

MAXIMIZING LOUDSPEAKER PERFORMANCE IN ROOMS

Part One: Why Loudspeakers Sound the Way They Do.

by

Floyd E. Toole, Ph.D.

Vice President Acoustical Engineering, Harman International Industries, Inc.

I am old enough to remember when having a tweeter in a loudspeaker system was a novelty. I owned 12-inch drivers with “whizzer” cones (thin paper cups attached to the centers of the diaphragms) to help give them some high-frequency output. In those days bigger was clearly associated with better and, even with huge speakers, the volume had to be turned up very loud in order to hear any real bass. I remember absolutely needing and using tone controls, and when “good tone”, not accuracy, was the generally recognized standard of excellence.

In those days, loudspeaker manufacturers routinely made tremendous leaps of faith. They sold woofers, midranges and tweeters to customers who combined them in any mixture of brands, using a crossover network that was not designed for any specific units. All of this might be put into a homemade enclosure following a design that was published in one of the monthly audio/electronics magazines. Back in the late 1950’s my father and I built several of these horrors, not knowing any better. One, called the “fold-a-flex” was an ingenious design that, by operating doors and sliding panels could be converted from a closed box, to a reflex, to a folded horn! Another, the “Karlson”, survived long enough for me to be able to properly test it, after I matured, acoustically. It could best be described as an “acoustic meat grinder”, guaranteed to make mincemeat from any drivers put into it! Those were the days when the quality of the sound increased in direct proportion to the effort expended in the hand-rubbed finish on the box.

Were these the “good old days” of audio? No, they were just the “old days” before we learned how to engineer loudspeakers. In those days, all loudspeakers sounded different from each other, and none of them were truly good. I think we were still a bit dazzled by the “miracle” of hi-fi, that it was actually possible to record and reproduce something that resembled music. If we choose an adjective to describe these days, I would pick “colorful”. The experiences, the mythology, the charismatic characters, the stories, the memories and the sounds of the loudspeakers, were all colorful. How a loudspeaker sounded, in this era, was as much a matter of accident, as it was design.

We have come a very long way since then, but progress has been slow. From instinct-driven, trial and error designs, we are now able to model much of the performance of a transducer on a computer, before building a prototype. We can explore different materials, different shapes and sizes, and different motor configurations, to optimize the frequency response, sensitivity, etc. Once we have a prototype, we use detailed acoustical measurements to judge quite accurately how it is likely to sound before we even play a note of music through it. If it turns out to be too “colorful”, we can do

scanning-laser-vibrometer measurements to determine what, precisely, is causing the timbral inaccuracy, so that it can be designed out of the next iteration.

Once we have suitable woofers, midranges and tweeters, we can integrate them into an appropriate enclosure, using more computer aids to design crossovers that are tailor made for just these drivers in just this enclosure. While all of this technical tweaking is being done, we continue with more acoustical measurements, which tell us whether we are on course to hit our target. This assumes that there is a target. There is, and this is one of the main differences between loudspeakers today, and those of the “colorful” era.

When it is ready, we listen. And what do we listen for? Color. Coloration in the sound of a loudspeaker means that it is adding its own personality to all of the sounds it reproduces. Color should be in the music, in the voices and instruments, in the virtuosity with which they are played – it should not be in the loudspeakers! Some colors are more annoying than others; some even can be pleasant with certain recordings. However, if one hopes to listen to a wide variety of musical and other sounds, it is unlikely that any added color will be acceptable with all of them.

When it has been put to the test, in rigorous controlled listening tests, we find routinely that listeners prefer the loudspeakers with the least coloration. When making notes of their reactions to loudspeaker sounds, we find adjectives like, boomy, tizzy, honky, nasal, thin, fat, resonant, shrill, and boxy, along with some more poetic descriptors, like “chocolatey bass”. The highest ratings tend to go to the loudspeakers with the shortest list of adjectives. Neutral, transparent, loudspeakers are good. So, how do we get them?

Before we get into the details of loudspeaker design, we need to stand back a good listening distance, and examine what we hear. We listen in rooms that have mostly reflective walls, floor and ceilings. Loudspeakers radiate sound in all directions and, after reflection, all of these sounds eventually arrive at the ears. Let us look at what happens in a typical room. The loudspeaker used in this example was one that listeners said had some timbral problems, and it is a good way to demonstrate some important fundamentals of loudspeaker design. The data are from one of my early AES papers (1).

Figure 1 shows some measurements, indicating that the design directive for this loudspeaker system was focused on the on-axis performance. In this, the engineers clearly succeeded, it is very smooth and flat, but equally clearly they ignored what happened off-axis.

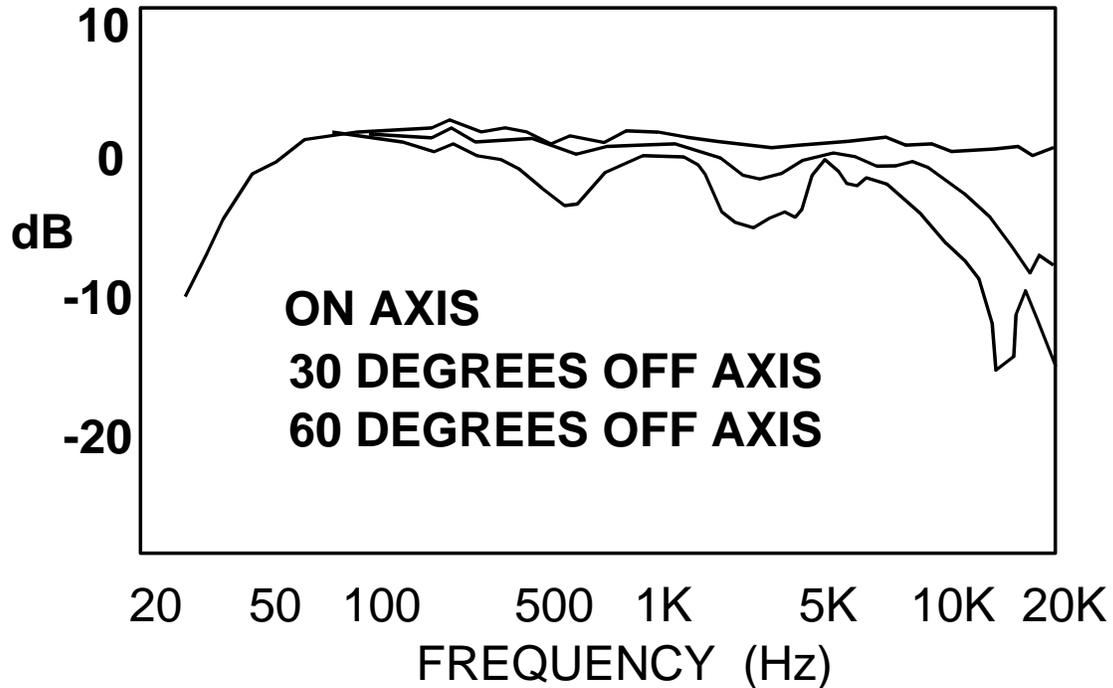


Figure 1. Anechoic frequency response measurements of a loudspeaker showing (top to bottom) a very smooth and flat on-axis response, and increasingly non-flat behavior at increasing angles off axis.

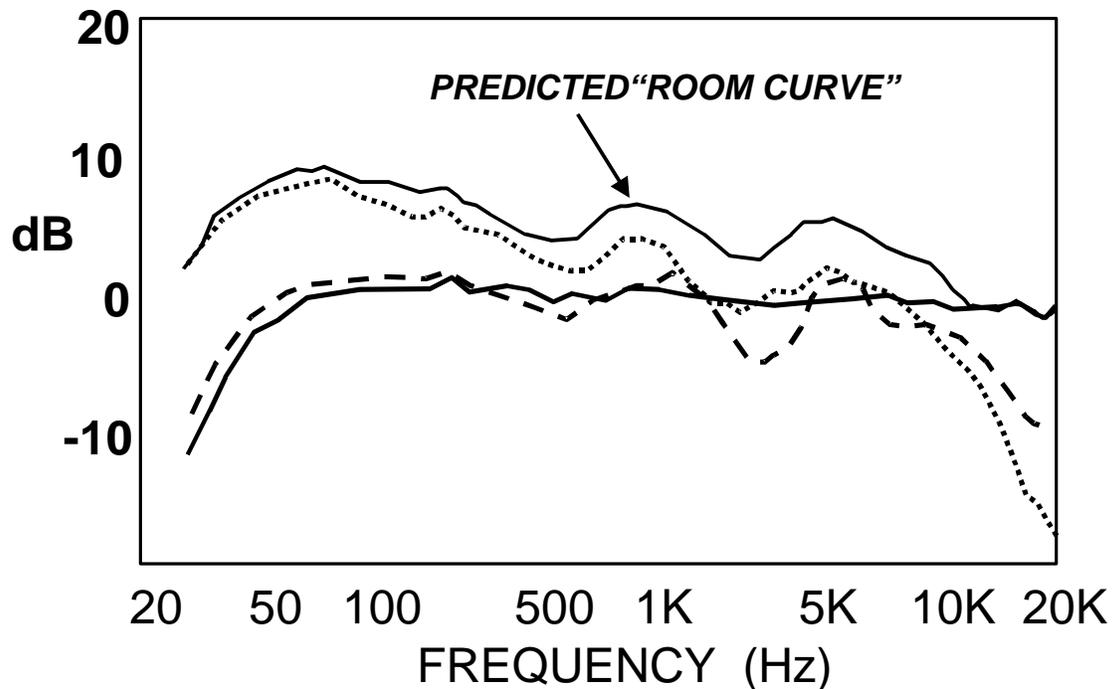


Figure 2. The sequence of sounds arriving at the listening area in a room. The first sound to arrive is the direct sound, represented by the on-axis frequency response (the solid line around 0 dB). The second sounds to arrive are the early reflections from the floor, ceiling and side walls. These are represented by an energy sum of anechoic

measurements made at the appropriate off-axis angles, corrected for propagation losses (the dashed line). The final collection of sounds is the reverberant sounds that arrive after multiple reflections from the room boundaries and contents. These are represented by a measurement of the total sound power radiated from the loudspeaker, corrected for the frequency-dependent absorption in the room (the dotted line). The solid curve plotted over the top of these three, is a sum of them, a prediction of what might be measured in a real room.

