Back

Speech Intelligibility – Part 3

The author concludes his series with a look at problems—and solutions—associated with reverb, resonances, and reflections.

PHOTO 3: Treatment too thin.

A classic method of measuring reverb is to clap two boards together and use a stopwatch to time how long the sound takes to decay. The better method of the day was to sound an organ pipe and time its decay. More sophisticated measures, of course, evolved as technologies advanced. Much of the early work was at 500 or 512Hz, which is still the design center frequency for music; for speech 2000Hz is preferable. The modern version of this is a small sound system, such as the "Sound Strobe" (*aX* 3/06). Now you can just twiddle a knob to get the impulse you need for listening tests.

REVERB PROBLEMS

Most of the computer-based audio measurement systems will also give a reverb time number or graph. They typically derive the reverb time from other kinds of test signals. When you use computer-based measurement systems to do this, you encounter that important warning: "To err is human, to really foul things up requires a computer!" (Anon.)

The reverberant field decay is usually presented as a small matrix or graphs of the actual frequency and the time it takes to decay by 60dB. A typical reverb time meter such as the Goldline GL60 model uses octave bands from 125Hz-4000Hz. It is different in that it measures only the first 20dB of decay, then presents the number as though it were 60dB of decay.

That is because when I designed it, I was aiming for 20dB of signal-to-noise ratio to ensure good speech intelligibility. I knew then that it was the first drop in the sound that affected this. The 60dB number is still the classic reference, so that was my way of presenting useful data to the less experienced user. The meter also has the "bad" habit of not settling on a single reverb number if the field does not have a uniform rate of decay.

The model of speech I use is a burst of sound followed by silence for twice as long. I had a recording of Andrew Carnegie speaking at the turn of the century in a large reverberant room. He spoke slowly by today's standards—about three syllables per second. A fast talker today will speak six. For my model I use a rectangular gate that is open five times per second. Thus the on time is 1⁄15 of a second and the off time is 2⁄15 of a second. Others who have studied actual speech use different time intervals. As wrong as they may be, my numbers work for me!

You can use a meter such as the Goldline to determine not only a matrix of reverb numbers, but also how reverb changes in a given room. The manual (not mine) suggests you use a pink noise source with the meter. The sound directly from the source drops off as you move farther away. The reverberant field in a room will continue to build until the absorption and leakage out of the room equal the rate at which energy is added. This is the steady-state condition.

Turning off the noise while you are far away from the sound source will trigger the meter and give you a classic reverb time number. Just remember, this is for only the first 20dB of the curve. There are some cases in which this does not happen. The meter bounces all over and does not settle on a single number, which is a strong indication that the room is not linear and there may be specific echoes that interfere with speech intelligibility. The room is most likely not very music friendly either.

If you produce a sound impulse right next to the meter, the number it computes will be much lower than if you repeat the experiment with the meter across the room. If you use a Sound Strobe or other impulse generator through an existing or portable sound system and move away with the reverb time meter, at some point the meter or your calibrated ears will indicate a reverb time of 1.5 seconds at 2000Hz. Note this distance, which is most likely the same distance you would get from the speech articulation tests for an acceptable loss distance.

Assuming that you speak at five syllables per second and you require about 5dB of decay before the next syllable to get acceptable understanding, then you would like to see a decay rate of 25dB in two-thirds of a second. The other third of a second is your sounding time. This gives you a classic good speech reverb time of 1.44 seconds, basically what the meter just showed you.

CRITICAL DISTANCE

As you move away from the constant sound source, the level directly from the source decreases. The level of the reverberant field for a steady source will reach a constant level throughout a room. The distance at which the direct field equals the reverberant field is called the critical distance, which you can measure with a constant noise source and a sound level meter.

