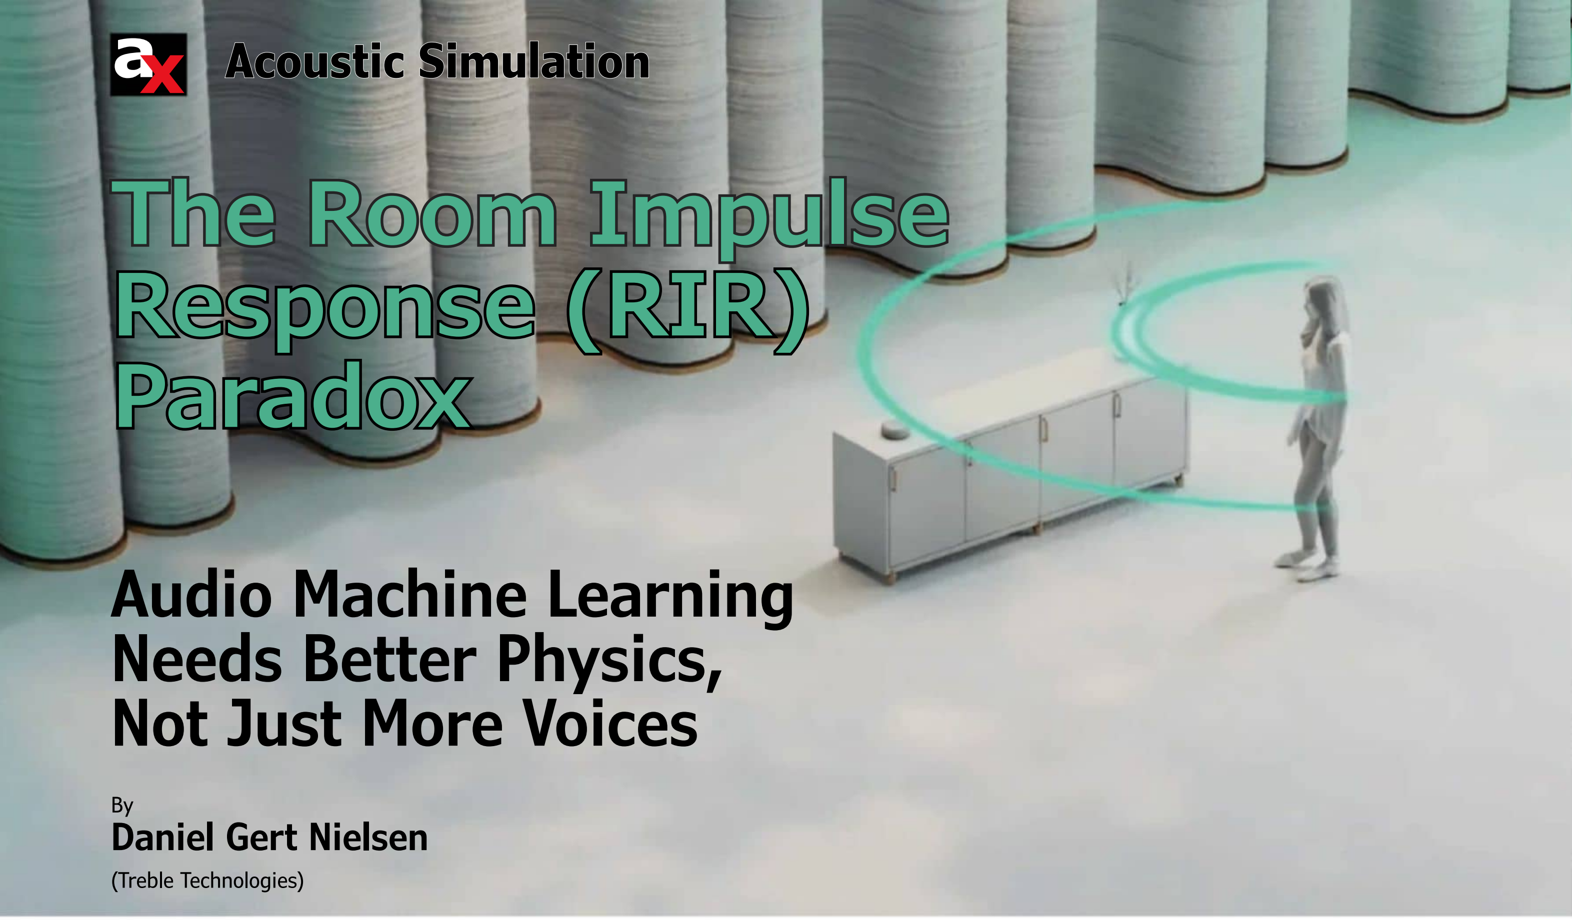

The Room Impulse Response (RIR) Paradox

Audio Machine Learning Needs Better Physics, Not Just More Voices

By
Daniel Gert Nielsen
(Treble Technologies)

The quest for superior audio machine learning (ML) has rightly emphasized the diversity of source signals such as speech. However, a critical bottleneck often remains unaddressed: the acoustic realism of the environments in which these signals are placed. This piece explores the disconnect between sophisticated source generation and often simplistic environmental simulation, arguing for a renewed focus on physics-based Room Impulse Responses (RIR) and highlighting how new simulation capabilities, such as those offered by platforms like the Treble SDK, are paving the way to truly advance the field.

The world of audio machine learning (ML) is buzzing. Advancements in deep learning have unlocked incredible capabilities, from hyper-realistic speech synthesis to robust noise cancellation and accurate sound event detection. A key driver behind this progress? Data. Lots and lots of diverse data. Recognizing this, the community has heavily invested in data augmentation techniques. One star player in this arena is text-to-speech (TTS). Sophisticated TTS models can generate vast quantities of speech data with diverse speaker characteristics, accents, emotional tones, and speaking styles. This is invaluable for training models that need to generalize across the rich tapestry of human voices. Convolutional approaches then combine this speech with noise sources to create training examples.

We see countless papers, articles, and discussions focusing on improving TTS diversity, naturalness, and control. And rightly so—the source signal matters.

But here's the paradox, a conundrum we at Treble Technologies have pondered deeply: While we obsess over perfecting the source (the voice), there's a comparative silence surrounding the medium—the acoustic environment through which that voice travels before reaching a microphone.

In the real world, sound doesn't exist in a vacuum. It reflects off walls, diffracts around corners, gets absorbed by furniture, and reverberates through space. This complex interaction shapes the sound profoundly, encoding information about the room's size, geometry, and materials. This transformation is captured by the Room Impulse Response (RIR).

For training robust audio ML models—particularly for tasks such as speech enhancement, dereverberation, source localization, source separation, and Direction of Arrival (DOA) estimation—the RIR is not just an ingredient; it's arguably the critical ingredient that determines the realism and effectiveness of the training data.

Why the Disconnect?

Generating realistic RIRs is fundamentally a physics problem. It requires accurately simulating wave phenomena, an area where shortcuts can lead to significant divergence from real-world behavior. **Figure 1** illustrates a complex acoustic environment, characterized by multiple interconnected rooms and various sound-influencing objects. Simulating such a scenario, mirroring common house and apartment configurations, with shortcuts or low-fidelity approaches would result in significantly compromised accuracy.

Consider, for instance, reflection and scattering. Sound waves striking a surface don't simply bounce off like a billiard ball. The nature of that reflection is dictated by the surface's material, its texture, and its geometry. A smooth, hard wall might produce a specular, mirror-like reflection, while a rough, textured surface, such as a heavy curtain or a purpose-built acoustic diffuser, will scatter the sound energy in multiple directions. This scattering is crucial for creating a diffuse sound field, which is characteristic of many real rooms and impacts the perceived naturalness and spaciousness of an environment. Failing to model this complexity can result in RIRs that sound sterile or overly simplistic, lacking the rich textural detail of actual spaces.

Then there's diffraction, the often-underestimated phenomenon describing how sound waves bend around obstacles and through openings. Think about hearing someone talking from an adjacent room even when you can't see them—that's diffraction at work.

In a typical room, sound diffracts around furniture, people, and architectural features. For audio ML, especially for microphone arrays, this is incredibly important. The way sound diffracts around the edges of a device casing, or around a person standing near a microphone, will alter the sound reaching each microphone element differently. Neglecting diffraction leads to an incomplete acoustic picture, particularly in scenarios where direct line-of-sight is obstructed.

The absorption of sound energy by materials is another critical factor, and it's highly frequency-dependent. Thick carpets and soft furnishings tend to absorb high-frequency sounds more effectively than low-frequency sounds, while thin panels might resonate and absorb energy in specific mid-frequency bands. This differential absorption shapes the timbre of the reverberant sound and the overall frequency balance of the RIR. An accurate RIR must capture how different materials in a room "color" the sound by selectively absorbing energy across the audible spectrum. Without this, simulated environments can sound boomy, tinny, or otherwise unnaturally balanced compared to their real-world counterparts.

Finally, all these interactions—reflections, scattering, diffraction, and absorption—combine to create reverberation. This is the dense tail of decaying sound that persists after the direct sound has ceased, giving a room its characteristic acoustic "signature." The length of this reverberation (often quantified as RT60) and its spectral characteristics are vital perceptual cues. A well-modeled reverberant tail lends realism and immersion, while a poorly modeled one can make an environment sound artificial or fail to represent specific acoustic conditions accurately.

Many current data augmentation pipelines rely on simplified RIR generation methods (e.g., basic image-source models) or use limited datasets of measured RIRs. While useful, these approaches often fail to capture the full complexity and variability of real-world acoustics. They might model simple shoebox rooms well but struggle with complex geometries, frequency-dependent effects, or the subtle but crucial impact of diffraction.

The Added Complexity: Multichannel Realities and Device-Specific Responses

The challenge intensifies significantly when we consider modern audio devices equipped with microphone arrays (**Figure 2**). For tasks such as Direction of Arrival (DOA) estimation, beamforming, source localization, or multichannel speech enhancement, we don't just need one RIR; we need a distinct RIR for each microphone in the array relative to the sound source.