Figure 2 shows a synthesis, from anechoic measurements on the loudspeaker, of the sounds arriving at a listener's ears in a typical room. The sounds have been divided into three categories, direct, early reflections, and reverberant sounds. For this example, the room boundaries were assumed to be relatively reflective. It can be seen that, at low frequencies the dominant factor is the reverberant sound (sound power). The direct (on-axis) sound is not really a factor. At the other extreme, at the highest frequencies, what we hear is absolutely dominated by the direct (on-axis) sound. In between, for a frequency range covering the most important instrumental and vocal sounds – from a few hundred Hz to a few thousand Hz – everything matters. All three components are clearly influential. So, if we are trying to characterize the sound of a loudspeaker in a room, by using anechoic measurements of frequency response, it is quite clear that we must measure absolutely everything! Reducing the description of frequency response to a single curve is overly simplistic. An on-axis frequency response tells only part of the story, as does a measurement of total sound power.

This is interesting, but it is still theory. What actually happens in a room? Figure 3 shows that the real world is more complex. At low frequencies, the room resonances and adjacent-boundary reflections dominate the communication of bass sounds from the loudspeaker to the listener, and different loudspeaker positions yield very different impressions of bass – from punchy “rock ‘n roll” bass (solid curve), through wimpy, inadequate bass (dotted curve), to quite good (dashed curve). Although not shown, moving the listener can achieve similar differences in bass performance. Our prediction did not consider this profound effect, although there are ways to do it for specific rooms. The conclusion, though, is clear – at low frequencies, position is everything. Even with “perfect” woofers, the room and the arrangement of sources and listeners within it, determine the quality of the bass we hear.

At middle and high frequencies, the result is strikingly different. Position matters much less, and our calculated room curve fits almost perfectly. From “sterile” anechoic data on the loudspeaker, we have been able to describe what happens in a room. Not bad.

Summing this up, it can be stated that there are two distinct domains in a listening room. Below about 400-500 Hz the room dominates what we hear and, above that division, the loudspeaker dominates. In the case of this loudspeaker, it is clear that listeners' complaints about colored midrange sound were caused by the way the loudspeaker was designed. Focusing on the direct/on-axis sound ignored the fact that what we hear in a room is very much influenced by sounds radiated in other directions.

It is worth noting here that equalization is not a solution for this kind of problem. Changing the shape of the room curve destroys the only good thing the loudspeaker had, its on-axis response. This loudspeaker can be made to sound different, possibly slightly better, but it can never equal the sound of a loudspeaker that is properly designed to begin with. To avoid coloration in its interaction with the listening room a loudspeaker must be well behaved both on and off axis. If it is designed to have relatively constant directivity, as a function of frequency, then it is possible for all three categories of sound, direct, early reflected, and reverberant, to exhibit similarly accurate timbral signatures. That is what makes good sound.

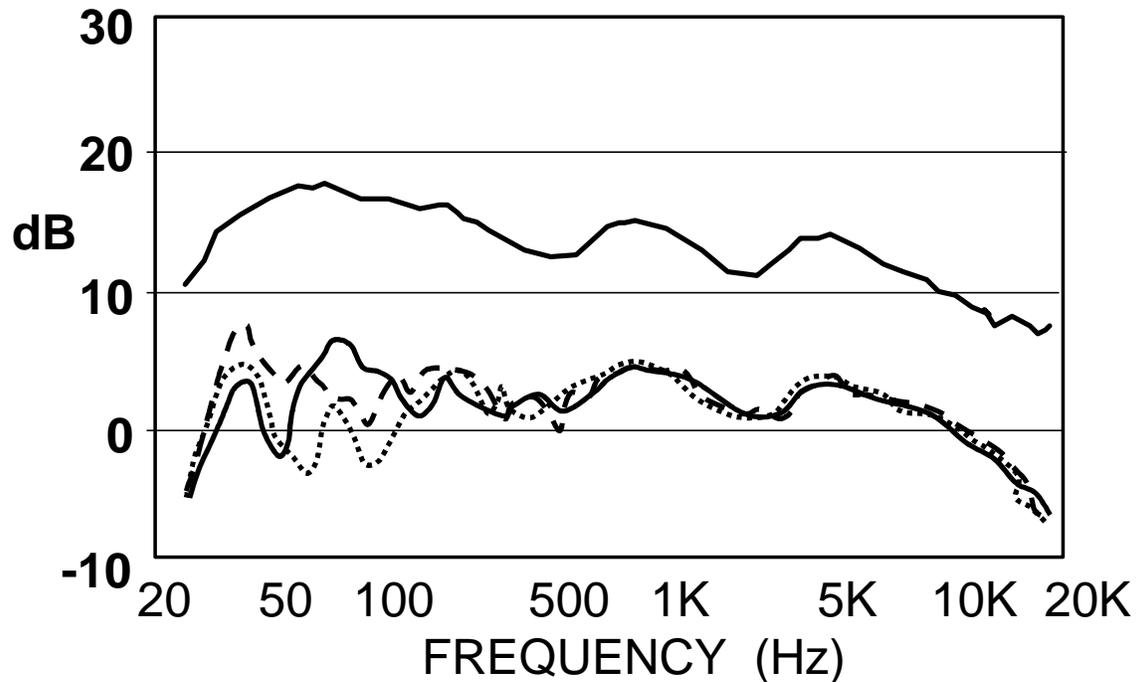


Figure 3. Measurements of the example loudspeaker in a room. The loudspeaker was placed in three realistic left/right channel locations and, for each one, measurements were made at four possible listening locations, all within a two-foot radius. Each of the bottom curves is an energy average of the four measurements. The top curve is the predicted room curve from Figure 2, which, for clarity, has been shifted vertically by 10 dB.

Now that we understand the “big picture” of loudspeaker / room interaction, let us get closer to the device itself. Loudspeaker systems use transducers. A transducer is more than a fancy word for a loudspeaker driver, such as a woofer. It describes a device that converts energy from one form to another – in this case, from electrical to acoustical. The electrical signal is, or should be, a waveform record of the sounds created by the artists and recording engineers who originated the “art”. The task of the transducer is to create an exact acoustical analogue of that waveform, and thereby preserve the integrity of the art.

One transducer cannot do the entire task, however. There are two reasons. First, music extends from powerful low frequencies to delicate high frequencies. It may be expecting too much for the same transducer that reproduces a chest-thumping kick drum or movie explosion, to reproduce the delicacy of violin overtones. That is why we have woofers and tweeters. Big diaphragms move a lot of air, to shake us and the room around us. While it is physically possible to make large diaphragms move at high frequencies, we may not want to try. There are two reasons.

First, from the loudspeaker / room investigation we learned that it is a good idea to strive for constant directivity over most of the frequency range. As sounds increase in frequency (go up in pitch), they become physically smaller - their wavelengths become shorter. Therefore, in order to maintain a uniform dispersion of the sound, it is necessary to progressively reduce the size of the radiating diaphragms as we go up in frequency. How many different sizes of transducers are used in a loudspeaker system is, in part, determined by the requirement for constant directivity. Every transducer becomes progressively more directional, favoring the forward direction, as it goes up in frequency. A 12-inch two-way design is not ideal because, at the crossover frequency to a typical 1-inch tweeter, 2 to 3 kHz, the woofer would be very directional, and the tweeter turns on with wide dispersion, because of its small size. The discontinuity in directivity is what causes the off-axis problems seen in Figure 1, where there are two such examples: when the woofer transitions to the midrange near 500 Hz and when the midrange transitions to the tweeter around 2-3 kHz.

The second reason not to try to do it all with one transducer, is that the diaphragms flex and resonate at certain frequencies. Ideally, we want a diaphragm that is perfectly rigid, moving as a piston at all frequencies. Large woofer diaphragms want to resonate at frequencies right in the really important part of the frequency range, where they would add serious boxy, honky and nasal colorations to voices and instruments. Consequently, we have a second reason to cross-over, or transition, to a transducer with a smaller diaphragm for the upper-bass and middle frequencies. At still higher frequencies, the same thing happens to the mid-range driver. Its diaphragm breaks up, and resonates at frequencies that cause high frequencies to sound harsh, brittle and strident. So we transition to a tweeter which, with good design, behaves itself to frequencies above our audible limit.

The first transducer diaphragms, conical in shape, were made of paper, and for many applications, paper is still widely used, albeit in a much-evolved form. However, paper is a “witches brew” of ingredients, extremely difficult to control in production, and subject to variations due to temperature, humidity, and fatigue. Consequently, engineers have sought the perfect “cookie cutter” material: stiff, light, strong, well damped, impervious to the elements, attractive, inexpensive and easy to duplicate, perfectly, in mass production.

Over the years, we have seen cones and domes made of various plastics, including the extremely popular polypropylene, fiber and fabric (e.g. silk, glass, Kevlar and carbon) reinforced composites, metals like aluminum, titanium, and beryllium, and laminates of

various substances. All of them work, some of them very well, but they represent two different approaches to reducing audible colorations due to resonances in the diaphragms.