Start at the farthest point and walk toward the sound source. You are at the critical distance when the level rises by 3dB. With some practice you can get close with just your ears and not-

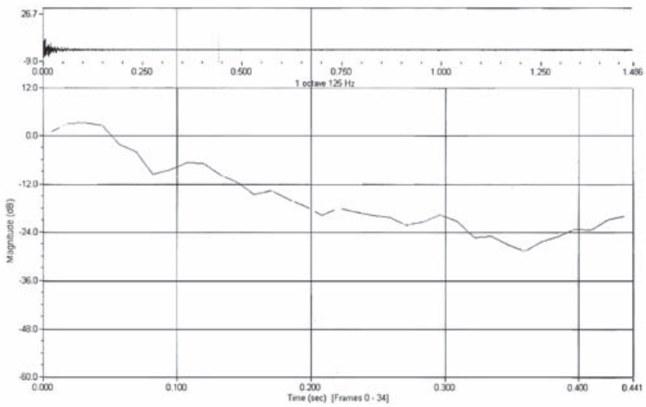

FIGURE 1: Pre-absorption reverb.

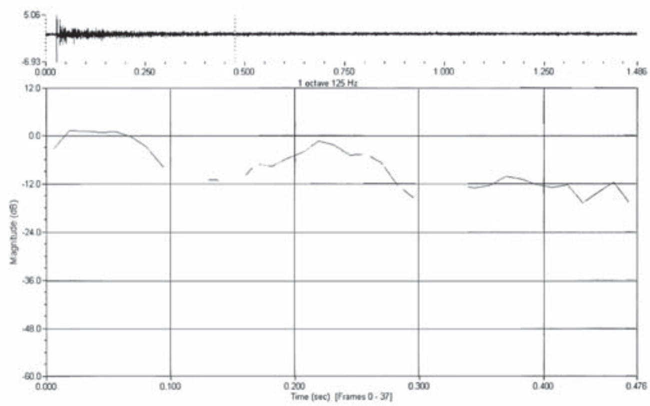

FIGURE 2: Post absorption reverb.

ing where the volume begins to increase. The rise means the sound level has doubled. One-half is the direct field, the other half is the reverberant field. In a really bad room this distance should be about the same as your speech articulation loss distance.

That is why when you move away from the pulsed noise the reverb time meter shows agreement with the articulation tests. This is the maximum distance for unaided acceptable speech in the room. This will be at least the critical distance. At the critical distance the signal-to-noise ratio is by definition 0dB for continuous sounds.

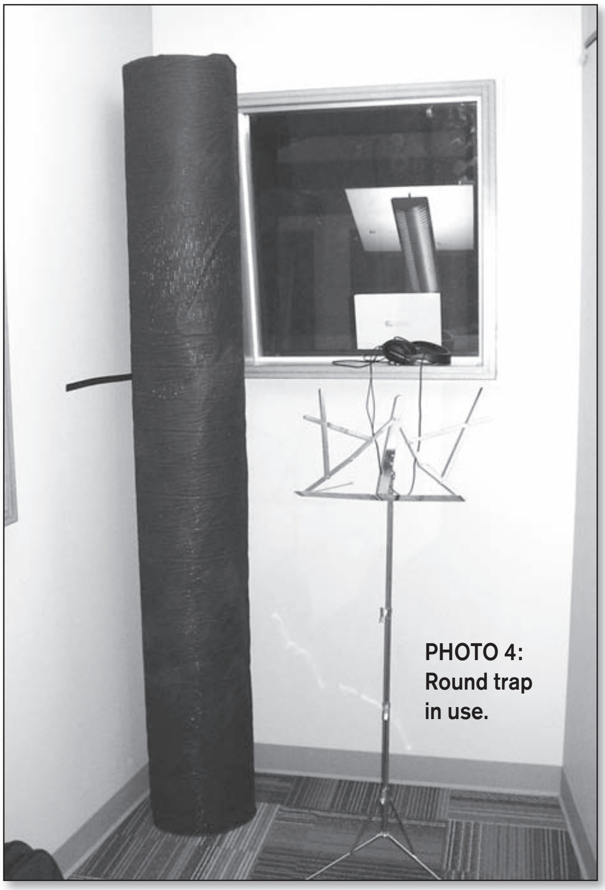

PHOTO 4: Round trap in use.

Because we speak in bursts, the reverberant field will not build to the same level as for continuous sound. With the model I use one-third on, two-thirds off—the ratio at the critical distance is at least 4.8dB signal to 0dB noise. This would be the worst case with an almost infinite reverb time.

If you know the reverb time is less than infinite, you can calculate how much decay will occur between syllables and add that to the signal-to-noise ratio. You also know that the talkers' speech drops off with distance squared. You can now calculate where the signal-to-noise level drops to 5dB in your room. You may wish to raise this goal depending on the expected use of the room.

