The Acoustical Design of Recording Studios

By Malcolm Chisholm

When one considers that the recording industry has been building and using studios for about 70 years, it is remarkable that so little basic theory has been published on the subject. To be sure, there are plenty of "here's how we did it" articles in print as well as a number of "here's how to do it" examples to be found in books and magazines, but none of these provide enough of the underlying design principles to enable a reader to duplicate the performance of such studios unless he also duplicates the studio. For that reason, while such publications are interesting and even entertaining, they are of little use to a studio owner who wants to improve an existing room or build a new one in a space different from the exemplars given.

The situation grows even more extraordinary in light of the massive amount of experimentation and research devoted to control rooms over the past few years.

Those efforts have resulted in enough published material to allow a studio to select from at least two demonstrably excellent generic control room designs, both of which spring from the same clearly expressed theoretical underpinning. While the general case design will need some cleanup and tuning to achieve optimum results, a studio owner can use the published theory to modify a given plan, adapt it to his particular situation, and come up with a fundamentally decent room. In short, we know how to build good control rooms.

We sure as hell don't know how to build good studios. In fairness, there are some designers who appear to know something of the subject, but they don't give away their stock in trade, so a studio owner is faced with the problem of separating the good designers from the good talkers. With the near future of his business at stake, that's a serious problem, made worse by the fact that even very good acousticians have been known to make very bad mistakes when dealing with recording rooms.

As an example, the two worst studios the writer has ever encountered were designed from scratch by a Ph.D named Sabin. (There were three of them.) The rooms were retreated within a couple of months, but we turned out some pretty marginal work in the meantime. Since this happened in the city's premier recording facility, marginal was a bad case of egg on face.

Doc Sabin was not at fault in that mess.

In fact, nobody was. The whole thing was a terrible mistake.

The mistake was to confuse a recording studio with a normal acoustical environment.

Acousticians ordinarily think of large rooms in terms of theaters and auditoriums, which have a definite sound source feeding a definite audience. That applies equally to control rooms, theaters, auditoriums, and almost everything else acoustical designers get into.

It does not apply to studios, which are completely different from most other rooms. A studio has any number of sources in the persons of the musicians, and an audience comprised of those same players. Multiple scattered sources, ditto listeners. Peculiar room.

Keeping strictly to acoustical performance, the primary function of a recording studio is to provide adequate isolation between microphones while allowing the players to hear each other as well as possible.

Acoustical isolation is by far the most often discussed of these two areas, but since the parameters involved are addressable by acoustical mathematics, producing satisfactory isolation levels is a fairly straightforward process. All that's needed is a knowledge of what constitutes adequate isolation and several pages of mathematical computations. Happily enough, there is a way round that last item.

Treating a room for multidirectional listenability is a good deal more difficult, as it is not a direct function of the room's global characteristics and therefore cannot be treated mathematically. General solutions are available, and they work nicely, but they have more to do with old fashioned intuitive acoustics than with the glitzy new computer aided stuff. Among other things, this means that the merits of a suggested treatment cannot be readily confirmed by punching up one's handy dandy number cruncher.

Getting on with it, the specific design parameters are:

1. What is a reasonable isolation level and how much treatment is needed to get it?
2. What constitutes acceptable listenability and how is that managed?
3. What's the catch? (There's a bear in every forest)
4. ((Sometimes more))

Item one. Acoustical isolation between instruments is a function of the degree to which the sound of one dies away before getting to the next. When the die off is inadequate the sound of one instrument falls through the mike of the next and trashes it. If it is excessive the musicians can't hear each other properly, which makes group playing difficult and ruins section sound. Everything in the real world is a compromise, and acoustical isolation is no exception.

The amount of acoustical attenuation for a given instrument in a room depends on the room's global characteristics. As with any radiated field, sound pressure levels diminish as the square of the distance from the source. Double the distance, lose 6 dB SPL. The equation holds for any distance in a perfectly dead room or out of doors. In a normal room, however, the walls reflect some of the sound. Since the source sound level diminishes with distance, at some point the reflections from the walls will equal the source level. Beyond that point the sound no longer dies away, and the level becomes constant at any further distance. The distance at which this transition takes place is currently called the Critical Distance. It has been called other things in other centuries, as it's existence has been known for a very long time. It is easy to observe, easy to measure, and a remarkably accurate indicator of a room's acoustical performance.

The Critical Distance (Dc) of a sound source depends on the reflectivity of the walls and how much wall surface the sound hits. As an example, a firecracker hung on a string in the middle of a room produces a spherical sound field which will bounce off all six walls of the room. Six, because sound has no sense of direction, and can't tell a floor or ceiling from any other surface. This spherical source is assigned a "Q" (figure of merit) of 1, meaning it has no directionality.