The first approach is to allow the diaphragm to resonate within its intended frequency band, but to rely on mechanical damping in the material to reduce its Q, or inclination to “ring”. The non-metallic materials generally fall into this category, and we all have spent many happy hours listening through them and to them. In the old days, this was called “controlled breakup”, and the idea was that by letting large cones break up, that they would have improved dispersion at higher frequencies, avoiding the need for a midrange driver. Well, most of them didn’t work very well, and multi-way systems have come to prevail where costs permit. We still use flexible materials in our systems but, through materials science, we have been able to muster much better control of the ways and degrees in which they resonate.

The second approach is to use very stiff diaphragm materials that move the resonances higher in frequency, high enough to be out of the frequency band over which we want to use the device. The traditional problem with metal cones and domes is that when they do break up – and eventually they do – they ring like bells. These are materials with low mechanical losses, and they characteristically have very high-Q, resonances, giving them their inimitably “metallic” timbres. However, there has been progress in this area as well, with some current designs being essentially free from resonances within their useful bandwidths.

The challenge is to push the resonances below the threshold of audibility. If they cannot be heard, they do not exist, and that is equally true for both the approaches described above. It is quite a challenge, since we humans are very, very, good at hearing resonances. The reason is obvious, when you think of it. It is because all of the sounds we are interested in listening to, voices and musical instruments, are composites of many resonances. It is the details of the resonances, their strengths and Q’s, that allow us to recognize different voices when somebody says “hello” on the telephone. It is differences in the collections of resonances that differentiate musical instruments playing the same notes; a violin note and a cello note have the same pitch, but very different timbres. A Stradivarius has a different timbre than an inexpensive student violin. We humans are designed to hear resonances, timbral differences, and in loudspeakers it should come as no surprise that listeners’ main complaints are about unwanted colorations caused by resonances.

As an example of just how fussy we humans are, Figure 5 shows frequency-response deviations caused by resonances of three different Q’s, when they are just at the threshold of detection when listening to orchestral music – one of the more revealing signals (2). Although they look very different to the eyes, they are equally audible to the ears. This is why psychoacoustic investigations are so very important. They help us to understand the often non-linear relationships that exist between what we hear and what we measure.

An important side note to Figure 5, is that it means that the popular verbal description of frequency response: 20 to 20,000 Hz \pm 3 dB, is meaningless unless it is accompanied by a

graph. One needs to know whether the deviation is a broad bump (serious) or a narrow spike (possibly innocuous). A tolerance of ± 1 dB would get my attention though, even without a graph. A further observation from this Figure is that measurements must be done with enough resolution, in the frequency domain, to reveal these audible resonances. Many measurements one sees in manufacturers' literature and magazines lack the resolution to show these aberrations at all, or with enough accuracy to enable one to interpret them realistically.

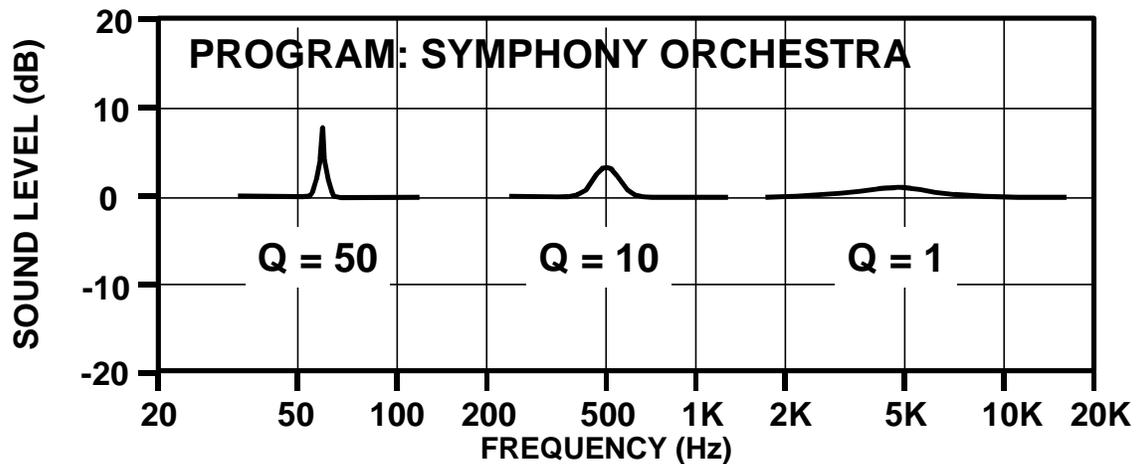


Figure 5. Frequency response deviations caused by resonances with different Q 's, when they have been adjusted to the levels at which they are just detectable when using symphonic music as a test signal. NOTE: Ignore the positions the resonances occupy on the frequency scale; they are about equally detectable at all middle and high frequencies. From reference 2.

So, having started this article with the crude beginnings of our industry, where are we now? Have we reached the nirvana of perfect transparency?

In order for that to be audibly true, we would need to trust that every recording was recorded without timbral distortions. That is patently not the case, and one of the main reasons is that studio monitor loudspeakers and rooms are at least as variable in quality as consumer loudspeakers and rooms. Monitors are used to judge the selection of microphones, and electronic processing, if any, that capture and create the recorded sounds. It is common for recordings to exhibit colorations, added during the process, that mirror the circumstances in the monitoring environment. If the monitor loudspeakers are too bright, the recordings tend to be dull, and so on. Until we can be assured that the monitoring audio system was as neutral as our home system, we will never be sure whether what we hear was done deliberately, as part of the “art”. Frustrating, isn't it?

In the meantime, loudspeaker engineers are gradually mastering the science of designing and building loudspeaker transducers and systems that approach the ideal. Figure 6 shows two bass-midrange cone transducers, in an enclosure, measured on axis. For the crossover frequency intended for this driver, about 2.5 to 3 kHz, it is evident that the

aluminum cone does a remarkably good job. However, when it resonates, it does so with typical metallic vigor – a high-amplitude, high-Q resonance at about 4.5 kHz, close enough to the crossover frequency to still be an audible threat. Also shown is a curve for a new ceramic / aluminum laminated material. This cone exhibits a higher resonant frequency, consistent with the higher stiffness of this material, but the resonance is better damped (lower Q), and has reduced amplitude.

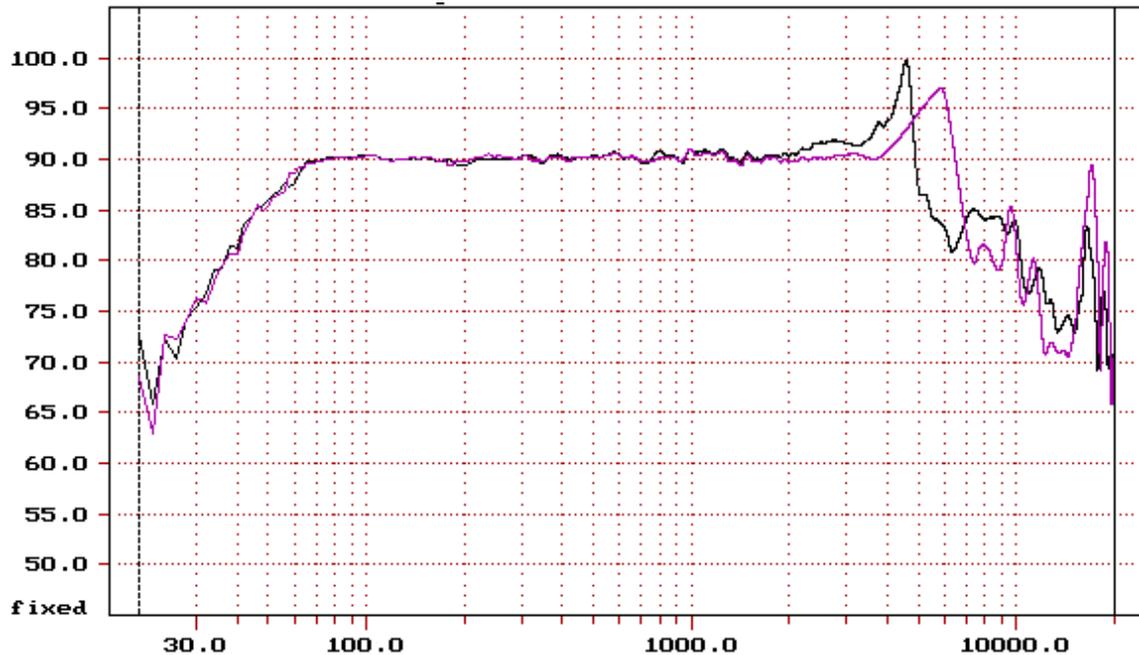


Figure 6. On-axis measurements of a 6.5 inch mid-bass driver with a cone made of aluminum (top curve) and of a ceramic / aluminum laminate (bottom curve).

When these drivers are combined with a crossover network, the result is shown in Figure 7. The high-Q resonance of the aluminum cone is still evident in the slope of the crossover, and it therefore still presents a threat of coloration. Additional filters in the crossover could help, but a better solution is the laminated cone, where the resonance is just barely visible in the lower portion of the attenuation slope – the resonance has been moved up in frequency and about 10 dB further down in amplitude. We know that this is below the threshold of audibility. This is a measured performance that comes dangerously close to the theoretical ideal.