This isn't simply about simulating multiple receiver points in abstract space. A physically

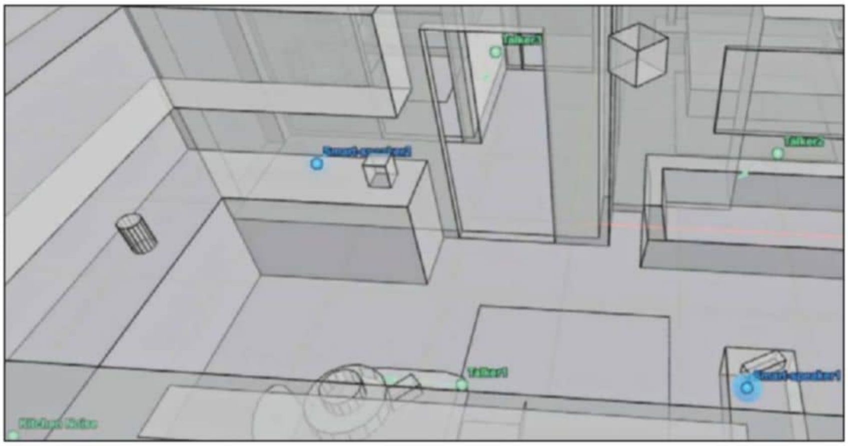

Figure 1: Complex acoustic scene with three talkers, noise sources, and two smart speakers modeled in the Treble SDK. Realistic environments enable high-quality data generation for both training and validation purposes of products and prototypes.

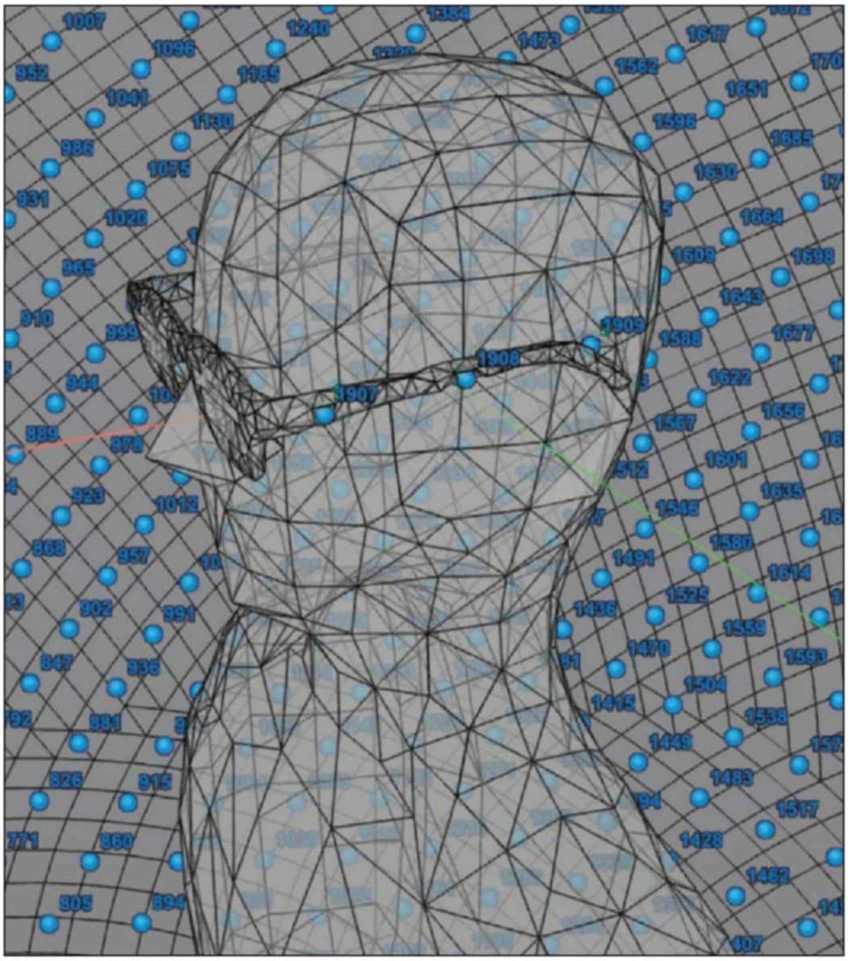

Figure 2: An example of a microphone array prototype on augmented reality (AR) glasses. Simulating devices enable users of the Treble SDK to get realistic speech signals for the individual microphone positions in real-life scenarios.

About the Author

Daniel Gert Nielsen is Senior Product Manager at Treble, a startup from Iceland, enabling acoustic engineers, designers, and developers to model sound with accuracy and efficiency within virtual environments. With a PhD in numerical modeling of vibroacoustics and loudspeaker optimization, Daniel Gert Nielsen has spent his career at the intersection of simulation and hardware. He was an early advocate for simulation-driven prototyping in AV device development and now leads the Treble SDK, a powerful toolkit for improving ML training and validation, DSP performance, and audio hardware prototyping. His focus is on utilizing advanced simulation technology to make practical tools that help engineers innovate faster and build better products.

Resources

More information on the Treble SDK can be found at:
https://treble.tech/software-development-kit

Validation studies are available at https://docs.treble.tech/validations

Discuss your project with Treble technologies at contact@treble.tech

accurate simulation, such as those we champion at Treble, must meticulously account for precise inter-microphone relationships. These relationships manifest as minute differences in the time it takes for sound to arrive at each microphone (Interaural Time Differences, or ITD, in the context of two microphones) and differences in the sound pressure level at each microphone (Interaural Level Differences, or ILD). These ITD and ILD cues are fundamental for human spatial hearing and are precisely what sophisticated DOA algorithms and beamformers are designed to exploit. If these subtle differences, often in the microsecond and sub-decibel range, are not captured with high fidelity in the training RIRs, the resulting ML models will be learning from flawed spatial information, significantly hampering their real-world performance. For instance, DOA estimation relies heavily on the accurate representation of phase differences across the array, which are directly determined by the time of arrival at each element. Even small errors in simulating these time-domain characteristics can lead to large errors in predicted arrival angles.

Furthermore, the influence of the device geometry is a critical, yet often overlooked, aspect. The physical casing and structure of a smart speaker, a laptop, or a headset act as acoustic obstacles. Sound waves scatter off the device surface and are shadowed by its bulk. This means the sound field is perturbed by the device itself before it even reaches the microphone diaphragms. Each microphone on the array, depending on its position on the device geometry, will experience a slightly different version of this scattered and shadowed field.

Simulating the RIR to the specific microphone positions on the actual device geometry is therefore crucial for realism. At Treble, we've seen firsthand how incorporating the precise CAD model of a device into the acoustic simulation dramatically improves the accuracy of the generated multichannel RIRs. Simplified RIR generation methods often completely neglect the device's own acoustic influence or use generic array configurations that don't match the target hardware, leading to training data that doesn't reflect how a real microphone array would perceive sound.

This leads to a critical bottleneck: We're generating increasingly diverse and sophisticated voice signals, only to convolve them with simplistic or limited RIRs (single-channel or inaccurately modeled multichannel ones). The result? Training data that lacks true acoustic diversity and realism, especially for spatially aware applications. Models trained on this data may perform well in simulated tests but often falter when deployed in the unpredictable acoustic environments of the real world. They become brittle, failing to generalize to spaces, source locations, and device orientations not adequately represented in their training diet.

Shifting Focus: The Power of Physics-Based Simulation—The Treble Approach

This is where our work at Treble Technologies comes into play. Our journey began with a deep conviction that advancing state-of-the-art audio ML requires a renewed focus on the physics of sound propagation, including the complexities of multichannel capture with device interaction. We saw a gap between the sophisticated algorithms being developed and the often-rudimentary acoustic data they were being fed.

The Treble SDK is the culmination of this vision, providing a powerful, Python-based environment built upon cutting-edge acoustic simulation engines. These engines are not generic wave solvers; they are specifically architected to tackle the intricacies of room acoustics and device interactions. We employ hybrid methods, often combining the strengths of geometrical acoustics for high frequencies and computationally efficient handling of large spaces, with full-wave Finite Element numerical methods when wave phenomena such as low-frequency diffraction around complex objects, including the device itself, become dominant and require utmost precision. This hybrid approach allows us to achieve a balance of accuracy and computational feasibility that is essential for generating large-scale datasets.

This empowers engineers and researchers to move far beyond simple rectangular room models. With Treble, they can simulate acoustically complex environments, importing intricate geometries (e.g., from CAD files) that include furniture, varied wall structures, and other real-world complexities.

Our engines accurately model the crucial wave phenomena—diffraction around edges, scattering from rough surfaces, and frequency-dependent absorption by materials—that define how sound truly propagates and interacts within these spaces. This means the generated RIRs inherently contain the richness and subtlety of real-world acoustics.

Crucially for modern applications, the Treble SDK excels at generating multichannel RIRs at scale, with device interaction. Users can define custom microphone array geometries, place them on imported CAD models of their specific devices, and simulate the acoustic response at each microphone position (Figure 2). Our solvers meticulously calculate the wave propagation to these exact points, inherently capturing the device's own scattering and shadowing effects, along with the precise inter-microphone phase and level differences. This is vital for training robust DOA estimators, beamformers, and other spatially aware algorithms because the training data now reflects the acoustic reality as perceived by that specific device. Imagine training a speech enhancement model for a particular smart speaker; with Treble, the RIRs used for augmentation will account for how that speaker's unique shape alters incoming sound.

Furthermore, the platform offers fine-grained control over acoustic parameters. Researchers can systematically vary room dimensions, define custom material properties (with frequency-dependent absorption and scattering coefficients), precisely position sources and receivers (including complex microphone array configurations on devices), and arrange objects within the simulated scene. This systematic control allows for the generation of highly diverse datasets tailored to specific research questions or product development needs, ensuring that ML models are exposed to a wide and representative range of acoustic conditions. This ability to parametrically sweep through configurations is key to understanding model sensitivities and building truly robust systems.

By leveraging accurate physics simulation, we can generate RIRs—both single and multichannel—that truly reflect the richness and complexity of real acoustic spaces and device interactions (**Figure 3**). When these high-fidelity RIRs are convolved with diverse source signals (e.g., those from advanced TTS or real recordings), the resulting audio scenes provide a much more robust and realistic foundation for training ML models.

The Path Forward: A Holistic Approach

The advancements in TTS for voice diversity are fantastic and necessary. But to unlock the next level of performance and robustness in audio ML, especially for applications leveraging microphone arrays, we must adopt a more holistic view of data generation. We need to pay as much attention to simulating the acoustic journey of sound to each microphone element, considering the device itself, as we do to generating the speech and noise samples.

It's time to move beyond the RIR paradox. Let's embrace the power of accurate acoustic simulation to build datasets that capture not just the diversity of voices, but the equally important diversity of the environments they inhabit and the specific ways our devices perceive them. By improving the physics engine behind our data augmentation, we can build more robust, reliable, and effective audio ML systems for the future. At Treble Technologies, we're committed to providing the tools to make this a reality, helping the audio community bridge the gap between simulation and the complexities of the real acoustic world.

