octal VR tubes.

About the Author

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

preamp tube for the preamp. Not having an LTSpice model for the 6SJ7, I used one triode section of the similar 12AU7 instead. The 855Ω resistor R20 represents the equivalent internal resistances of the 5U4G rectifier and the power transformer.

Photo 3 shows the output wave from the simulation, at a power of 23.26W. The test signal had a frequency of 50Hz, in order to show up any possible differences due to the use of a VR tube to set the screen-grid bias of the output PP pair. **Photo 4** shows the fast Fourier transform (FFT) for that output wave; the total harmonic distortion

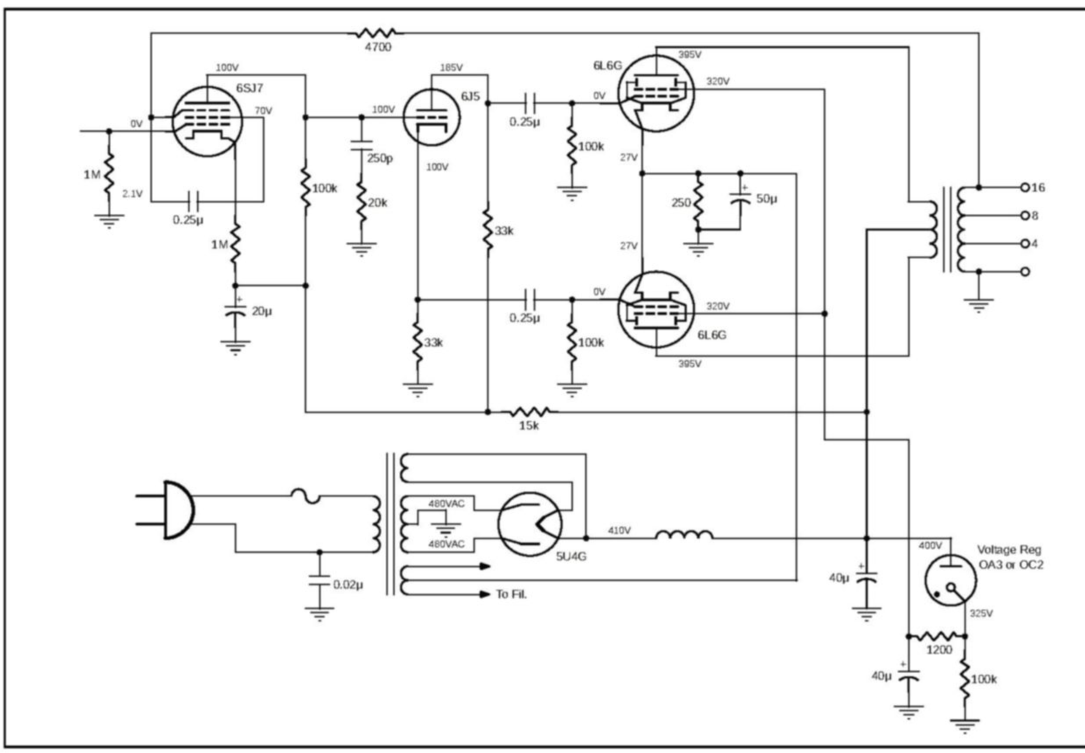

Figure 3: The Altec a333a power amp used an OA3 to drop the screen grid voltage of the output tubes.