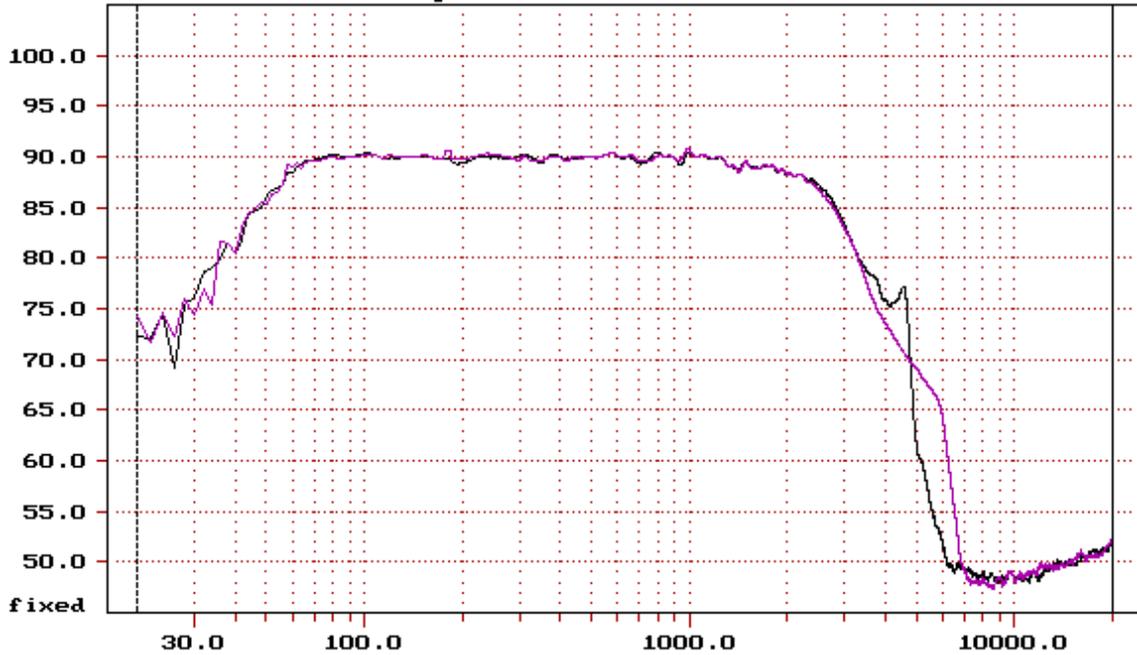


Figure 7. The transducers shown in Figure 6, after the crossover network is added.

In addition to the traditional frequency responses, only a sample of which have been shown here, transducer engineers confirm the performance of their designs by looking at the how the surface of the cone moves at different frequencies. The device that enables this is called a scanning laser vibrometer. Figure 8 shows that all points on the surface of the cone of the 6.5-inch driver are moving in unison – like a piston – at a frequency of 3.5 kHz, just below the first cone bending resonance frequency. It behaves in this way at all lower frequencies as well.

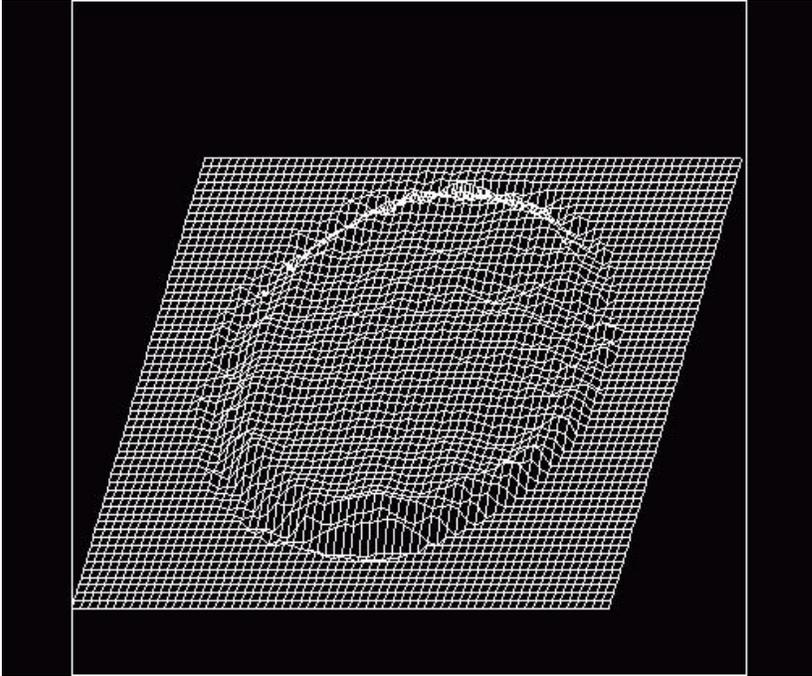


Figure 8. A scanning laser vibrometer measurement of cone movement for a 6.5-inch ceramic/aluminum laminated cone, operating at 3.5 kHz.

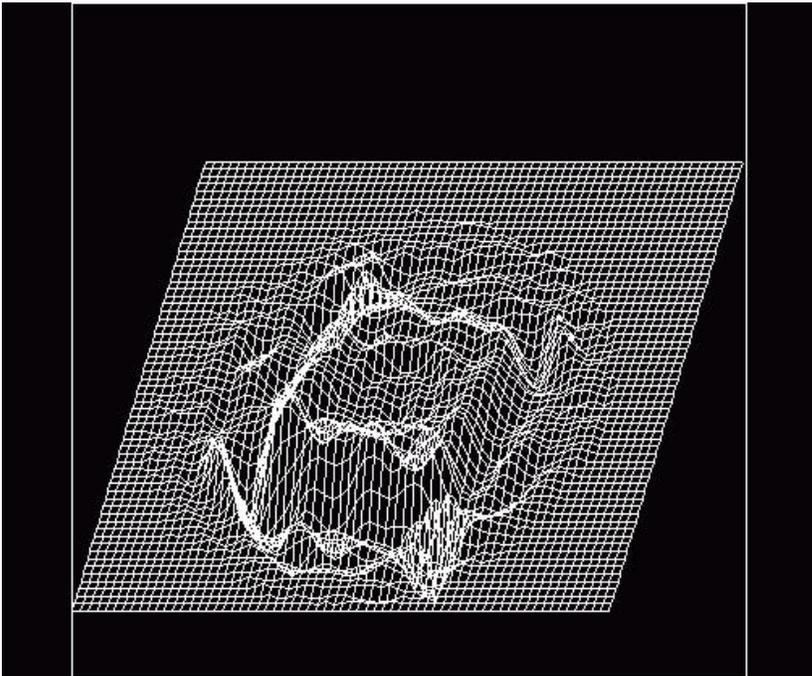


Figure 9. A scanning laser vibrometer measurement of cone movement for a 6.5-inch fabric-reinforced composite cone, operating at 3.5 kHz.

Figure 9 shows a comparable measurement on a cone made of fabric-reinforced composite material that, because of its much lower mechanical stiffness, is seriously into resonance at the same frequency. Different parts of the cone are going in different directions at the same time – in the trade jargon, it is in “break-up”. Such cones can still sound good, but only if the material is designed with sufficient mechanical losses to damp and attenuate the resonances to below audibility.

The point of these examples is to illustrate just how far we have come from the “old days” of loudspeaker design. Here is a transducer that was first designed in a computer model, which was used to develop an engineered cone material, all of which resulted in a mid-bass driver that is absolutely free from audible resonances within its intended range of operation from about 40Hz to 3 kHz. The transducer engineer has done a fine job, but the system engineer’s job is not yet done. He must ensure that the crossovers maintain the integrity of the sound output in the transitional regions between woofer and midrange, and midrange and tweeter, and he must be careful not to let enclosure resonances and diffractions undo the excellent performance of the drivers.

The remaining large issue determining how a loudspeaker sounds in a room, is the matter of room resonances. But that is a separate story.

REFERENCES

1. F.E. Toole, “Loudspeaker Measurements and Their Relationship to Listener Preferences”, J. Audio Eng. Soc., vol. 34, pt.1 pp.227-235 (1986 April), pt. 2, pp. 323-348 (1986 May).
2. F.E. Toole and S.E. Olive, “The Modification of Timbre by Resonances: Perception and Measurement”, J. Audio Eng. Soc., vol. 36, pp. 122-142 (1988 March).

MAXIMIZING LOUDSPEAKER PERFORMANCE IN ROOMS

Part Two: Acoustical Design and Equalization

by
Floyd E. Toole, Ph.D.
Vice President Acoustical Engineering, Harman International Industries, Inc.

This overview is intended to set a context within which readers can apply the more detailed technical information in the paper: “Loudspeakers and Rooms - Working Together”. Some people come to this topic thinking that, as is some other things, that there must be a simple way to do acoustical design, a kind of cookbook, that anybody can understand. I wish that were so, because it would simplify all of our lives. As it is, achieving truly good sound in rooms requires knowledge of how sound behaves in rooms, and some effort – or more than a little bit of luck.

Multichannel audio has become a reality in home theaters, and it is rapidly becoming a presence in music and games. Much of the technology in audio, that we know today, has evolved during the nearly fifty years we have spent with two-channel stereo. For example, designs for loudspeakers and rooms have been developed that specifically serve that medium. With only two channels, many listeners found benefits in spraying the sound around the room, using multidirectional loudspeakers, and in tailoring the reflected sound field with strategically placed absorbing, reflecting and diffusing surfaces. Stereo remains an area ripe for experimentation, mainly because only two loudspeakers restrict what is possible in creating realistic senses of direction and space. Stereo needs help, and in experimenting with various enhancements, we often create something that is different from what was intended, although it may be thoroughly engaging and entertaining.

Multichannel audio changes the rules significantly. There are now multiple real sound sources surrounding the listeners. The front channel loudspeakers include a center, thus theoretically eliminating the need for a “stereo seat” - multichannel audio should not be an antisocial listening experience; it is intended to be shared. As things stand, I am sad to report that several of the new multichannel music recordings avoid using the center channel, and the phantom center image forces listeners back to the stereo seat. I hope that this remnant of the stereo era will pass as recording engineers learn how to use the center channel tastefully, perhaps with some new production tools. Of course, customers must also make a commitment, and use center channel loudspeakers that are up to the task of doing most of the important work in a multichannel system.