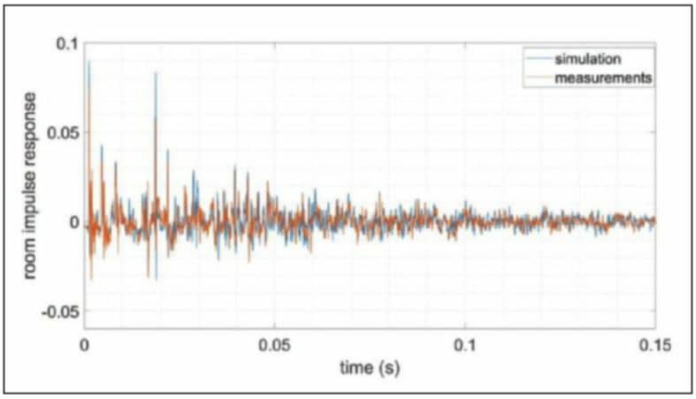

Figure 3: Impulse response from a furnished room—comparing measurements and simulation. One of the many validation studies that can be found on https://docs.treble.tech/validation.

Anechoic Chamber Design Fundamentals

This article is a detailed and experience-founded overview of some of the physical, acoustical, electrical, and other fundamental considerations that are essential in the acoustic design, planning, and implementation of anechoic chambers.

By
Neil A. Shaw
(Menlo Scientific Acoustics)

Specialized test spaces have progressed since Leo Beranek and Harvey Sleeper first published their article at a time when anechoic chambers did not yet exist. The need to measure acoustic instrumentation led Beranek and Sleeper to develop an anechoic chamber at Harvard that was free from weather conditions—at that time many measurements were done outdoors in what could be called a hemi-anechoic environment. These outdoors activities could annoy the public when measuring loud sources, and did not provide a quiet environment.

In this article, I expand on an original presentation that I gave to Audio & Loudspeaker Technologies International (ALTI) and the LA Acoustical Society of America (ASA) chapter, providing some of the things to consider from the acoustic design point of view and more in the planning for, and implementation of, walk-in chambers. Here I will also share additional details from my work with some 80Hz and 160Hz chambers, as well as random related thoughts.

Acoustic Overview

Many people have explored techniques to measure some loudspeaker performance parameters in a non-anechoic environment, but for some measurements you need a controlled environment—especially for products with low distortion and self-noise.

Anechoic chambers where the lower working frequency is 160Hz and below are undertakings that meet the criteria for working with nitroglycerine: "Know what you're doing before you do it." This is more important when people are in the chamber during tests.

From an acoustic standpoint, these are some of the criteria:

(1) The working floor area.
(2) The effective working volume to accommodate the devices under test and/or the required fixture(s).
(3) The low-frequency cutoff, which determines the effective wedge depth and the lowest frequency for "anechoicness."
(4) The overall chamber size—a function of anechoic wedge size, working volume, airspace, vibration isolation, and enclosure wall dimensions.
(5) The ambient/background noise level.

The most basic parameters are:

- The size of the largest DUT that will be used in the chamber. This lets you determine the minimum working volume where: $V_{object} \leq 0.05\ V_{chamber}$. For some chambers, a fixture will be built in the chamber (e.g., for psychoacoustic research), and this will determine the minimum working volume.
- The frequency range for tests in the chamber. This determines the low-frequency, *f*, cutoff for the working volume and the effective size

of the anechoic wedges as the size of an anechoic wedge is determined by the low-frequency cutoff. The most widely accepted approximate formula to determine the cutoff frequency of the chamber is $f = c/4l$, when c is the speed of sound and l is the length of the wedges.

The speed of sound is proportional to the absolute temperature:

$$c_{ideal} = \sqrt{\gamma \times R \times T}$$

where c_{ideal} is the speed of sound in an ideal gas (m/s2), γ is the adiabatic index ≈ 1.4 (for diatomic molecules), R is the molar gas constant ≈ 8.3J/(K·mol), and T is the absolute temperature in degrees Kelvin.

For an 80Hz chamber, at standard temperature, 293.15K (77°F) per NIST, the speed of sound is 1135.35 ft/s, the cutoff wavelength is $\lambda_{80} = c/f$ = (1135.35 ft/s)/(80 s^{-1}) ≈ 14.19 ft, which requires a wedge length l ≈ 3.55 ft (42.6 inches). So the first cut minimum size of the enclosure (MSE) is:

$$MSE = (2 \times l) + \sqrt[3]{V_{chamber}} + (2 \times \text{chamber panel thickness}) + 4"$$

where the 4" allows mounting the wedges to the chamber.

To maintain a constant and controlled environment some conditioned air circulation is required, which is also needed to deter mold. Air absorption in the chamber is a function of the relative humidity and is a good reason to have the temperature and humidity controlled via the HVAC system (**Figure 1**).

Space Planning

One thing that the real estate division may not know, or does not consider in almost all cases for audio, is the noise floor in the chamber needs to be ideally 10dB below the lowest sound levels of interest. A box-in-a-box construction is commonly used to mitigate external noise sources, and the outer box must fit in the proposed space (experience shows that the obvious is not always achievable).

To improve the noise isolation from external sources, the greater the distance between the inner wall of the outer box and the outer wall of the chamber, the better the low-frequency isolation, which increases the overall space needed. The overall height of the base building where the chamber is built is important. Will there be enough space between the lowest part of the building ceiling and the top of the chamber to install the structural beam, the fire sprinkler system, and the HVAC silencers? (See **Figure 2**.) For some projects the silencers can be mounted to the chamber sides.

For some projects, a pit will be needed—and this may require waterproofing, drainage, a sump pump, and structural design for the pit walls and floor (which needs to be massive enough for the vibration isolation for the chamber to be effective).

What are you going to measure in the chamber? Some ISO and other standards can allow a hemi-anechoic chamber to be used while other standards require a full anechoic chamber. I'll repeat the nitroglycerine warning—"know what you are doing before you do it"—and more than that—know what you want to do before you do it! And build it.

For instance, do you need to use a hemi-anechoic chamber due to a factor other than the acoustical criteria? In some chambers, a solid floor is needed since the equipment is very heavy, such as snowmobiles, autos, and the like (**Photo 1**).

Things to Consider

The setup can affect the measurement. Sources and receivers near the wedges can be impacted when close to the outer limit of the anechoic volume. **Photo 2** shows a device under test (DUT) in

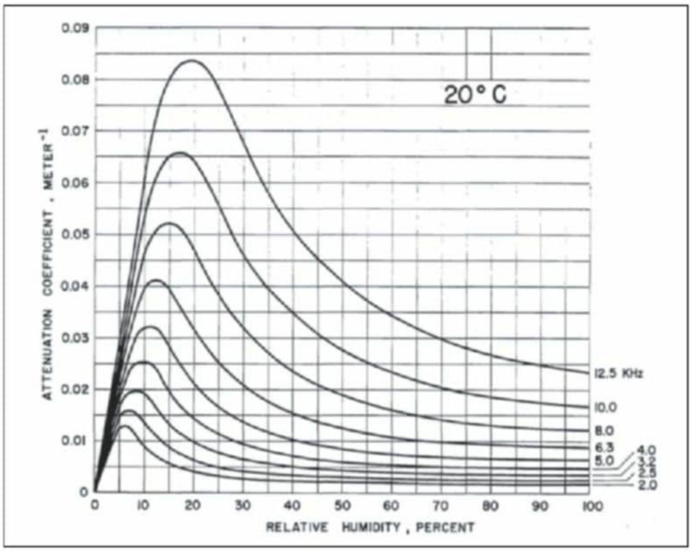

Figure 1: Values of the total attenuation coefficient m versus percent R. H. for air at 20°C and normal atmospheric pressure for frequencies between 2kHz and 12.5kHz at third-octave intervals. (Image source: Harris, NASA, 1967)

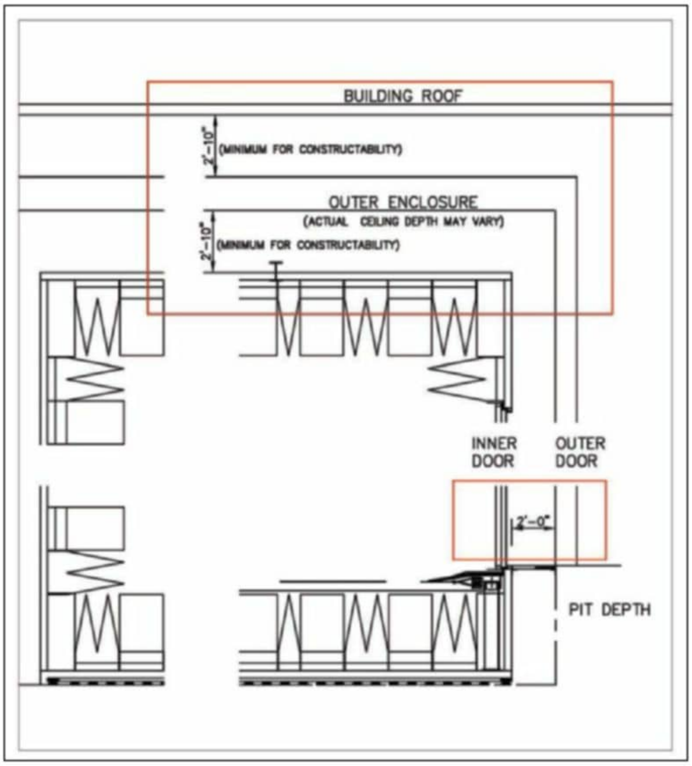

Figure 2: A chamber section. (Image source: Eckel)

Photo 1: A hemi-anechoic chamber. (Image source: Eckel)

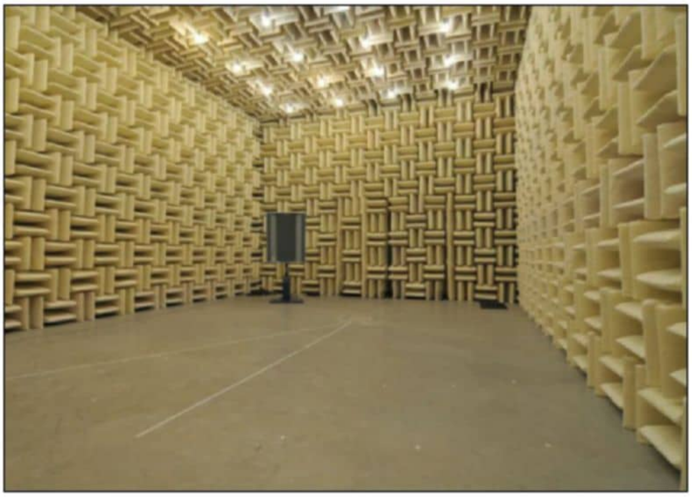

Photo 2: This is an Anslem Goertz chamber. (Image source: A. Goertz)

a corner and the measurement is near the opposite corner—once again, know how the test setup can affect the measurements.