Hang the firecracker against a wall, and it radiates a hemispherical pattern. That's a Q of 2. Halfway up the wall and in a corner it's a half hemisphere, and the Q is 4. On the floor and in a corner, Q=8. Q represents the beam width of the sound source. The higher the number, the narrower the beam.

The narrower the beam, the less wall surface is struck by a source's sound. Therefore, the higher the Q, the longer the Dc. And the higher the surface reflectivity, the shorter the Dc.

It follows from the above that low Q instruments will have the shortest Dc's, and the poorest isolation. As it happens, low Q describes both the human voice and the entire rhythm section. A moment's thought will explain that. Bass, piano, guitar and drums were used to accompany the human voice for several centuries before mikes and such were invented, and were designed to match it. They match quite well, which leaves us with a kit of Q 2.5 instruments as the basis of isolation design.

The most difficult instrument in terms of isolation is the voice. Not because it's so soft, but because of limiting. Unless a studio wants to turn out 1940's records, there is no choice but to limit vocals, and the limiter costs about 12 dB of isolation as it pulls up the consonants in the singer's words. What this amounts to is that the vocal channel should show something approaching minus 20 dB when the vocalist is quiet; 12 dB for limiting, and 6 to 10 to clear the consonants and allow a little dynamic range for the singer. Since other instruments work nicely with a clearance of 6 dB or so, adequate vocal isolation becomes the criterion for acoustic design in studios.

Vocal isolation is made a little easier by the small size of the instrument, which allows miking at a half foot without running into serious proximity effects, and generally presents about 87 dB SPL on mike. Hardly thunderous, but the peak level differences between voice and the rhythm instruments are not as great as commonly assumed. It's limiting up the minus 12 dB consonants that gives rise to vocal isolation problems. Still, since other instruments are 6dB or more over the voice's 87 dB, the room characteristics have to lay for about 26 dB of acoustical loss.

The problem is made harder by a simple but nasty fact.

A SOURCE GOES CONSTANT VOLUME AT IT'S Dc, AND THE VOLUME IS THE SAME EVERYWHERE IN THE ROOM.

Distance makes no difference in fallthrough. Directionality has no effect, except to mud up the fallthrough. Dead flats don't work. Hyper cardioid mikes don't work. Nothing works. The levels of the rhythm instruments have to fall about 20 dB before going constant volume, or you can't work a vocal anywhere in the room.

Item 1 is 20 dB.

More is nice, but getting much over 20 in a small room involves so much treatment that the studio turns into an anechoic chamber, with fuzz covering every wall.

Which brings up item two; listenability of the room.

Totally fuzzed walls return no sound to the players, who respond by playing louder. And much worse. It's very difficult for a group of musicians to work in concert (pun intentional) unless they can hear each other. Outdoors, with nothing otherwise coming back to the players, stage monitors are used to supply the sound of the group to the group. An engineer can use the studio playback speakers for the same purpose, and it works surprisingly well, but in both cases the players hear the mixer's balance, not their own. They play in one balance and hear another, which creates some subtile but nasty musical corruptions.

As an example, no mixer will let a solo ride at too low a level. It's the mixer's job to maintain a proper balance, and mixers do their jobs. So if a musician plays a tentative first solo, the mixer raises it's level as needed, and it plays back in proper balance. After a few takes, the soloist gets used to the idea that his solos will come out right no matter what he does, and lays back on all of them. It's easier to play soft. The mixer also adjusts to the situation, and bumps the level for each solo. All this sounds pretty good at the time, but after a few days both parties discover that the solos don't sound like solos. They sound like lifted fills. That's because a solo is generally a high energy item, and when a player lays back rather than putting out the energy, the solos lack drive and intensity.

Technically speaking, this is a matter of harmonic content. When an instrument is played hard or loud, the energy shows up as an increase in harmonics, and the result is a loud sound. When not, not, and artificially boosted soft solos just don't make it.

While this is one of the less obvious problems involved, musicians who can't hear their overall sound well enough to maintain solo and section balances during performance are very unlikely to play at their full potential, and they need to hear themselves directly. That's especially true if the mixer is tricking up the sound as it goes through the console to the tape.

Since the usual studio setup points the musicians and their instruments at the control room, the obvious (and normal) way to supply direct feedBack is to bounce the player's sound off the control room wall.