The surround side/rear loudspeakers have been matters for experimentation, especially for the reproduction of movie sound tracks, but the appearance of multichannel music changes the focus somewhat. In the days of Dolby Surround/ProLogic, film sound was well served by multidirectional (including dipole) surround loudspeakers. However, listeners are finding that the digital discrete 5.1-channel systems seem to work well with five identical loudspeakers, and this certainly is the early trend in multichannel music. The key factor here is that the perceptions of direction, space, depth and so on, are really very much in the hands of the artists and recording engineers. At the reproduction end of the chain, the multichannel audio system should be a faithful delivery system for their creations. There is less need for the customer to be creative, in order to hear good and exciting sounds.

With this little perspective behind us, let us dispense with a popular question: “aren’t some loudspeakers better for movies, and some better for music?” This is really the wrong question. The real question of the day is: “aren’t some loudspeakers better for stereo, and some better for multichannel audio?” There are good loudspeakers, and not-so-good loudspeakers of all design configurations. However, it is reasonable to assert that multidirectional designs, intended to enhance the spatial illusions in two-channel stereo, may not be ideal for the main loudspeakers in a delivery system in which the spatial illusions are supposed to be under the control of the artists and engineers. As for the loudspeakers designed for multichannel applications, it would be a terrible embarrassment to any competent loudspeaker designer if his products were not equally good sounding in multichannel performances of James Taylor or Fleetwood Mac in concert, Strauss waltzes, the 1812 overture, Armageddon and Terminator II. The only additional consideration in movies, is the need for higher sound levels, especially of deep bass, if one’s taste runs to blockbuster adventure films with their spectacularly over-the-top pyrotechnics.

GETTING GOOD SOUND IN A ROOM. HOW IS IT DONE?

The science of room acoustics has mainly developed in the context of live performers in concert and recital halls. Relatively little scientific effort has been put into understanding sound in small rooms, especially as it relates to sound reproduction. The irony is that far more music is listened to at home, than in concert halls. Still, there has been significant progress, and we are beginning to understand some of the things we can do to ensure decent sound quality in the semi-infinite range of room sizes, shapes, arrangements, and furnishing variations that exist. It sounds as though it might be difficult. Well, it is not nearly as complicated as “rocket science”, but neither is it a totally straightforward “cookbook” exercise. You will have to do some work, and think a bit.

Step One: Start with a good room

That is, of course, if there is a choice. Most often we must work within an existing space, or one that has been designed with things other than acoustics in mind. There are notions that some room dimensional ratios – length to width to height – offer special advantages.

Personally, I have not found this to be so, and I think I know why. It is because the theories and calculations leading to these preferred ratios assume several things that are not true to our realities.

- First, it is assumed that the rooms are perfectly rectangular, with perfectly reflecting, perfectly smooth, flat walls. This rarely happens, and if it did, we would probably want to do something to change it. Such rooms are unpleasant listening spaces.
- Second, it is assumed that all calculable room resonances are equally important. They are not. In terms of their impact on audible characteristics, it is abundantly evident that, in most rooms, the axial modes have the strongest “voices”, with tangential and oblique following behind. I have encountered a few rooms, with relatively massive, stiff walls, where one or two low-order tangential modes are audible problems, but that is all.
- Third, it is assumed that all of the calculated room resonances are equally energized by the sound sources, and are equally audible to listeners. This could only be true if we only had one sound source, on the floor in a corner, and if we positioned our head at another three-surface intersection. Such a notion is preposterous! In reality, there may be two or more sources of low-frequency sound. Two physically separated woofers, even if they both are in corners, do not energize all of the room modes equally, or at all. If they are not in corners, the modal excitation can be very selective indeed. Likewise, listeners do not stick their heads in corners. Out in the middle of a room, the coupling to room modes is extremely selective, and that is one of the biggest problems we have to deal with.

So, why did this whole business of special room ratios get started? Actually, it began decades ago, very scientifically, with serious minds trying to optimize the performance of acoustical reverberation chambers, where it was intended to conduct precise measurements of sound power. From there it got picked up and elaborated unrealistically to include rooms, like ours for sound reproduction, where those ideas simply do not apply.

Now, this does not mean that room ratios are irrelevant. It is a good idea to avoid a perfect cube, rectangles with simple (whole number) dimensional ratios, and long corridors. Beyond that I believe that, if you know what you are doing, it is possible to create excellent sound in rooms that are in gross violation of the “rules”, and just as it is possible for truly mediocre sound to exist in supposedly “good” rooms.

In truth, the most problematic rooms that I have encountered, were ones that came too close to the first of the “ideals” listed above. The room boundaries were very hard, very dense, and very flat. The result was that all of the room modes were extremely powerful, high-Q, and very “resonant”. Consequently, the resonant peaks were very high, the cancellation dips very deep, and the “booms” went on forever.

In order to be good, a room must have some low-frequency sound absorption, and if this is not to be found in the room boundaries themselves, then it must be added. A few inches of resistive absorbing material, such as fiberglass or acoustic foam, has no effect at low bass frequencies. Low-frequency absorption is most effectively done with panel, or

membrane, absorbers. When large surfaces, including the room boundaries, floor, walls and ceiling, move in response to powerful bass sounds, they are behaving as membranes and they are absorbing sound energy. This absorbed sound energy cannot contribute to room resonances, and as a consequence, the resonances are weakened. This is a good thing. One can buy or build membrane absorbers, although getting them to be effective at very low frequencies can be a challenge. Many of the devices in the marketplace are not very effective below about 100 Hz, just where things get interesting. Be sure to check that the absorption coefficient is high in the frequency range where your problems are. If possible, one can anticipate the problem and build the interior structure of the room so that the room boundaries are somewhat flexible. It turns out that one layer of gypsum board on single wooden studs is not a bad compromise – and it is inexpensive. A layer of acoustic board under the gypsum board can add mechanical damping, without adding much mass or stiffness. Some people have promoted varying the stud spacing to “detune” the mechanical wall resonances. By doing this, you can certainly anticipate the wrath of your builder, for making his job much more difficult. Much the same result can be achieved by occasionally doubling the studs, and by ensuring that the wall surfaces are not perfect flat slabs – a good thing from the point of view of diffusion.

This done, another very important benefit will have been realized, and that is to improve the uniformity of the bass sound over a large listening area. By reducing the “Q” of the room resonances, the pressure peaks are lowered, and the pressure nulls are not as deep, making good bass possible at more than a few locations. Multichannel audio is to be shared.

Step Two: Start with good loudspeakers – ones that are “room friendly”.

What we hear in a room is controlled by different factors at different frequencies. At low frequencies it is the room that dominates, but at middle and high frequencies, it is the loudspeaker itself – its frequency response and directivity – that dominate sound quality. Little-to-nothing can be done, with an equalizer in a room, to fix a loudspeaker that is fundamentally poorly designed. Here the solution is to start with a loudspeaker that is designed to be “room friendly”. It may come as a surprise to you, but not all are.

An example of this that many of you may remember is the fashion, a few years ago, of building rooms that were acoustically live at the listener end, and acoustically dead at the loudspeaker end. The inspiration for this appears to have been the need to improve the sound of a, then popular, studio monitor loudspeaker that misbehaved dreadfully in its off-axis response. The only way to deliver good sound quality was to absorb the sounds that would normally have reflected off the floor, walls and ceiling. This is the definition of a loudspeaker that is hostile to ordinary rooms, and that required a major overhaul of the listening space to make it acceptable. This ludicrous exercise managed to survive for some time, even in the normally-rational world of professional audio. It is totally ridiculous for home audio. Not surprisingly, when subjectively evaluated as a “hi-fi” loudspeaker, in a normally reflective room, this loudspeaker was enthusiastically disliked.

The real solution, for professionals as well as consumers, is loudspeakers that deliver similarly good timbral accuracy in the direct, early reflected and reverberant sound fields. This can be described as a loudspeaker with a flattish, smooth, axial frequency response, with constant directivity (which together result in flattish, smooth, sound power). Then it becomes an option, whether the room is acoustically damped, or not. If reflected sounds are absorbed, the listener is placed in a predominantly direct sound field, making the experience more intimate, and the imaging tighter and more precise. If the reflections are allowed to add their complexity, the overall illusion is altogether more spacious and open, to many listeners, more realistic. In part, this is a matter of taste. In either case, a room-friendly loudspeaker will yield timbral accuracy. So, at middle and high frequencies, the proper solution to getting good sound quality, is to choose good loudspeakers to begin with.

Step Three: Improving Bass Performance – working with standing waves

At low frequencies, the situation is very different, and the room, and the arrangement of sources and listeners within it, dominates the quality of bass. Of course, the woofer must be capable of the necessary quantity of sound, with low distortion, over the necessary frequency range. Taking control of this means that it is necessary to get a bit technical, in order to understand how the subwoofers couple their energy into the room resonances, and what listeners hear. There are several forms of computer aids that make life easier, and there are some simple manual methods that can take us most of the way.

Measurements of what is happening in the room are absolutely necessary if you are to be successful. However, they must be of the right kind – more detailed than is available from the normal 1/3-octave real-time analyzers. Use a high-resolution measurement system, such as MLSSA, SMAART, TEF, or even old fashioned swept or stepped tones, set up to give at least 1/10-octave resolution (2Hz resolution at 20 Hz). Measure what is reaching the listening positions when all subwoofers are simultaneously active.