Does placing the sound source in a corner affect the measurement? Can you always place the DUT in the best working portion of the anechoic volume? This requirement affects the required anechoic working volume, the overall size of the chamber enclosure, and the overall volume and dimensions to fit it where it is to be built!

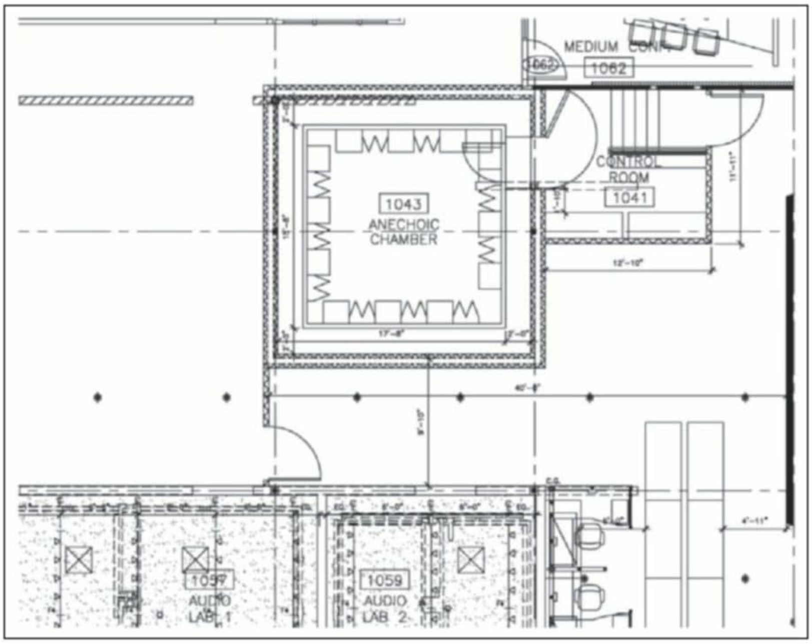

Figure 3: A space planning example of an above grade chamber. This is a 160Hz chamber placed in a large warehouse space near other research spaces. We were able to place it on grade and still accommodate the outer enclosure, but we needed to cut the slab and pour a large mass slab (equal to at least 10x chamber weight) due to the vibrations from non-audio equipment in the nearby audio labs. Note the door swing. If the control space could have been larger, the door may have been more in the center of the chamber. The platform at the top of the stairs just meets code and loading is only via stairs. And the chamber outer door swing is impaired by the fortune teller room 1062 and the control room space. (Image source: Callison)

Will live subjects (including humans) be in the room? if so, there are a whole lot of codes to deal with—egress, alarms, HVAC, and more—and just wait for the Fire Marshall and Building Department inspectors to let you know more about the more. Even without live subjects (including humans) conditioned air is needed to circulate in the chamber to keep the air at a relative humidity level that will not promote the growth of mold and other things. And then there is space planning.

Did I mention space planning? Try not to put the ore crusher (listening lab where high level audio can be present) next to the diamond cutters (anechoic chamber where the ambient noise criteria (NC) with the HVAC on is low, say at NC 10). A few space planning examples are shown in **Figures 3-5**.

For many projects, an outer box is needed to meet low ambient NC in the chamber. The typical anechoic chamber panelized wall can be up to STC-52 or so, which may not be enough to meet NC-15 criteria, let alone even lower ambient noise requirements.

Placing the chamber in an outer box can provide better sound isolation, but special care in the construction of a high sound isolation assembly is paramount or the money will not be well spent (**Figure 6**). The higher the needed transmission loss, the more care is needed. For a mediocre assembly, STC-50, such as that required for many multi-family resident projects, a leak 0.001% of the surface area can cause a loss of 3 points. For higher STC walls, an even smaller leak will become even more deleterious.

Once again, make sure it will fit and that it can be built—before you lease it or build it. For the 80Hz chamber section (Figure 5), a pit needs to be excavated to depress the chamber and to a lesser extent the control room.

And don't forget about the equipment on the anechoic chamber roof or side walls (HVAC, sprinklers, etc.) that might need to be installed—it can be tight.

Figure 4: A space planning example of Isolation. This anechoic chamber was located at the last minute near a high sound level listening room. This was solved with money—knowing the expected spectrum levels in the listening room and the background noise required in the chamber, a multi-leave partition between the spaces was designed with multiple layers of damped gypsum panel on the leaves, air gaps between leaves, an air gap between the chamber and the nearest leaf was developed for the calculated attenuation. (Image source: Callison)

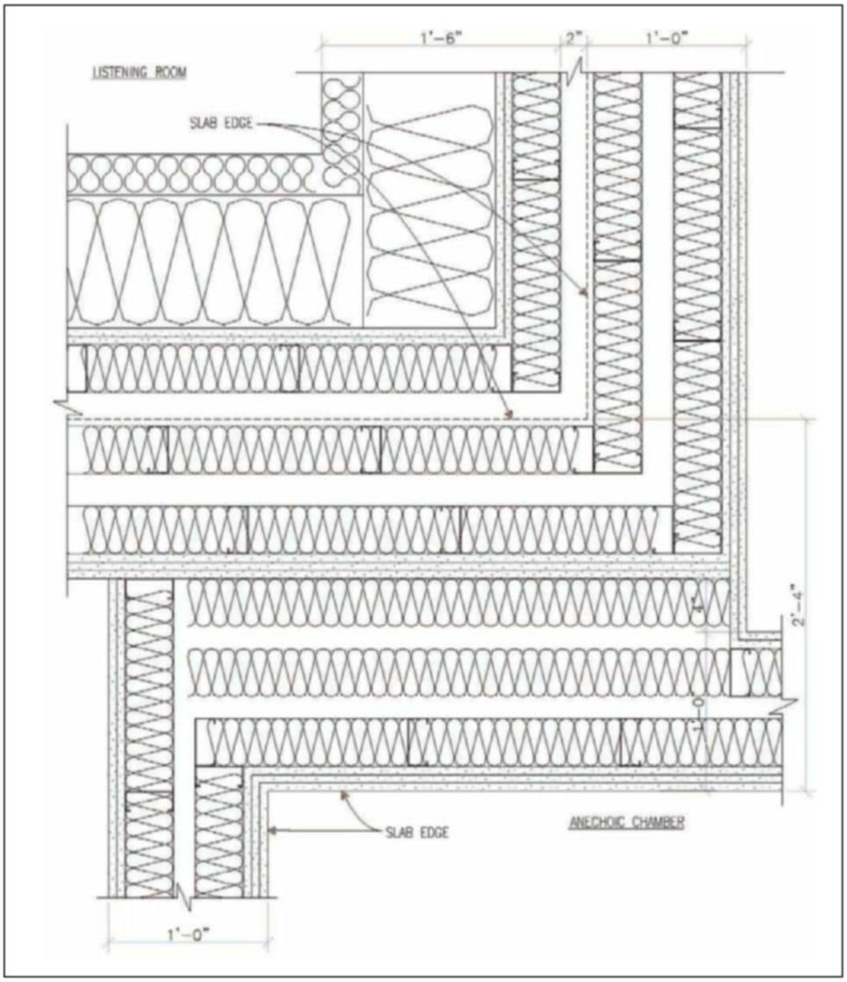

The air between partitions acts like a spring. The greater the distance between the walls—the clearance between the chamber exterior and the outer box interior—the smaller the spring constant of the air—which improves the low-frequency isolation.

Vibration isolation of the chamber and connections to the chamber is always a good idea to decrease vibration noise by noise that can be generated by the chamber vibrating. For this chamber there were two nearby highways to contend with. A site survey to determine the expected low-frequency noise generated from the floor is a good idea—it's cheaper and easier than fixing it after the fact.

The overall wedge size needs to be considered with care. When you have space planning restrictions and ambient noise criteria to meet, inches here or there for the overall chamber, the outer box dimensions, as well as the required air space, can have implications. Many chambers also have some space between the wedge bases and the enclosure for wedge mounting. Don't forget this.

The earlier the coordination with the chamber manufacturer the better. Once again, let the real estate department know what you need in terms of overall (gross) space (and adjacent property and spaces—don't let them situate you next to the cement plant!). One client leased three different buildings before understanding the volume of space needed was not available in the areas they leased. Another leased a building was next to a rail line—which we were able to mitigate—just add time and money. Again, money can solve many, but not all problems. The cheapest acoustics solution is proper space planning.

When you have an outer box, the clearances can be tight. Coordination with the base building mechanical engineer is best before they get into the mechanical system during the design development (DD) phase. When this is done during the construction document (CD) phase, it may be too late. Your acoustical consultant can determine the maximum noise at the connection from the building

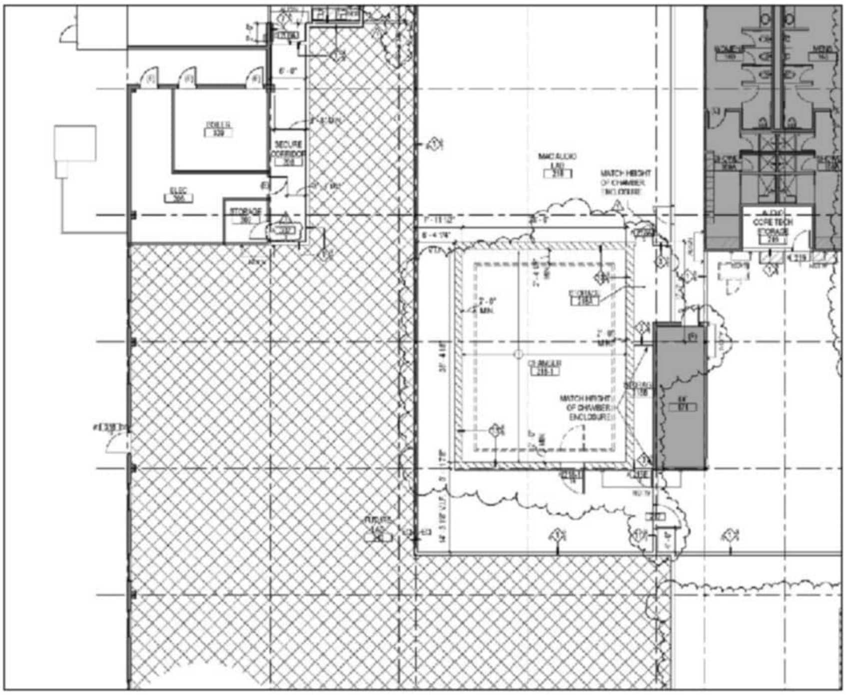

Figure 5: Another space planning example with a big 80Hz chamber. This large 80Hz chamber was surrounded mostly by unimproved space. Its future use was undetermined except for general office and support space. The problem here was the parking lot to the left. It had speed bumps and when tractor trailer trucks crossed the bumps the vibrations were sensible in the location where the chamber is located. So, in addition to needing to excavate a pit, the bottom of the pit was a separate isolated massive slab—and 2" deflection springs were used for the chamber. (Image source: HGA)

HVAC system to the chamber silencers (**Photo 3**). Be sure to pass this on to the mechanical designer and discuss how they are going to meet these criteria. The acoustical consultant can usually have this ready quite early in the design process if the ambient noise requirements of the testing protocols proposed for the chamber are known.