The critical distance can change if you approach from different directions when you measure the critical distance with a pulsed noise source. It is possible that this is due to room geometry or construction producing a nonuniform reverberant field. The most common cause is that the sound source has directional characteristics.

You can use this ability of loudspeakers to be more directional than a plain human talking to your advantage. The idea is to put sound on the audience only where it will be absorbed. That way you do not feed energy into the reverberant field. The directional loudspeakers will increase the effective speech intelligibility distance. It is possible to deliberately design very directional loudspeakers.

One theory states that you can make a sound system overcome bad acoustics. As some ancients believed, the gods made us vain so they can enjoy the fall! You can make some improvements—horrible to bad or bad to okay, for example. But horrible to good will just not happen with loudspeakers alone. You must improve the room.

One of humans' annoying traits is our preference for simple answers. Is the room good or bad? Exactly what does it take to be good? In this case, you could hire one of my acoustician friends, pay them lots of money, follow their advice, pay attention to the details, and you will have a good room. But for those of you who want to know more. . .

A CHURCH STORY

A local large church had a problem with speech not being well understood after a recent building renovation. The church, as far as I could tell, was originally built with an interior finish of plaster and stone. When this proved to be far too reverberant, the interior plaster ceilings were sprayed with a sound-absorbing plaster and the main ceiling covered with absorbing panels. This left the room adequate to good for speech but very unpleasant for music.

When 30 years had passed and no one left could remember the original problems, the room was renovated to improve the musicality of the room. No acoustician was consulted. The designer plastered over the absorbent tile ceiling and painted the absorbent plaster. Now the room had lots of reverb and a few folks incorrectly thought it was improved, but it was now good for neither speech nor music. In addition, the room was much noisier.

The noise problem involved an organ blower chamber isolated from the remodeled room by a masonry wall except where it had been cut to accept a heating radiator. Only a piece of sheet metal blocked the visible opening. Closing the opening with a double layer of drywall and sealing reduced the noise.

Many folks proposed installing large and very expensive loudspeakers to project speech the length of the room. Three different systems were tried and none proved adequate. I had a few ideas regarding the problem, so I contacted Sam Berkow (Sia Acoustics-Jazz at Lincoln Center, Hollywood Bowl, and Smaart Software) to visit the church. He took measurements using Smaart software and determined the major problem was that all the sound energy was not decaying uniformly. In fact, the energy at the 250Hz octave band was decaying so slowly that it was masking speech in the room. In addition, he thought that the curved roofs on the transepts (side seating areas) were the primary cause.

He suggested adding enough absorption to drop the reverb time of the

THE DECIBEL

The decibel is one of those wonderful bits of jargon that we use to confuse the uninitiated. Please do not read further unless you already know the secret sign and password.

The base unit is the bel, which is the logarithm base 10 of the ratio of the sound being measured to the threshold of hearing. Thus, if the sound energy is ten times louder than our threshold of hearing it is log(10/1)=1. If the sound is 100 times louder, it is 2 bels, 1000 times louder is 3 bels, and so on. Just count the number of zeroes after the 1.

You can hear over a range of greater than 12 bels (120 decibels), which, when written out, is 1,000,000,000,000 to 1. Writing 120dB is much easier.

The decibel becomes confusing when you have, for instance, only twice the energy. The math works out to 10 (to convert from bels to decibels) × logarithm (2/1). Use a calculator to do the logarithm, which works out to a two-fold increase, or 3.010299957. . . decibels, which you can round off to 3. With ten times as much energy (1 bel), you get 10 decibels. If you double the sounds' energy again, it cannot be 20 decibels because that is 100 times as loud.

You now know the secret of logarithms: add the 3 from the doubling to the 10 from the ten-fold increase to determine that the energy increases by 13dB. Thus, 40 times as much energy would be 3 + 3 + 10, or 16dB. Eighty times, of course, is 3 + 3 + 3 + 10, or 19dB.

If you know that 20dB is 100 times and 19dB is 80 times, then what is 1dB? Well, 100 divided by 80 is 1.25. The log of 1.25 is .0969100. Ten times that is close enough to 1 that you should see how this works. The precise answer is 10 to the exponent (1dB/10 decibels to bel), which equals 1.258925412. . ., or 1.