If the room is a simple rectangle in shape, the resonant modes are easy to calculate, at least the axial modes, which are usually the biggest problems. Calculating the frequencies at which the resonances occur is just a start; it is then necessary to determine where in the patterns of pressure peaks and nulls (the standing waves) it is best to locate the subwoofers, and where to position the listeners. It will be found that maximizing the pleasure and minimizing the pain involves some tradeoffs. Figure 1 shows a printout of a useful program that illustrates the sound pressure distribution along each room axis, allowing you to choose the listener locations to avoid the worst peaks and dips. Subwoofers need to be in high-pressure regions, preferably against the wall, or better still in a corner in order to excite the room modes. The PC program that does this is available, at no charge, from the author: ftoole@harman.com.

If measurements indicate that there is too much energy at a resonant frequency, try to move the listener towards a null in that particular standing wave pattern; too little, means moving out of a null, towards a peak. By these trial and error methods it is often possible to alleviate many of the problems, smoothing and flattening the frequency response.

Standing Waves					
Enter dimensions, in feet	1st	2nd	3rd	4th	
Height: 8.00ft	70Hz	141Hz	211Hz	281Hz	
Width: 18.30ft	31Hz	61Hz	92Hz	123Hz	
Length: 21.50ft	26Hz	52Hz	79Hz	105Hz	
Room volume = 3148CU ft					

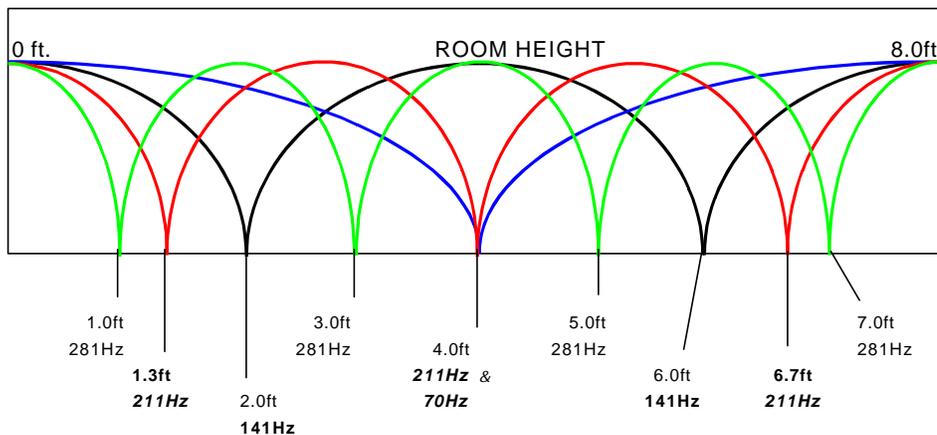
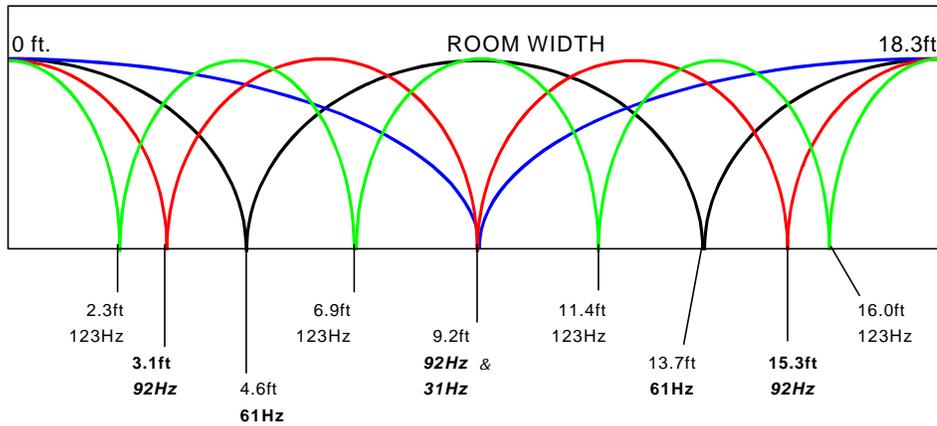
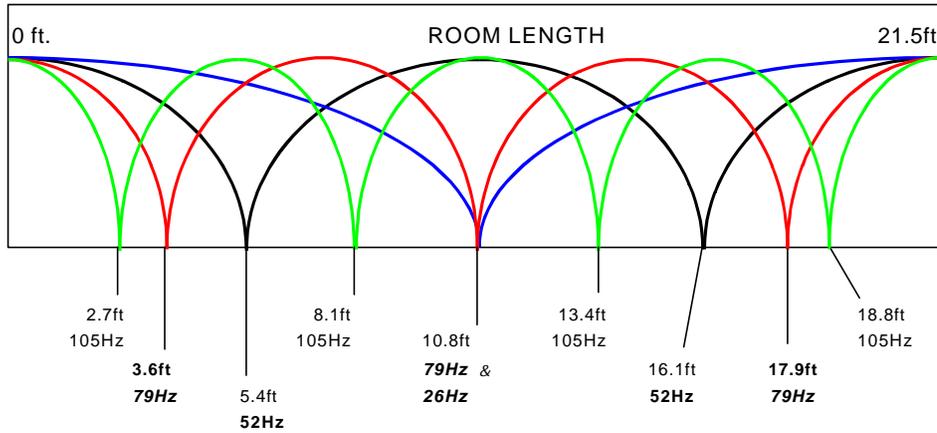


Figure 1 The output of a program that calculates the axial modes of a room, and plots the pressure as a function of distance along each of the principal axes.

If the room is not rectangular in shape, or if there are large openings in one or more of the walls, predictive calculations will not work well, or at all. In this case, one must revert to measurements, and trial and error repositioning of subwoofers and listeners. This is not a good situation to be in. Non-rectangular rooms do not eliminate room resonances; they just make it impossible to calculate them simply.

With the best rooms, and intentions, perfection may be elusive. Given the practical constraints of real environments, and the limits on listener positions, viewing distances and angles, acoustical manipulations may not be enough to eliminate all room resonance problems. In fact, in my experience, it is rarely enough.

Step Four: Improving Bass Performance – equalization does work!

It is in these situations, when you have exhausted the acoustical possibilities, that equalization of the right kind can be very helpful. However, it must be done intelligently, since there are some things that equalization can correct very well, and other things that it is a mistake to try to fix.

Not everybody agrees with equalization, accusing it of introducing “phase shift”, and other nasties. Well, there is no doubt that equalization has acquired a bad reputation over the years, but from the perspective of what we know now, it has been absolutely deserved. There are four principal reasons:

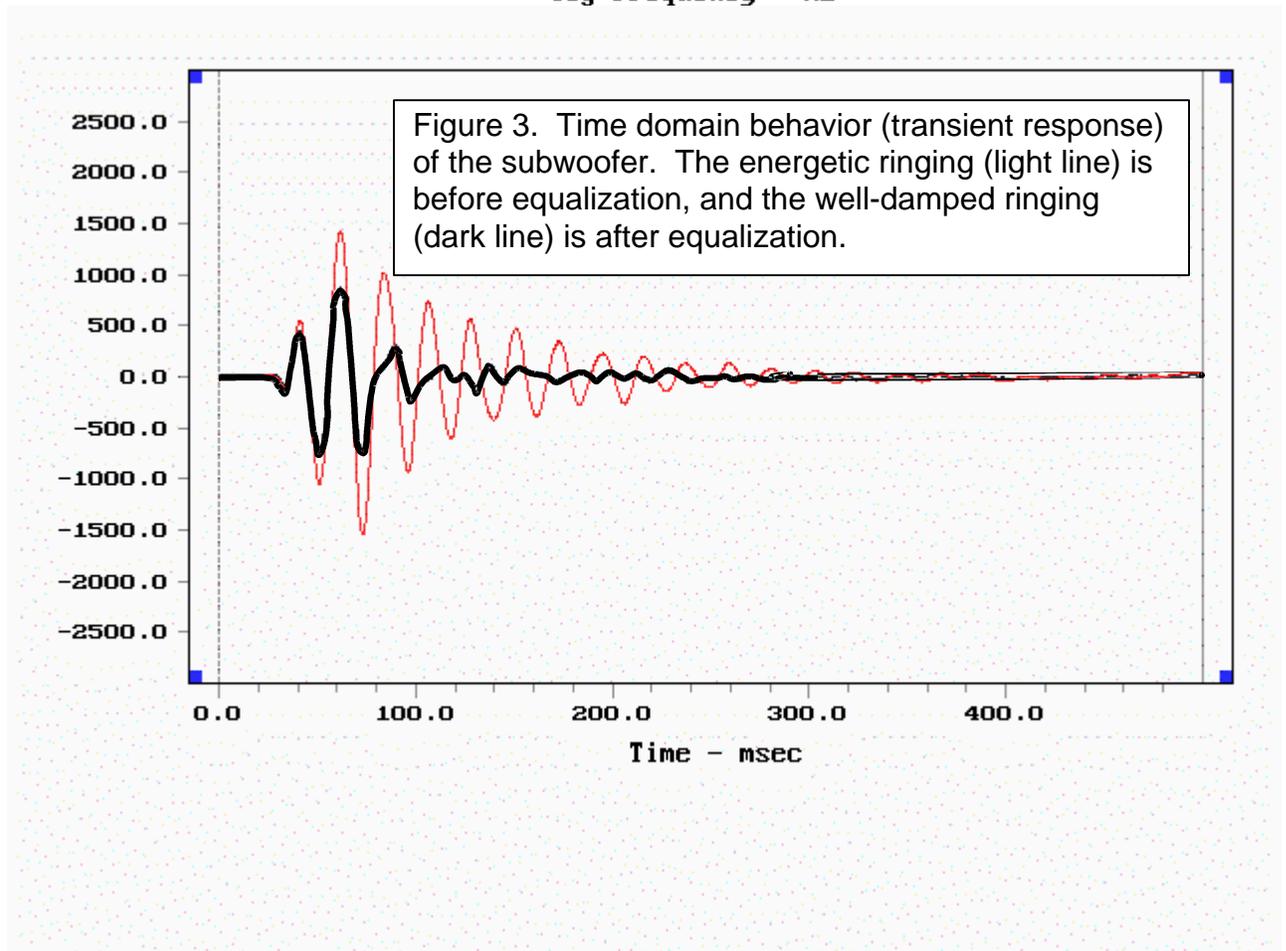
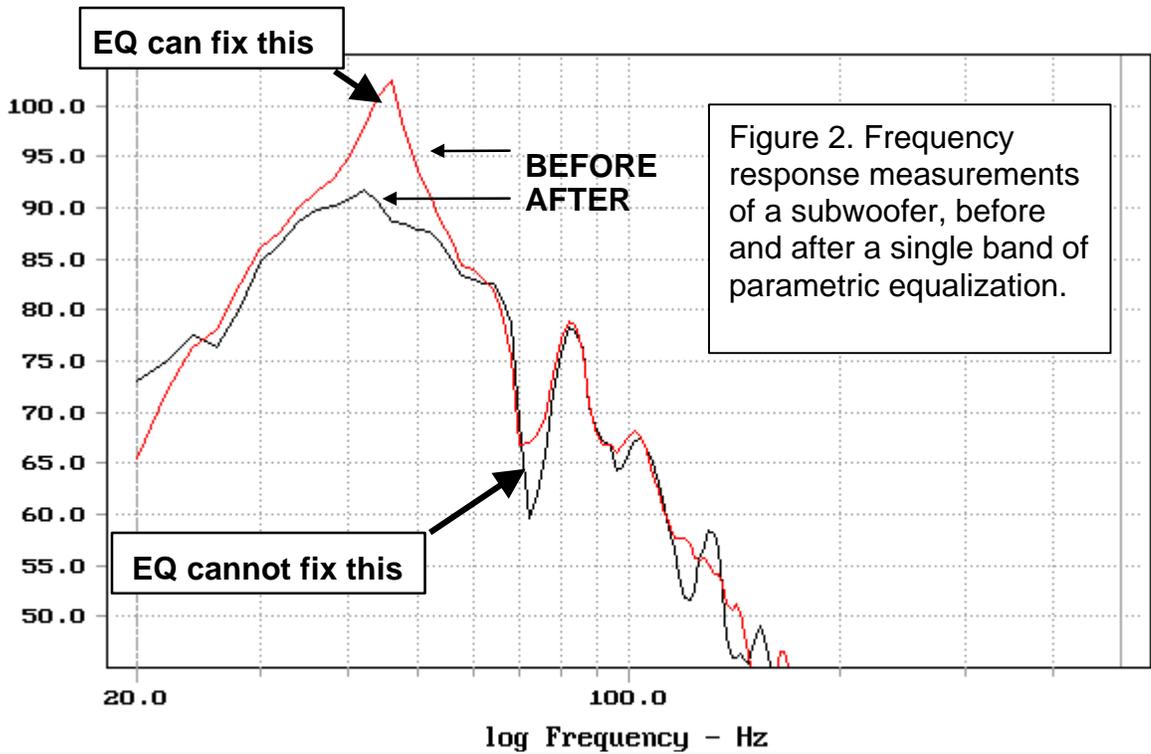
1. The popular measuring instruments, 1/3-octave real-time analyzers, do not have enough resolution to describe the problems accurately.
2. The popular equalizers, 1/3-octave “graphic” equalizers, do not have enough resolution to address the problem resonances specifically, without doing a lot of “collateral” damage.
3. Attempting to fill deep frequency response dips caused by acoustic cancellations or nulls is an absolutely futile effort, because no matter how much sound energy one pumps into a room the cancellation persists. All that happens is that amplifiers clip, and woofers distort, or worse, destruct. The only solution to this kind of problem is to relocate the loudspeaker or the listener, whichever is sitting in the null.
4. Equalization is attempted at too high a frequency. Low-frequency room resonances behave like minimum-phase phenomena, and addressing them specifically with parametric filters is a true solution. Above a few hundred Hz, the situation is very different, because we are using steady-state measurements to examine a complicated combination of direct and reflected sounds – time domain phenomena. The measurements may show “comb filtering” that is alarming to the eyes, but the ears hear only the natural sounds of a room – not necessarily a problem at all. If the reflections are perceived to be too energetic, the solution is not equalization, but rather the addition of some strategically placed sound absorbing or diffusing devices. As stated earlier, if there are obvious sound quality problems at middle and high frequencies, the only true solution is a properly designed, room friendly, loudspeaker.

INTELLIGENT EQUALIZATION

So, how do we do “intelligent” equalization? The first step is to work with high-resolution measurements that can show you what is really going on, the 1/3-octave real-time analyzers simply do not cut it. The ability to average measurements made at several locations within the listening area is a big help, because it will tend to attenuate the interference dips that equalization cannot fix, and help bring into focus the room resonances, that equalization very effectively addresses. Then the task is to decide what to change with equalization. A safe place to start is to use the equalizer to pull down peaks, and to avoid trying to fill holes. A broad, gentle, depression might be filled, but do it in stages, listening to see if something positive is actually happening. It is wise not to add more than a few (say 3 to 6) dB of boost. If you do add boost, remember that each 3 dB doubles the power requirement from the amplifiers, and the loudspeakers. Everything will be working much harder. Preferably, try to find acoustical ways to fill holes, and use the equalizer to smooth out the peaks. If there is a persistent notch, try to identify which mode is involved, and whether the loudspeaker or the listener is in or close to a null. Move the suspected element a foot or so away and see if there is improvement. The room mode analysis program is a big help in this – assuming you have a rectangular room. If all attempts fail, be content that at least the resonant peaks are gone, and that narrow dips are much more difficult to hear.

Here is an example, taken from the hands-on workshop held after this course at the 1998 CEDIA convention. The room we were given to work with was clearly going to be a problem (which we liked), because it had rigid masonry walls (good strong, high-Q resonances), and dimensions that were in simple ratios to each other (8 x 12 x 24 feet). With a TV at the end wall, a good viewing distance of 12 feet put the prime listening location close to the mid-point of the room. This puts the listener near a null for the first-order length mode ($1130 / 2 \times 24 = 23.5$ Hz). It is difficult to get very concerned about this, because there is almost no useful information at this low a frequency, but we might try to ensure that the listener’s ears are just ahead of or just behind the half-way point of the room. However, at the second-order modal frequency, 47 Hz, there is an abundance of audio information, and here the listener is sitting in a broad pressure peak.

Figure 2 shows what we measured for the subwoofer, by itself and, as predicted, there was a prominent peak right around 47 Hz. When we listened, the bass was flabby and boomy, with a “one-note” quality. Even movie explosions sounded faked. To address this problem we dialed in a single parametric filter, set to 47 Hz, with the appropriate bandwidth, or Q, and simply turned the resonance down. Room resonances at low frequencies behave as “minimum phase” phenomena, and so, if the amplitude vs. frequency characteristic is corrected, so also will the phase vs. frequency characteristic. If both amplitude and phase responses are fixed, then it must be true that the transient response must be fixed – i.e. the ringing, or overhang, must be eliminated. Figure 3 shows that this is so. Equalization of the right kind can work. Notice that we completely ignored the acoustical cancellation dip at about 73 Hz.



So, what would have happened in this case if we had used the traditional methods based on 1/3 – octave measurements and equalization?

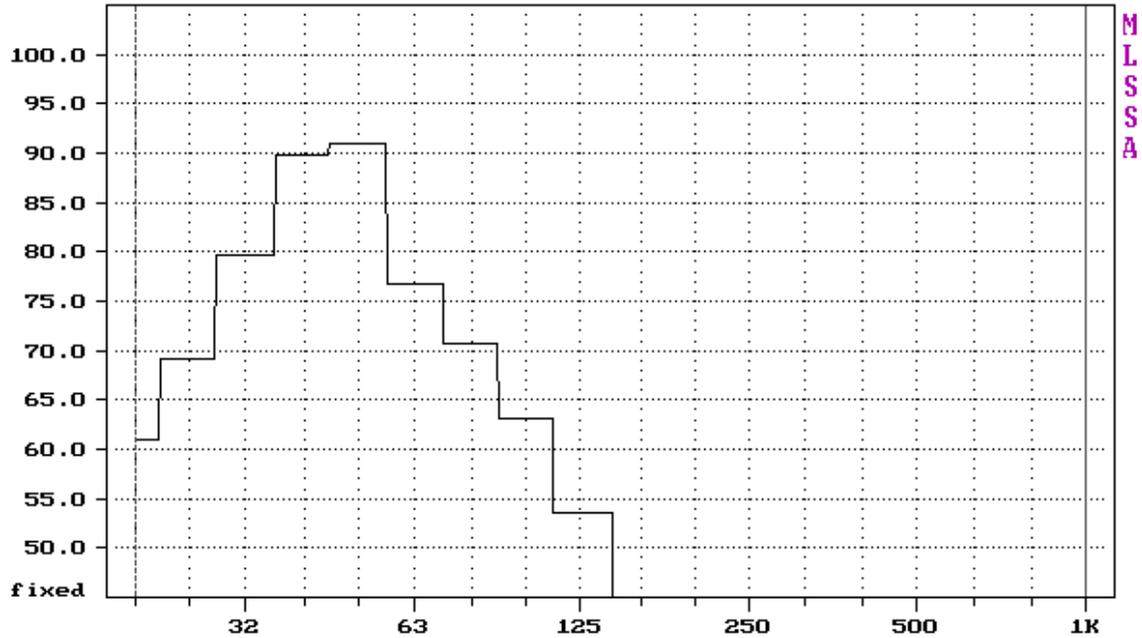


Figure 4. A 1/3-octave version of the unequaled high-resolution frequency response curve in Figure 2. Note the absence of any hint of a high-Q resonance at 47 Hz, and the lack of any evidence of an interference dip at 73 Hz. It doesn't look all that bad, really.