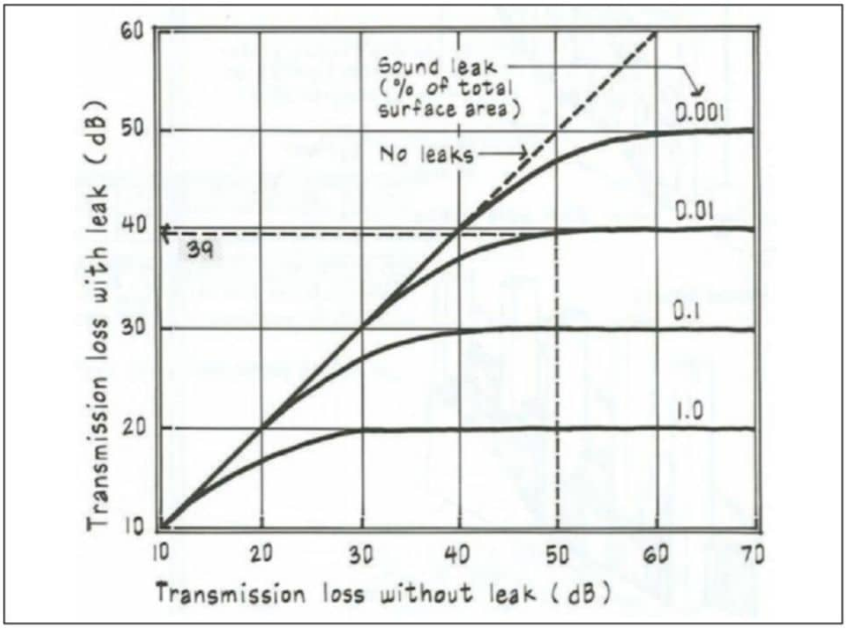

Figure 6: Sound leak chart. (Image source: Egan)

Photo 3: HVAC side silencer. We always need to remember the HVAC! (Image source: Eckel)

About the Author

Neil A. Shaw is a Principal with Menlo Scientific Acoustics, Inc. His projects range from performance and production spaces, sacred spaces, corporate meeting and conference facilities, academic buildings, and cruise ship vessels to research laboratories and include the Getty Villa, the Rose Bowl, Microsoft, Apple, Disney Cruise Lines, and the Government of St. Lucia. He studied Electrical Engineering at the Cooper Union and received his BS Engineering and MS Engineering degrees from the University of California, Los Angeles. He has served on the faculty at the Thornton School of Music at the University of Southern California and at the Southern California Institute of Architecture. He is a Fellow of the Acoustical Society of America and of the Audio Engineering Society.

In terms of space planning, it is always best practice to place the chamber at grade on an isolated slab at least 10 times the gross weight of the chamber. A large chamber on upper floors is not recommended.

And be sure that construction debris is removed! Debris can vibrate and generate noise due to low-frequency noise in the chamber—STC 52 panels do not have good low-frequency isolation.

The penetration work through the roof for sprinklers (and lighting) need coordination as these are usually installed by sub-contractors other than the chamber manufacturer's installer. Flexible connections are also needed for electrical and low voltage connections, and perhaps for other services (device air, device fuel, etc.).

Don't forget the chamber ceiling structural beam—the HVAC silencers cannot run perpendicular to the beam in low clearance situations. In any case, there is a maximum length for the silencers to allow connection to the building HVAC system (**Photo 4** and **Photo 5**).

Most measurements in the chamber require electrical power and low voltage connections. There are quite a few conduits here (**Photo 6**)—all with flexible connections for vibration isolation. Local codes will require electrical and electronic cables for different systems and with different voltages be segregated. In some instances, only the lighting and fire alarms are needed.

Passthroughs between the chamber and the control room can be used to provide a path for DUT power and measurement cables. Also note the vibration isolation loop for the sprinkler black and orange pipes.

Are you measuring a snow mobile or a loudspeaker? For the loudspeaker and electronic products being able to measure the self-noise of the device is important. As a well-known and respected consultant from Northen California asks: "Which has less self noise: a silk dome or an aluminum dome?"

Figure 7 shows a multitone signal where the tones peaks represent the frequency response and the dirt between the tones is the noise. A quiet background lets you see more of the "dirt" between the tones and then compare changes to the loudspeaker soft and hard parts—one change at a time is better to see the more cost effective change.

For a loudspeaker the choice for multitone testing tones includes, of course, the number of tones, the frequencies for the tones, the amplitude for each tone, and the phase for each tone. The tone frequency, amplitude, and phase is important—especially the frequency selection so any harmonics and intermodulation products are not hidden by a tone.

Photo 4: Sprinkler in the form of a duct. Coordination is important! The return air (RA) silencer and a conduit to a chamber light is shown here, this would have been quite interesting if it was the supply air (SA) silencer. (Image source: MSAI)

Photo 5: Sprinkler vibration isolation. Many chambers require sprinklers. Note the flex connections to the chamber roof mounted sprinklers. The yellow rope was removed when the construction was done. (Image source: MSAI)

If we are looking for the “dirt” between the tones, we know we need an acoustic environment that is quieter than the “noise,” and, what is often overlooked, the electrical and electronic noise may need to be even lower. Make sure the measurement equipment noise floor is substantially lower than the acoustic noise floor and the acoustic noise floor is lower than the DUT noise!

For some devices the electrical power noise can be important. This can be a factor when the DUT power supply rejection ratio (PSRR) is not sufficient to reject the noise in the AC power.

The ambient noise in a building typically increases as the frequency decreases and can also vary over time, so the ambient noise should be measured before and after a test, so you know if the noise was sufficient below the test measurement levels.

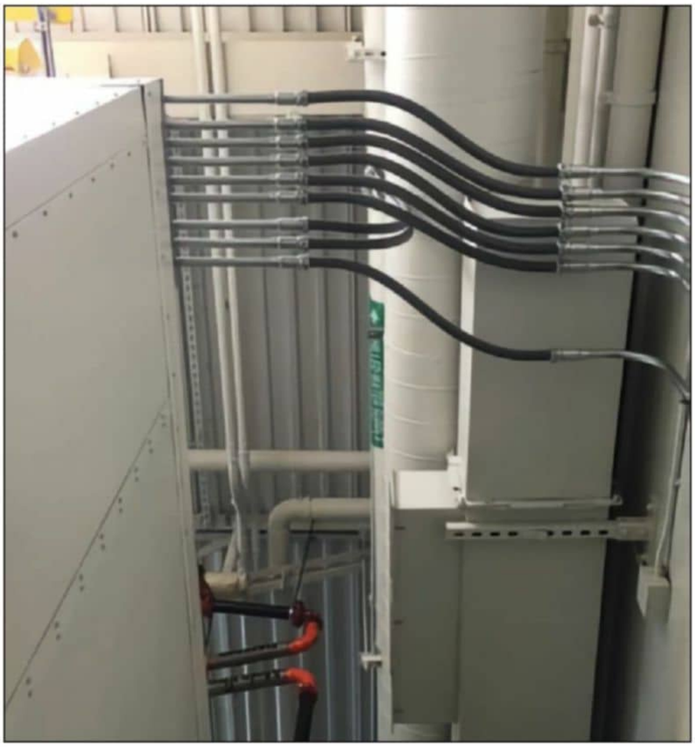

Photo 6: Lots of conduit. (Image source: Eckel)

The graph in **Figure 8** shows the noise floor for two 1/2” measurement microphones—an excellent general-purpose microphone, the Bruel & Kjaer 4189 with a ZC 0032 microphone pre-amplifier; and a special built very low-noise microphone, the Bruel & Kjaer 4199, with laser welded pre-amplifier—connected to a popular Type 1 sound analyzer, the Bruel & Kjaer 2250/2270, compared to the ISO Minimum Audible Field (the dashed green line) and some noise extrapolated NC curves (the dotted orange, blue, and purple lines).

Noise criteria (NC) is commonly used to describe the noise levels in a space and is the common metric for mechanical system noise in the United States. The 4189/2250’s self-noise (solid blue curve) is 17dBA, the 4955-A/2250’s self noise is 4dBA (solid red curve). Note the 1/3 octave NC curves below NC-15 are extrapolated. The overall A-weighted sound level, for noise that follows these curves from 63Hz to 8000Hz, is 15dBA for NC-0, 19dBA for NC-5 (extrapolated) and 23dBA for NC-10 (extrapolated). PCB, DAS, and Bruel & Kjaer, among others, offer some microphones with even lower noise floors.

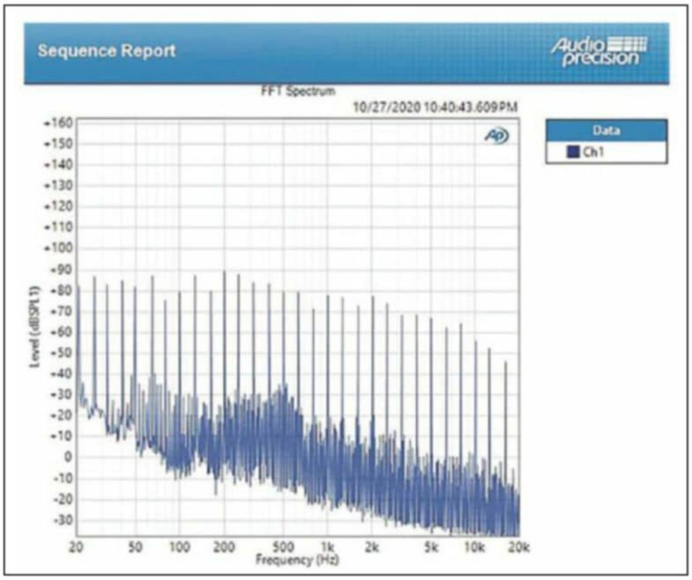

Figure 7: Multitone signal measurement. (Image source: MSL)

In many projects the electrical power distribution is large and has many parts. These various and sundry parts serve everything from rotating equipment like chillers and other HVAC equipment, kitchen equipment, lighting, convenience power, IT servers, and more, including your chamber.

If low and very low noise measurements are contemplated, power quality analysis is a good idea to find out what sort of electrical junk is going into the DUT. Some DUT power supplies have poor PSRR and will pass the dirt in the AC power right through to the DUT output.