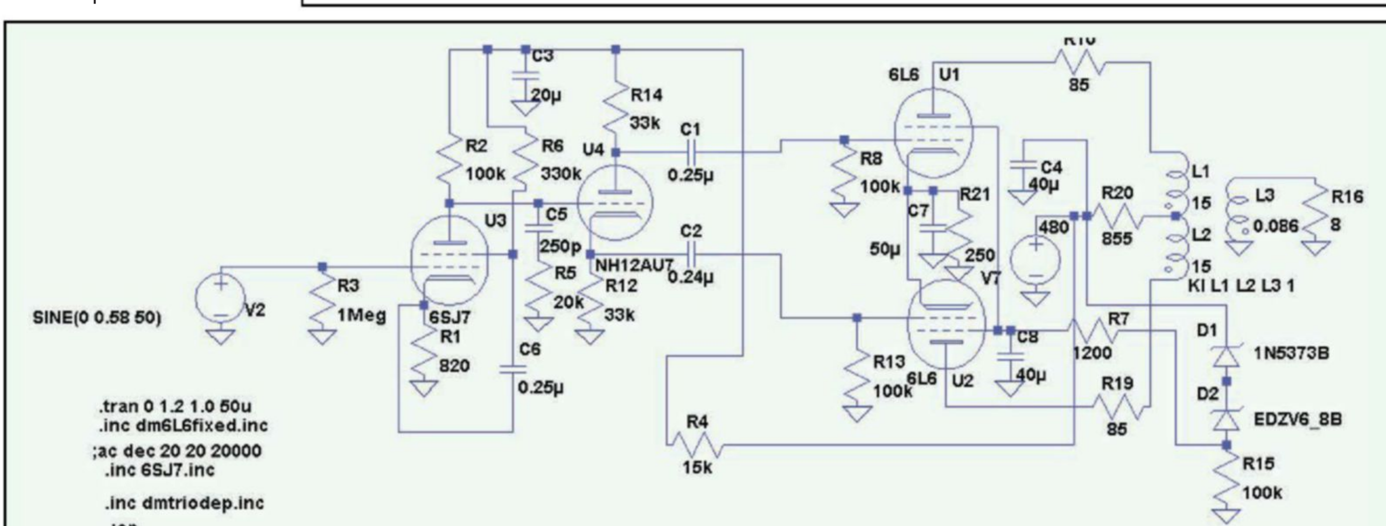

Figure 4: This is the LTSpice model of the Altec a33a amp used for analysis.

(THD) is a very low 1.14%, well under the level of 2% Altec claimed for that amp at 20W.

Photo 5 and **Photo 6**, respectively, show the screen-grid signal waveform and FFT. Although the waveform looks like a fairly clean 50Hz sine wave, the FFT shows that the 100Hz second harmonic predominates. Both even- and odd-order harmonics are present, with even-order ones being dominant. But the peak-to-peak (P-P) signal voltage on the screen grid is only 0.6V.

The manufacturer's circuit uses the VR tube to set the screen grid bias at 75V below the plate voltage. It is not intended to stabilize the screen grid bias at a particular voltage. The more common approach to setting screen grid bias below the output tubes B+ is to use a dropping resistor. In either case, decoupling capacitor C8 is used to place the screen grid supply voltage near AC ground. (If it were truly at AC ground, no signal voltage would appear on the screens.)

I modified the LTSpice model to use a 7102Ω resistor instead of the VR tube (Zener diodes in the previous model). Frequency responses for the Zener model and the resistor model are virtually identical. **Photo 7** and **Photo 8** show the output waveform and FFT, respectively, of the resistor model. With the same input voltage, the output power is about 0.7 dB lower than with the Zener diode model, and the distortion is the same at 1.14%.

Due to there being no change in THD compared to the VR circuit, I am not including the screen grid waveform and FFT. It is not obvious why the extra expense of a VR tube and socket instead of a dropping resistor would be used for this small improvement. Increasing the input signal voltage to give the advertised 27W output power revealed no new perspective.

Let's consider other classic hollow-state power amplifiers that used VR tubes in the same way. The Altec a340a amp (**Figure 5**) is a somewhat higher-power unit using push-pull (PP) 6550s to produce 40W at less than 0.5% THD, and it uses an OA3 in a way very similar to the Altec a333a.

The power amp used in the Leslie models 122 and 122R speakers for Hammond organs used an OC3 and two dropping resistors to drop the high B+ from 420V to 260V to supply the plates of the 12AU7 balanced driver amplifier. The PP 6550 tubes put out a bit over 40W maximum.