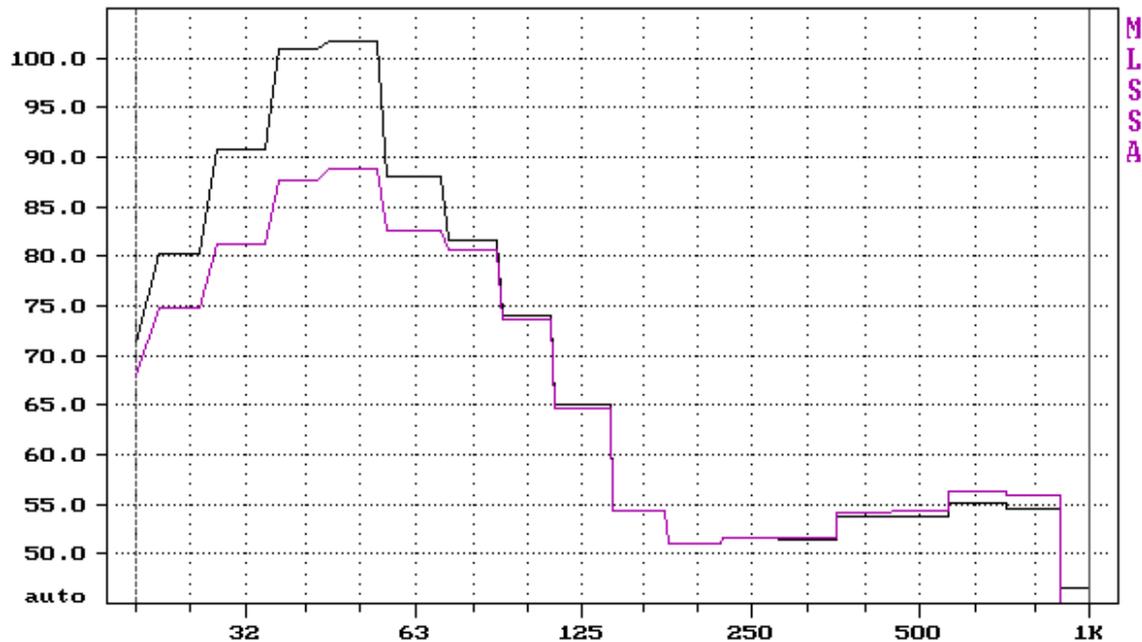


Figure 5. We think we can improve the shape a bit, so we dial in some attenuation at the two highest peaks, using a 1/3-octave graphic equalizer (lower curve).

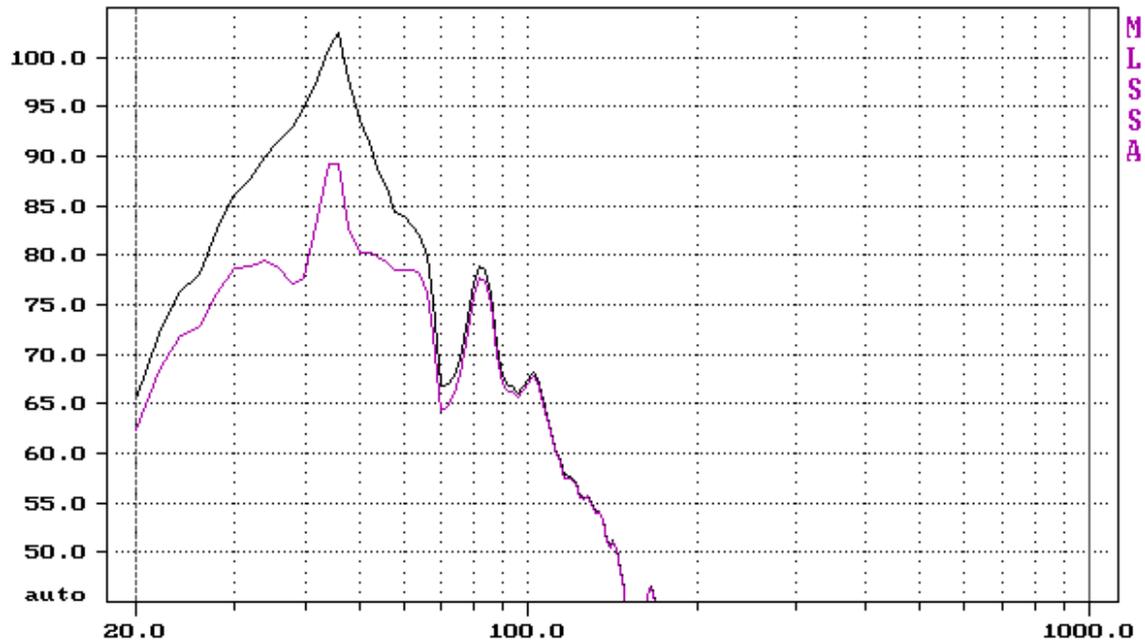


Figure 6. These are the high-resolution frequency response measurements of the subwoofer before (top curve) and after (bottom curve) equalization with the 1/3-octave graphic equalizer.

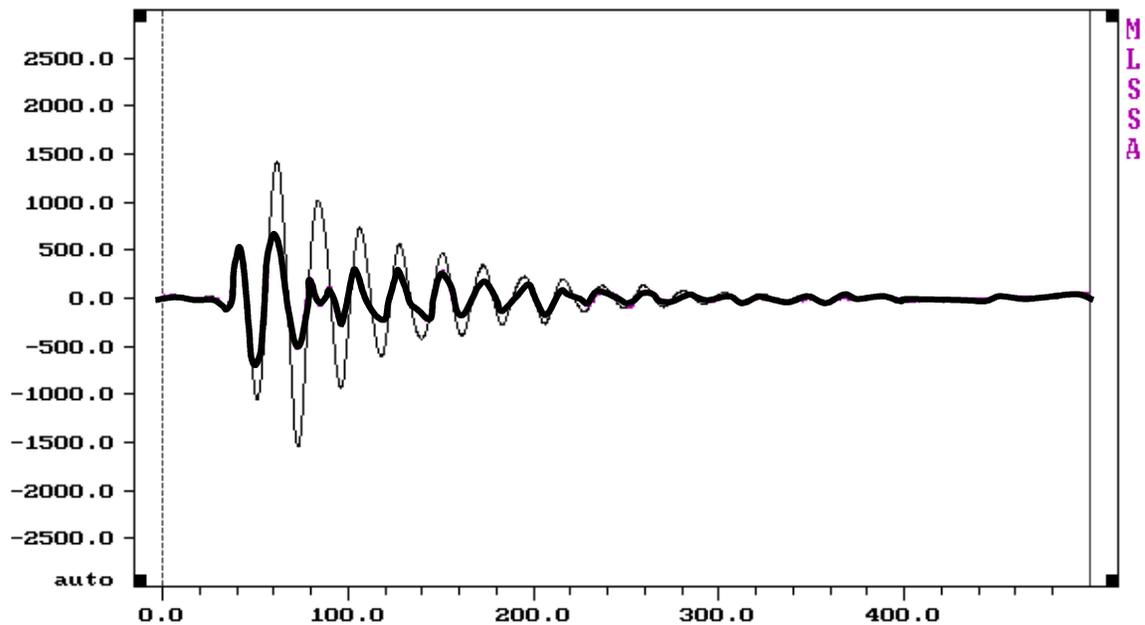


Figure 7. Time-domain behavior of the system before (light line) and after (dark line) equalization with a 1/3-octave equalizer.

The evidence of this series of measurements is quite clear. The 1/3-octave measurements lulled us into a sense of security by presenting data that did not look all that bad. Certainly there was no evidence of the sharp peak and dip that was obvious in the high-

resolution measurement. If we thought we could tweak the performance a little by using a 1/3-octave graphic equalizer, we did not succeed. It can be seen that, just as the 1/3-octave measurement failed to show any evidence of the high-Q resonance at 47 Hz, the 1/3-octave equalization failed to get rid of it. The annoying ringing, or bass “boom” was almost as strong after, as before, the equalization. It is no wonder that equalization has acquired its bad reputation. We threw away some good bass energy, and left the annoying boom.

To be fair, it is possible for a resonance to have a frequency that coincides with the center frequency of one of the 1/3-octave filters, and for that resonance to have a Q that can be addressed by a 1/3-octave filter at the same frequency. In that case, the resonance would have been attenuated, just as in the first example. However, a fortunate instance of that kind is not enough to justify using a system of such extremely limited usefulness.

SO, NOW YOU ARE A ROOM ACOUSTICS EXPERT

Perhaps that is an exaggeration. However, there is no doubt that the majority of people in the audio business do not know most of the facts and techniques you have just read about. They rely on the roll of the dice, or purely subjective trial and error, to get good sound. Often they do not succeed. This is not an acceptable way to run a business, much less an entire industry.

Given that the room is the final audio component, and it is the one over which the audio manufacturers normally have no control, any improvement is a tremendous asset to both loudspeaker manufacturers, and customers. Those people who know how to elicit good sound from loudspeakers in a room – every time – have an enormous advantage. They are the ones who can truly serve their customers, by delivering something tangible: great sound.

No, these few pages have not made anybody an expert, but they do contain the essential information and methods that, with study and practice, can make a person into one. It is an important start. For a more detailed analysis of loudspeaker / room interactions, read “Loudspeakers and Rooms – Working Together”.

03/13/00