This goes back to the test criteria. What is nearby? How is power distributed? For your facility you may need two sources of AV power—clean and very clean.

This should be done in the design phase, and at least make provisions for doing it later, since it will be more expensive and difficult to do when there is no room in the electrical room to accommodate the additional equipment.

A review of any shop drawings (which should show what and how the design will be built) before construction and scheduled observation visits during construction is needed to ensure that all transformer

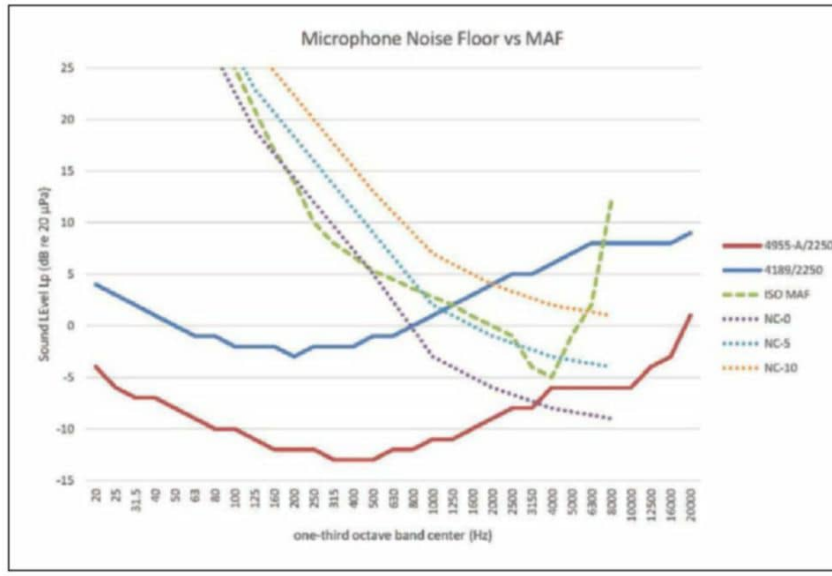

Figure 8: Microphone noise ISO minimum audible field noise criteria chart. NC 0, 5, and 10 extrapolated down from NC 15 (Image source: Bruel & Kjaer/MSAI)

Resources

ANSI/ASA S12.2-2019 “Criteria for Determining Room Noise” Acoustical Society of America, Melville, NY.

L. L. Beranek and H.P. Sleeper, Jr., “The design and construction of anechoic sound chambers,” *Journal of the Acoustical Society of America*, Volume 18, pp. 140-150 (1946).

M. D. Egan, *Architectural Acoustics*, McGraw Hill, New York, 1988.

and conduit vibration isolation is installed and installed correctly (**Figure 9**).

Final Thoughts

Get the criteria for the chamber's use right in the programming phase, be sure you have the right location when you start the design development phase, fully coordinate with the other design disciplines in the construction document phase, closely monitor the build during the construction phase, and then verify during commissioning.

As Branch Rickey said: "Luck is the residue of design and is governed by causes which are generally in the power of the man himself to govern."

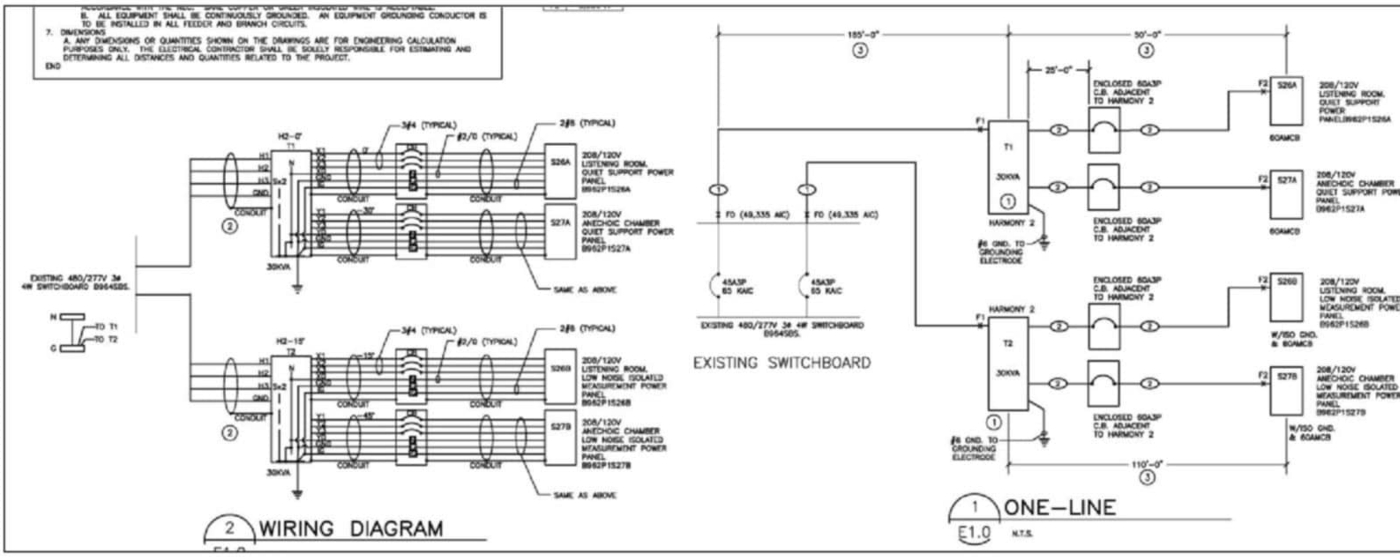

Figure 9: A clean and very clean power single line. (Image source: MSAI)

A Pocket Checklist for Anechoic Chambers

1. What is the minimum size (length, width, and height) of the room for the chamber? Put it in writing and then give the gross dimensions to the real estate.
2 What are the maximum length, width, and height of the equipment to be tested? Put it in writing and then give this to the designers as well as real estate.
3. What volume of air (in CFM) is required for testing and to also keep the temperature and humidity constant. Don't forget the heat dissipation of the DUT and humans. Put this in writing and give it to the mechanical engineer and the real estate. This demands a dedicated HVAC system for the chamber and associated control room solely under the control of the chamber personnel. You don't want to have a zone from the building system that turns off at 5 PM.
4. How much power does the electrical and electronic equipment that is to be tested need? How much power does the control room measurement equipment need? Put this in writing and give it to the electrical engineer and real estate. This may require a separate transformer or transformer from the building electrical room.
5. What lighting is required and for situations where there are live subjects (including humans) what is the minimum light level or wattage needed for the test?
6. What about noise from the light fixture(s)? You can always turn the lights off, unless there are human subjects where you will need an exit sign.
7. What is the number and size of doors (and very rarely) windows? This is important as the doors (and windows) can cause a variation in the uniformity in the anechoic volume.
8. What is the maximum ambient noise level versus frequency in the chamber?
9. What is or will be the ambient noise surrounding the test chamber? Best to know the un-weighted (dBZ) levels for the 1/3 octave bands of interest, and for at least one band below and one band above, as well as the overall level in dBA. Are there any tones in the noise? Is it intermittent?
10. Again, what is the low-frequency cutoff?
11. What fixtures are needed (e.g., IEC baffles, chain motors, liquid cooling, engine exhaust)?
12. What are the fire protection requirements? What does code require (e.g., wet pipe, dry pipe, inert gas)? What can be used and what happens when it is used?
13. CCTV camera system?
14.Access control?
15. Smoke/fire detectors?
16. Can the cable floor support the DUT? If not, then stations that can be used to support the fixtures are needed. Although there is usually a fine transondant cloth below the cable floor to catch smallish objects that are commonly used to test gravity, a drop cloth is recommended when assembling the fixture. Also, a grabber/magnet tool that fits through the cable floor mesh will be needed!
17. Can you get the fixture or pieces into the chamber? Doors can be a sound leak.

Acoustical Effects of Floor Finishes

In any room where acoustics are important—be it a studio, a rehearsal room, a lobby, a conference room, a classroom, an auditorium, a lecture hall, a sanctuary, or whatever—the effects of the floor finish must be considered. This article explains why.

By
Richard Honeycutt

Common floor coverings include ceramic tile, marble and other stone tiles, concrete, wood parquet, vinyl composition tile (VCT), luxury vinyl tile (LVT), linoleum sheet, carpet, carpet tile, and others. The first ones are acoustically hard materials, providing little acoustical absorption. **Figure 1** shows the sound absorption coefficient (α) of common hard floor finishes plotted against frequency.

When you concentrate on the low absorption coefficients (below 0.1), you may tend to discount them as insignificant. But remembering that the total absorption in Wallace Clement Sabine's equation is the product of absorption (α) times the area of the material, you see that for a sufficiently large area, a material with a low α can make a significant difference.

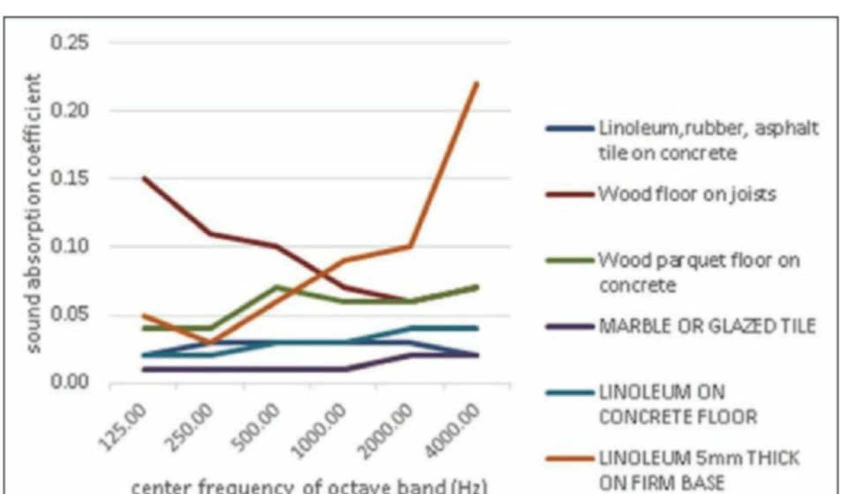

Figure 1: The acoustical absorption of various floor materials differs widely.

A high-school gym with dimensions of 110′ (L) × 90′ (W) × 33′ (H) will have a floor area of 9900ft². Part of the floor will be covered by seating, leaving perhaps 6200f² exposed. The wooden floor will have an α similar to that of wood parquet on concrete, so the total acoustical absorption in the 4kHz band will be 6400×0.07=448 sabins.

If no acoustical absorption is applied to the walls or ceiling and the gym is unoccupied, the reverberation time (RT) will be as shown in **Figure 2**. While the floor's contribution to total acoustical absorption is not large, neither is it negligible, even for an acoustically hard floor. Several acoustically hard floor finishes do have similar sound absorption behavior (e.g., VCT, LVT, and linoleum behave similarly as far as absorption is concerned.)