When these classic amplifiers were designed, standard AC line voltage in the US was defined as 95V to 115V— nominally 115V with a +10% tolerance—rather than 125V as it is today. In order to see whether the VR tube was used simply to

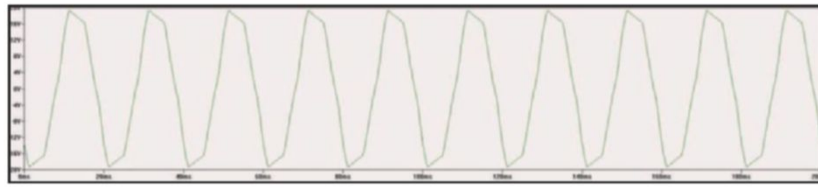

Photo 3: The output wave at 23.28W shows low distortion.

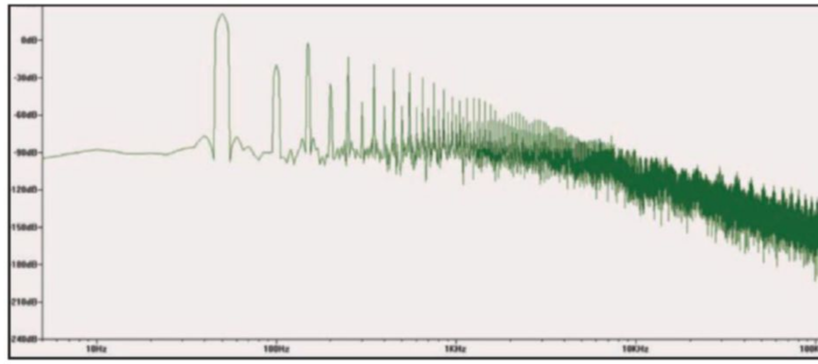

Photo 4: Due to the use of the PP configuration, the FFT shows mainly odd harmonics.

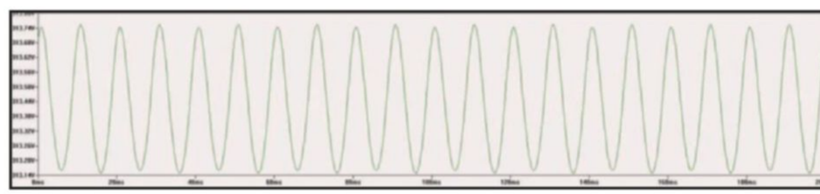

Photo 5: The signal on the screen grids is a steady 313.3V + 0.3V.

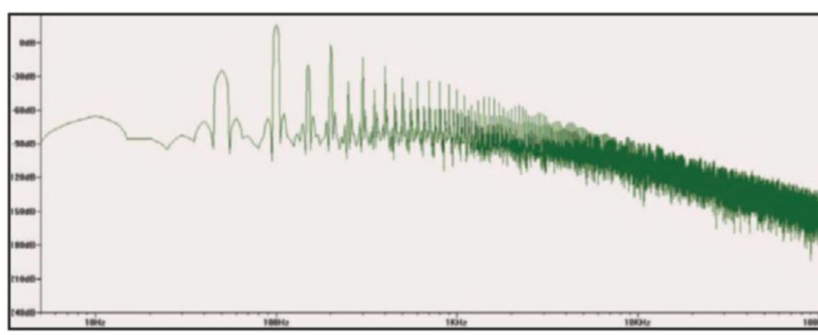

Photo 6: The FFT of the screen signal shows little 50-Hz fundamental.

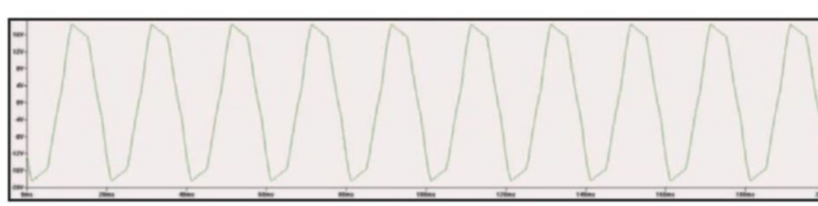

Photo 7: The output wave of the resistor model shows the same THD as the Zener model did, and with the same input signal level, the output power is about 0.7dB lower at 19.43W.