Another consideration is the accuracy to which you can measure. It is fine measuring a voltage on a digital meter that is accurate to four digits. A change in barometric pressure of 1″ will change some sound levels by about .2dB. Wind and air movement also affects the readings. That is why sound pressure levels are rarely used with a resolution of greater than 1.0dB.—**ES** room to 1.5 seconds with a reasonable crowd; in particular, putting absorbers in the transepts. As soon as we added two absorbers to a transept, it became immediately clear that he had identified the source of the resonance. Treating both transepts and adding material along the side walkway ceilings provided the acoustic treatment needed to change the room from very bad to good. In addition, because of the location of the absorbers, the aesthetics of the church remained unchanged. The moral of the story—find out exactly what the problem is before trying to fix it.

RESONANCES

Any two parallel walls will bounce sound back and forth. Most rooms have at least three pairs of such surfaces. If the walls are good reflectors, this will show up at the frequencies which have a half wavelength that is a multiple of the spacing distance. There are other kinds of resonators besides parallel surfaces.

It is easy to understand that a wind instrument such as a flute uses the air motion inside it to produce a tone. Sound travels at a speed of about 1132′ per second. If you have the air bouncing between two acoustic nonlinearities, it will form a standing wave. Even random air motions will excite the wave and add energy to it as it bounces back and forth. The other common form of resonator is like a whistle in which the air rolls around in a circle and reinforces the resulting sound wave.

All rooms have many resonances. So the concern involves resonances that occur at particular frequencies to which we are especially sensitive or are very strong.

Finding out you have a problem resonance is simple. It will show up as a frequency or band of frequencies that do not decay as rapidly as the surrounding bands of frequencies. You can measure this with test equipment or use a musical instrument such as a piano or organ. You just listen for notes that hang around too long. This is sometimes called "color."

Figuring out which part of the room causes the problem resonance is a bit harder. Once you are able to produce or measure the resonance on demand, you can move a large absorber around until you find out where it has the most effect. Sometimes you will need to use many absorbers and cover most of a wall to produce a noticeable effect. The difficult times occur when you have a bell tower that may be causing problems and you want to try material on the ceiling 120´ above you!

This may sound difficult, and it is. Finding out the cause of a problem is more art than science. Although there are some instruments that help, they are not common or readily available. The good news is that from my experience, resonances that cause a real problem with speech intelligibility are present in about 10% of the rooms I examine. Many small resonances at the right frequencies can enhance music.

The common mistake in trying to fix a resonance is to place foam on the walls until the problem goes away. If the foam is not thick enough to absorb the frequencies of interest, it is not doing what you want and is just killing the mids and highs. This also colors the sound. If you want to use movable panels, be sure they are really big and thick.

Once you have identified the area that causes the resonance, you can add just enough absorption to damp it. This is different than trying to add absorption to get a desired reverb time. Because of the nature of resonance, the absorption has twice as much or more effect than it would if it were just reverb.

A second approach would be to change the surface. Building out, adding angled surfaces, or features of interest such as reliefs, or even using shelving to break up the characteristics that allow a single frequency to build are often used. If you can't absorb it, bounce it around enough so that it spreads out and dissipates.

REFLECTIONS

Of course, by now you are thinking, this is nuts. "How can this be right if you don't need to use a computer or any ludicrously expensive test gear?"

That's right! With simple measurements and calculations, you can achieve good reverb times in a room with nonuniform absorption and nonlinear reverb time curve. The only equipment you need is a pair of ears, paper, and a pencil. A calculator helps.

Of course, sometimes this is not enough. You have until now been looking at statistical averages of the sound field. Another problem area to address involves echoes or sharp distinct reflections that are not muddled in the statistical average.

If you add bookcases to the side and front walls, a small room would sound much better even with the same reverb time. The decay curve would now be much smoother due to the diffusion you have added.

Helmut Haas (1958) determined that a single echo can reinforce speech intelligibility or destroy it. He concluded that if the repetition of a sound comes within 30ms, you perceive it as part of the original sound, perhaps adding a bit of spaciousness or volume to the original. The second sound must be 10dB louder than the first to be perceived at all.