Real-World Examples

I am a member of a group that meets weekly in an acoustically hard room: VCT floor, cinderblock walls, and plaster ceiling. A few years ago, we moved to a different meeting room having dimensions of 50′ × 30′ × 10′. At our first meeting, we noticed that the room was excessively reverberant. Before the next meeting we added a 20′ × 30′, 3/8″-thick carpet.

The acoustical comfort and speech intelligibility improved noticeably. **Figure 3** shows the initial RT and the RT after the carpet was added. Although carpet is often over-rated as an acoustical absorber, if carpet of sufficient thickness is used, and/or the carpet is installed over an underpad, it can be very helpful. As you will also notice from Figure 3, carpet is not very helpful at low frequencies.

As you can see from **Figure 4**, carpet thickness and installation details have a significant effect upon the sound absorption coefficient. These plots apply to carpet on a wood floor. Carpet on concrete can only provide a sound absorption coefficient of about 0.12 at 125Hz, although the α can be much higher in the mid-frequency octave bands.

I was recently engaged to analyze the effects of a proposed change from a carpet-and-wood floor to ceramic tile in a large multipurpose auditorium. **Figure 5** shows the measured RT and the target average and maximum RT (with no occupants) for the programs presented in this room. The carpet was 3/8″ thick, laid without an underpad on a wood floor.

Computer modeling of the auditorium did not predict the same RT characteristics as measurements revealed. This is not uncommon and is a result of the variability of carpet construction and installation details. However, it was clear that the acoustics the room did not support speech intelligibility or music with the present floor finish. The large amount of sheetrock, which absorbs mainly low frequencies, caused the room to sound "thin" due to low RT at low frequencies, and the seat and seat-back cushions caused the room to sound "dull" because of low high-frequency RT.

To properly recommend acoustical treatment for an existing room, I usually measure the RT in the room, then build a computer model and calibrate it to the room so that its RT (unoccupied) reasonably matches the measured values in the octave bands from 125Hz through 4kHz. I do this calibration by modifying the α vs. frequency values of any materials whose published values could potentially be suspect.

In this case, the variability of published carpet sound absorption coefficients seemed to be most likely to contain inaccuracies. The seat cushions were second. Having completed the calibration, I analyzed types, mounts, and locations of acoustically absorptive material to achieve the optimum RT in the various octave bands.

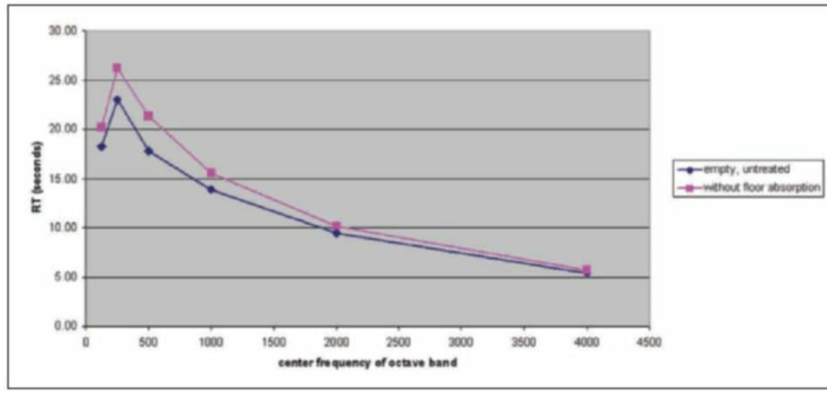

Figure 2: Ignoring the "low" absorption of a floor can lead to errors in predicting RT.

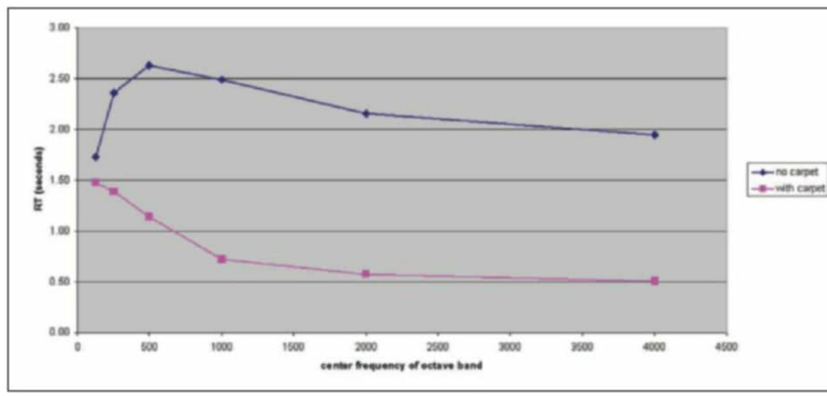

Figure 3: Adding an area carpet to a conference room can help control RT.

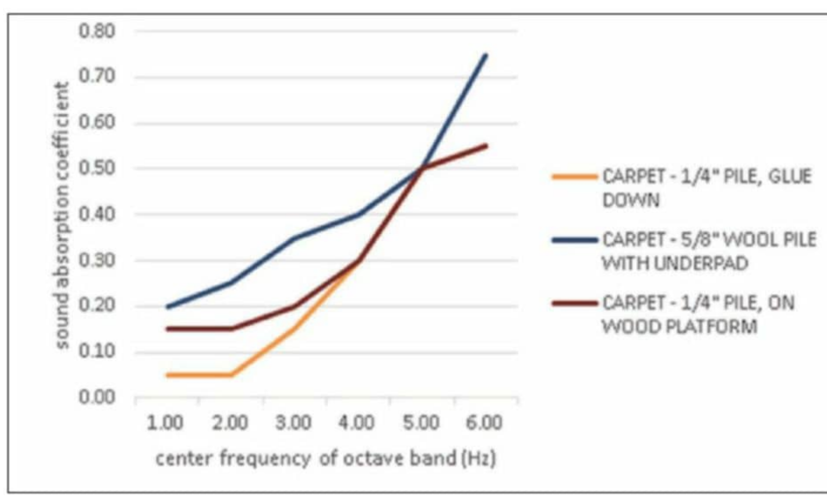

Figure 4: The weight of carpet greatly affects low-frequency sound absorption.

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.

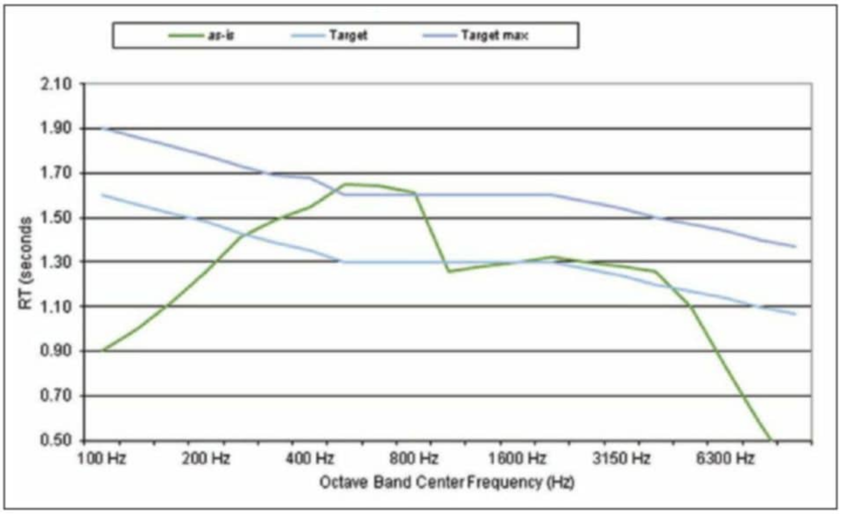

Figure 5: The measured RT in the unoccupied auditorium with an existing carpet-and-wood-floor only falls within the target RT range for a couple of midrange octave bands.

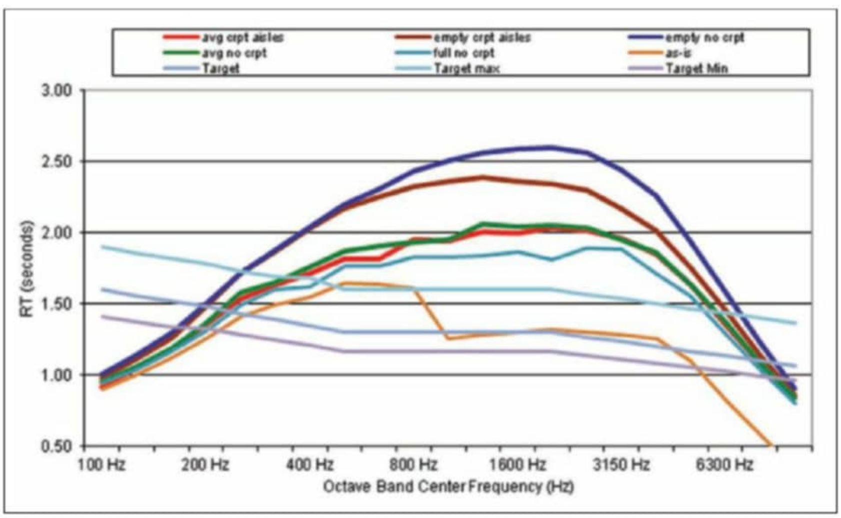

Figure 6: Analyzing the computer model of the auditorium with a variety of floor finishes did not reveal an optimum solution, if only the effect of the floor finish was considered.

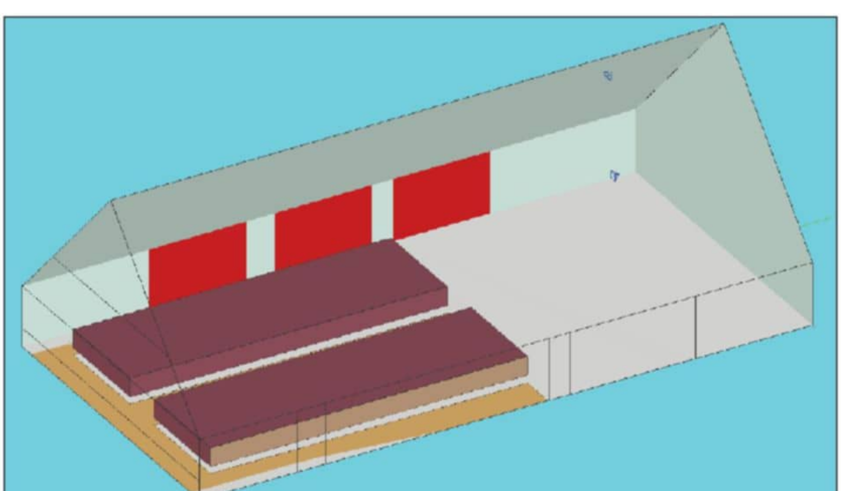

Photo 1: The pale orange on the floor represents 3/8″-thick carpet, and the red rectangles show the location of 1 1/8″ cloth-covered fiberglass panels.