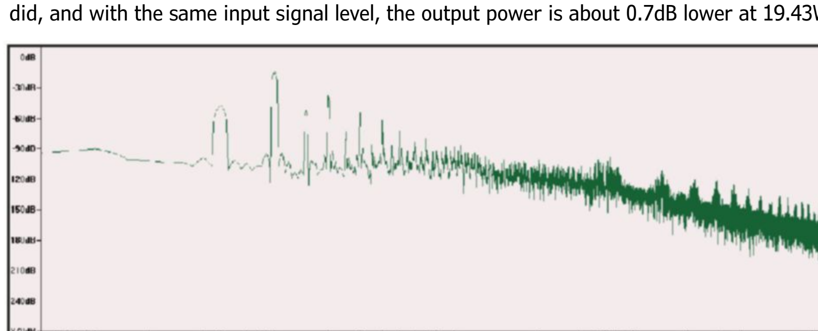

Photo 8: The FFT of the output wave shows a very slightly higher THD for the resistor version.

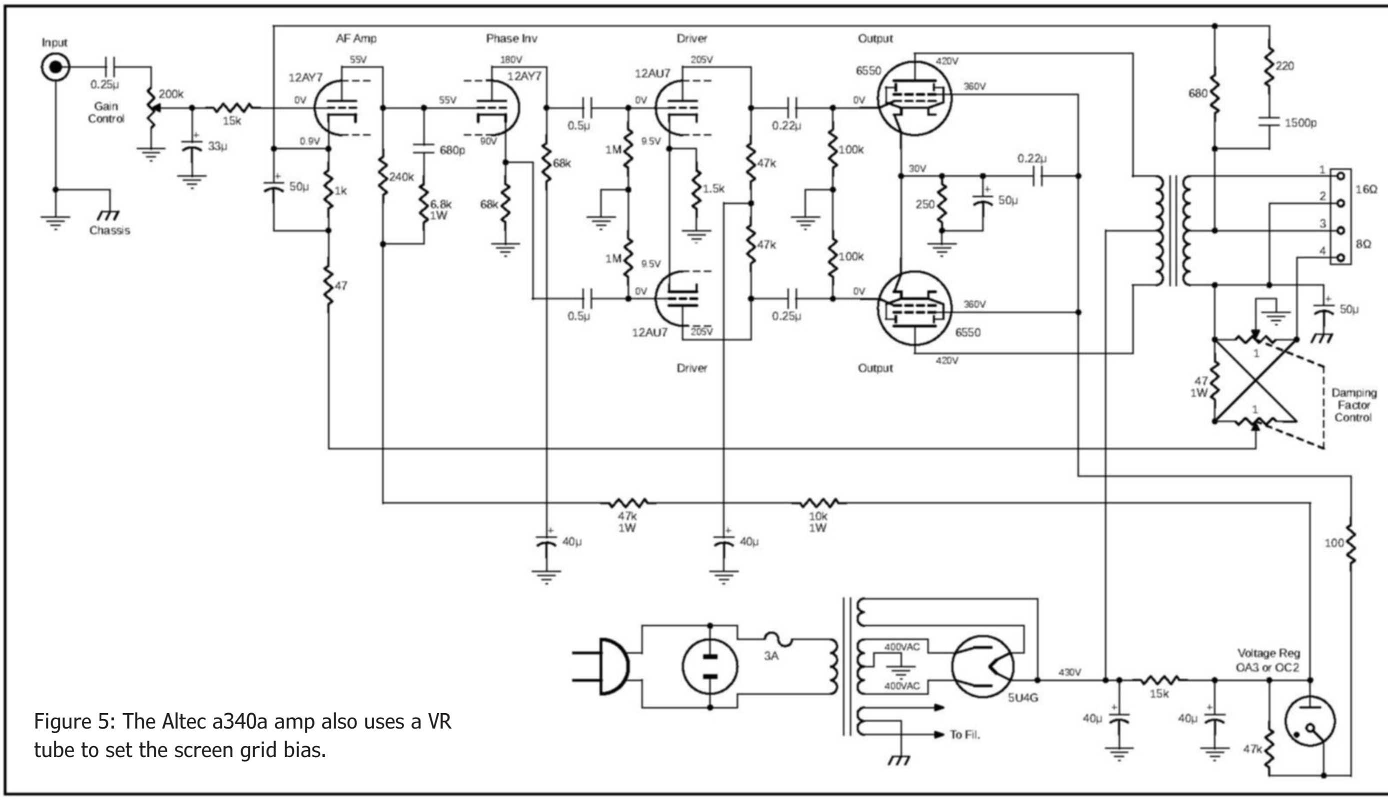

Figure 5: The Altec a340a amp also uses a VR tube to set the screen grid bias.

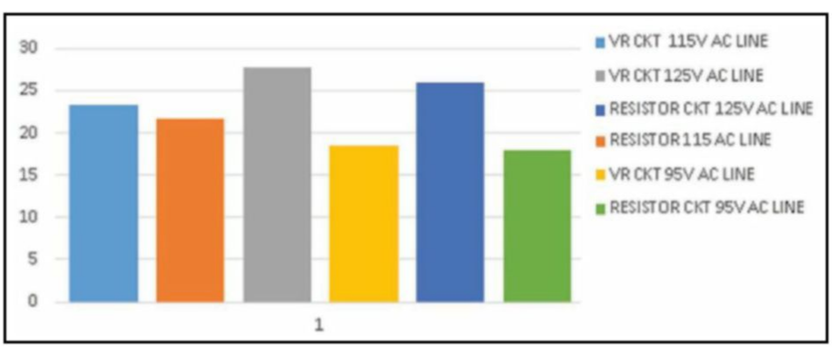

Photo 9: The VR circuit reduces variation in maximum power output caused by line voltage variations, compared to the circuit using a dropping resistor.

stabilize the screen grid bias, I re-ran the simulation with the high-voltage supply (V7 in the model) set at 10% lower and then at 10% higher than the nominal value. The results are shown in **Photo 9**.

The use of the VR tube results in a higher maximum power output, 27.78W at a THD of 1.09%, at the 125V line voltage, and 18.44W at 1.02% with a 95V line. The resistor circuit's output power and THD at 125V and 95V, respectively, were 25.95W/1.16% and 18.02W/1.09%. Although the small increase in output power and decrease in THD under low-line and high-line conditions are still not impressive, they do provide some advantage in "specsmanship." And the glowing VR tube really looks cool, too!

Corrections for *audioXpress* October 2022

In the October 2022 issue of *audioXpress*, one of the equations (located at the bottom of the first column on page 44) in the article "Triode Common Cathode Stages, the Oracle Equation, and Thodosian Objections," by Christopher Paul was incorrect.

The correct equation is shown here:

$z_{ka} = (2\times R)||[2\times r_a/(\mu+2)] = 2\times[R||\{r_a/(\mu+2)\}]$

Also in the October 2022 issue of *audioXpress*, Gary Galo's article "Sumiko Phono Cartridges and RS78 78 rpm Stylus" incorrectly listed the price of the the Sumiko Rainier cartridge, which is actually $149, and the Rainier replacement stylus, which is really $89.

Our apologies for the errors.

ELECTRONIC TRANSFORMERS

- Original Equipment
- Replacement
- Amateur

TRIAD MAGNETICS, INCORPORATED
460 HARLEY KNOX BLVD., PERRIS, CA 92571
www.TriadMagnetics.com