When the second sound comes after this period, it disturbs the speech. The effect is greater for faster speech. Some forms of late energy can even destroy speech intelligibility. Increasing the volume of the speech does not change this effect because the echo also increases. The direction of the echo does not have a great effect on sounds that originate from the front.

As a demonstration, I ask two singers to sing a capella the same piece while walking away from each other. When they get to around 30–35´ apart, they find they cannot sing together. There is a slight variance due to the tempo of the piece.

If you have ever tried Sound Strobe on a loudspeaker that has a path length difference at the crossover frequency, under some conditions you can hear a distinct double click even though the path length difference may be less than a foot. Haas used speech for his model; it does not encompass dissimilar sounds or frequency effects. It does work great for speech reinforcement design if you do not act silly trying to stretch the results of his work.

RAY TRACING

When you try to predict what the echo pattern of the sound will be like, you use the ray tracing model of sound dispersion. Simply put, you draw arrows pointing out from the sound source. When the rays strike a wall, ceiling, or other surface, you draw them bouncing onward as though they hit a mirror. Continue this path until the rays (arrows) reach where the listener(s) should be.

You can then assume that the length of the ray is directly proportional to the time it takes the sound impulse to reach the listener. It is useful to plot the first direct ray, the second rays that bounce once off a wall or ceiling, and the third rays that bounce off two surfaces. This is usually enough to get an idea of what the ray response will be like. You can do more if you like, but the task becomes almost impossible for five to seven bounces.

You can also scale the amplitude of the sound rays by multiplying the inverse square law value by the absorption coefficient of the surface materials. Sometimes the reflecting surface is small enough or angled so that it reflects only part of the ray and spreads out part of it. The wavelength or frequency of interest determines what size of non-uniformity becomes effective.

Sometimes the room is too small to use the ray simplification at the lower frequencies. I stop using the ray model if the room is too small to allow seven wavelengths of the frequency of interest. This allows you to produce a graph of amplitude versus time.

If you want to cheat a bit, you can also plot the reverb time curve on the same scale. Because the reverb time is usually based on steady-state energy excitation and the ray tracing assumes a unit impulse, you may wish to move the reverb time curve down a bit on the X axis. If you plot enough ray bounces, you will find it merges into a giant jumble that about matches the classic reverb time and that is where you can merge the curves. Computer programs are useful for this, but you can still do this by hand!

All the unit spikes up to 30ms are combined as though they are one great sound. Some folks combine the rest of the mess for a second value. When the first value is 5–10dB greater than the second, speech intelligibility will be good. The exact value is influenced by your exact methodology, the noise level, a few other semi-mythical issues, and the intended audience.

When the combination curve shows a nice drop of about 5–7dB for the initial break and a smooth linear tail without any major jaggies, you probably have a good room for music. When the reflections coming in past 30ms are stronger than desired, you can add absorption to the reflecting surfaces. One mistake is to treat one sidewall and not the other. The reflections off the sidewalls contribute to the spaciousness of a room. If sidewall treatment is needed, use it sparingly and split it up onto both walls and stagger it. Sometimes this will decrease the overall reverb time too much.

The trick, then, is to use a diffusing surface—anything from columns along a wall, moldings, or even strips of absorbent material. This will spread out the reflection, decreasing its level and adding to the reverberant tail.

Sometimes you wish to reinforce the early energy of the first 30ms, which tends to improve the intimacy of a space. This is where a low ceiling angle to maximize projection to the audience works wonders. I once was involved in a project with acoustician David Klepper, who designed a glass reflector system that was on a remote control. At the first use Mimi Lerner was the soloist.

I did not like the sound and asked her what she thought. She was diplomatic but agreed to sing a very brief bit (conserving her voice for the second performance). I lowered the reflectors by about 3´ and heard an amazing difference as the path length locked into the magic 30´. The second performance was much better!

In the 1940s studies of movie theaters surprisingly concluded that some of the "good" examples had worse reproduction system frequency response than some of the "bad" examples. An excellent example that the nature of the sound reflections, reverb, and noise can be more important than just "fidelity."

One question that comes up from time to time is, "If we know so much these days, how come all the old concert halls are good halls and many newer halls are not as good?" The answer is simple: When you build a bad concert hall you either fix it or tear it down and build another one! I hope you have picked up enough useful information to improve the acoustic spaces you care about and will not need to resort to the wrecking ball! aX