The computer model analysis showed that simply replacing the existing floor finish with ceramic tile would result in an unacceptably high RT, with poor acoustics for music and only fair speech intelligibility (**Figure 6**). Although the RT profile improves when the auditorium has average or full occupancy, RT overall remains too high.

Since the predicted RT remained too high after changing to a tile floor unless some acoustical treatment was added to the sanctuary, I examined possible wall treatments to help control RT. With carpeted aisles and average occupancy, the RT slightly exceeded the target maximum from about 200Hz to 3kHz.

Applying 1″ cloth-covered fiberglass to portions of the rear side walls resulted in a slightly excessive RT below about 400Hz with RT tapering below the target at higher frequencies. Changing the side wall treatment to 1 1/8″ cloth-covered fiberglass eliminated most of the low-frequency excess while adding an insignificant 0.1-second dip in RT between 2kHz and 2.5kHz.

I recommended applying 3/8″-pile carpet behind the rear pew and in the aisles. This also has the advantage of helping to silence heel-tap noise when congregants enter or leave during a service.

In addition to carpeting the aisles, I recommended treating the side walls in the rear using three segments of 1 1/8″ cloth-covered fiberglass from 3′ above the floor to the bottom of the soffit on both sides (**Photo 1**). This type of fiberglass has a 1/8″ thick hardened layer just under the cloth, improving low-frequency absorption and supplying impact resistance (**Photo 2**).

As discussed in previous Sound Control articles, a given total area of sound-absorptive material behaves like a larger area if it is segmented into smaller blocks, due to diffraction effects. **Figure 7** shows the resulting predicted RT plots. Although the RT is a bit lower than the average target above the 6300Hz octave band, this treatment will increase the STI to 0.67, bringing it from the present value just at the border between FAIR and GOOD to a value in the upper GOOD range.

Theoretical Background

Sabine's analysis of the effects of acoustical absorption led to the development of his famous equation that allows the prediction of the RT of a room from a knowledge of its physical dimensions and the acoustical absorption in the room:

$$RT = \frac{0.161V}{\sum S\alpha}$$

where V = the volume of the room in cubic meters, and the denominator of the fraction signifies the addition of the products of the area (S) and absorption (α) of each material. This is a statistically based equation, rather than a theoretically based one.

As early as the 1932 publication of Vern Knudson's *Architectural Acoustics*, a problem with Sabine's approach was noted. In Appendix II of Knudsen's book, R. F. Norris pointed out that Sabine assumed an absorption coefficient of unity for a perfect absorber such as an open window;

Photo 2: A thin hardened layer of fiberglass under the cloth adds impact resistance and improves low-frequency absorption.

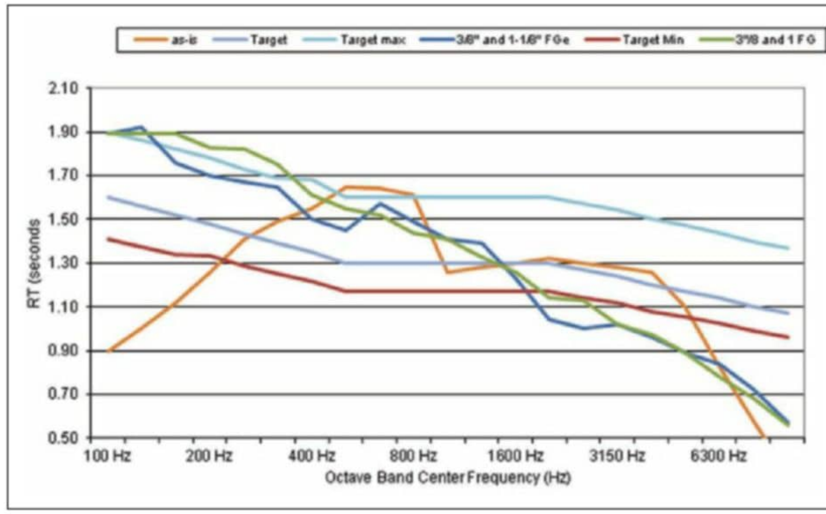

Figure 7: Using a combination of partial carpeting and absorptive wall treatment, we can significantly improve the auditorium's acoustics. The red rectangles represent 1 1/8" cloth-covered fiberglass panels.

whereas, that value used in the Sabine equation does not yield a prediction of zero RT. Denoting Sabine's absorption coefficient as *a*, and the "true absorption coefficient" as α, Norris derived an equation based on the mean free path of a sound wave in a reverberant room, from which he defined the true absorption.

Carl F. Eyring noted, as had others, that Sabine's formula worked better for live rooms (little absorption) than for dead rooms (lots of absorption). Norris and Eyring published a new statistically based equation that proved more accurate for acoustically dead rooms (low RT) than the Sabine equation. [1]

All the statistical equations for RT are based on the assumptions of homogeneity (at any point in time, the reverberant sound level is the same at all points in the room), isotropy (reverberant sound is equally likely to arrive from any direction), and diffusiveness (sound from any point can spread equally in all directions). In practice, this means there is no focusing of sound, and the acoustical absorption is equally distributed on all surfaces in the room.

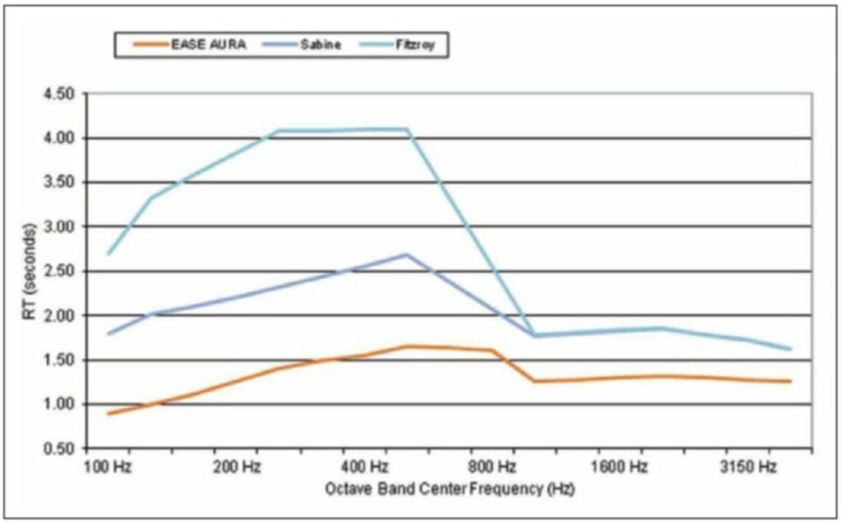

Figure 8: Statistical and computer ray-tracing RT predictions yield different results for rooms not matching the Sabine criteria.

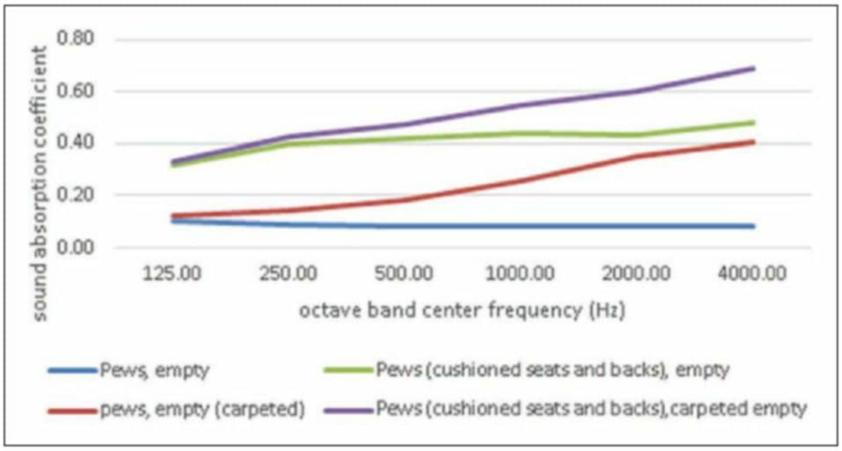

Figure 9: The sound-absorptive effect of carpeting differs depending upon whether the carpet is covered by seating.

In 1959, Daniel Fitzroy published "Reverberation Formula Which Seems to Be More Accurate with Nonuniform Distribution of Absorption" [2], in which he introduced an empirically derived equation for calculating RT that took into account varying acoustical absorptivity in three orthogonal directions (x, y, and z):

$$T = \frac{S_x}{S} \times \frac{0.161V}{S(-1n(1-\alpha))} + \frac{S_y}{S} \times \frac{0.161V}{S\left(-1n\left(1-\alpha_y\right)\right)} + \frac{S_z}{S} \times \frac{0.161V}{S\left(-1n\left(1-a_z\right)\right)}$$

where S_x, S_y, and S_z correspond to the areas of absorption in planes perpendicular to the x, y, and z axes and α_x, α_y, and α_z are the corresponding absorption coefficients. The effectiveness of this formula helps to overcome the inaccuracies imposed by lack of isotropy and diffusiveness.

Figure 8 shows a comparison of the results of Sabine, Eyring, and Fitzroy's statistical acoustics RT predictions, and EASE AURA ray-tracing-based predictions. The Fitzroy equation matches the (presumed accurate) EASE AURA results more closely than the Sabine equation. In fact, in rooms matching the Sabine assumptions and having an average α of 0.3 or lower, the Sabine method, and EASE and other computer ray-tracing methods yield very similar results. This room does not match the Sabine criteria.

Combined Effects of Carpet and Seating Upon Acoustical Absorption

If a carpeted floor is covered by furniture that extends all the way to the floor, the covered area of carpet cannot contribute to the total sound absorption in the room. **Figure 9** shows the sound absorption coefficient plotted versus frequency for empty church pews, both cushioned and uncushioned, with carpeted and uncarpeted floors. If you compare these curves with the absorption data for carpets not covered by seating (shown in Figure 4), you will notice that the seating greatly affects the sound absorption of the area, as does cushioning of the seats. For fully occupied seating areas, the absorptive effect of carpeting is pretty much swamped out by the absorption of the occupants.

References

[1] C. F. Eyring, "Reverberation Time in 'Dead' Rooms," *Journal of the Acoustical Society of America* (JASA), Volume 1, pp. 217–241 (1930), https://doi.org/10.1121/1.1915175

[2] D. Fitzroy, "Reverberation Formula Which Seems to Be More Accurate with Nonuniform Distribution of Absorption," *Journal of the Acoustical Society of America* (JASA), Volume 31, No. 7, July 1959.