Control room walls are left reflective as a matter of conventional wisdom, and are even somewhat optimized by stacking the musicians cases against the wall under the control room window. The cases offer a fair degree of dispersion to the strong boundary layer sound traveling along the floor to the control room wall, and integrate it before reflecting it back to the players. That won't work with a rug on the floor, but improves things quite a lot otherwise. Given that a primary function of the control wall is to supply a live surface to the musicians, it can be made more effective by using some of the techniques employed in the backs of control rooms. Some of these would involve substituting RPGs for the stacked cases, retreating the wall for maximum reflection, moving the control room window to the upright position, and installing a reflector above the window angled to bounce even more sound back to the rhythm section. The combination of flat and dispersed reflections has been shown to be optimum for critical listening, and if it's good enough for engineers, why not supply it to the people who are doing the actual work in a studio?

In any case, the amount of acoustical treatment in a studio is limited by the need to leave the major part of the control room wall reflective. And there advantages to a live floor in allowing solid boundary layer sound at the control wall, in addition to making it easier to move things around in the studio.

The side walls are far less critical. Because of that, they are commonly either left untreated or given some kind of uniform treatment. Neither is a good idea.

Flat, straight walls have been known to be acoustically unacceptable for centuries. That is partly because sound reflects off such walls as a sort of flat smack, which sounds bad, and partly because it bounces so strongly. If the side walls are either untreated or evenly treated the sound will ricochet around the room like a ball on a billiard table until it finds an open mike to get into. That was the problem with Doc Sabins' room. It is probably possible to control the results by putting an absorbent flat behind every mike in the room, but it's a tedious process, and interferes with player communication. Much better to clean up the bounce. Since the villains in the piece are flat, evenly treated walls, the obvious remedy lies in knobbing up the walls and installing absorptive treatment in patches.

Both objectives can be accomplished by hanging live sided boxes filled with fiberglass on the walls. (See drawing.) Floor to ceiling treatment is unnecessary, as mikes rarely point up. The boxes should start high enough off the floor to clear chairs and other clutter leaned against the wall, and will generally top out at eight or nine feet above the floor.

A box with reflective sides will act as a disperser, and at a foot or so deep will disperse sound down to about 550 Hz. Not ideal, but not bad, and at a foot the boxes are pretty manageable. They are normally spaced at two to four feet apart, leaving the walls reflective between them. This presents a combination of dispersion, absorption and reflection to both the musicians and the mikes, and cleans up the billiard ball syndrome quite nicely while presenting an optimum listening environment to the hard working types in the studio.

The back wall can be treated in the same way in small rooms, although it is best to leave the back as live as possible, as reflections from it give the players a sense of being in a room rather than working with their backs to a vacuum.

In cases where a great deal of absorption is needed, the wall area above eight feet and below about two can be totally treated without ruining the generally live sound of the room, as the ear only needs a little encouragement to think it's in a normal environment.

The ceiling is another matter, and needs be almost entirely dead, because it is almost never high enough to establish a decent modes structure. The standard literature lists acceptable room proportions of up to two to one as an extreme case, and the vast majority of ceilings are well over that. As always, the best way to deal with an unsolvable problem is to eliminate it, and since a non-reflective surface generates no modes structure one way or another, dead ceilings are the norm in most studios.

The ceiling also presents the largest area available for serious treatment, especially as it can be totally absorbent without making a room sound dead. Short, yes. Dead, no. Unless a ceiling is extremely low the ear ignores it preferring to take it's cues more or less horizontally.

It is critical that the ceiling treatment be acoustically flat in it's absorption. Given an ordinary grid hung 16 inches below the structural ceiling, flat response can be accomplished with 1-1/2 inch fiberglass ceiling panels or with thinner panels and a fiberglass batt overlay. It is wise to check manufacturer's literature for exact specifications, as the low end absorption of the ceiling must extend far enough into the bass range to avoid the common fault of acoustical treatment that soaks out the top end of the room while leaving the low end live. The that kind of treatment results in a muddy room with terrible isolation problems in terms of bass, floor tom, and bass drum.

The need for flat low end response applies to all room treatment unless the studio has big windows or it's walls are so flimsy as to transmit or absorb bass by vibrating to it. Even so, bass attenuation will seldom exceed 30%, leaving 70% to be supplied by other means. While there are any number of bass absorbing devices which can be built or purchased, they are inconsistent in operation, inefficient except in corners, and very difficult to analyze as to the number and size required. On balance it's more practical to install the general treatment in such way as to absorb uniformly from bass to cymbals. Controlling high frequency reflections is easy, but bass absorption is largely a matter of absorber depth, and it takes considerable thickness to get flat down to 60 Hz. Hung ceilings manage it with thin panels and the 16 inches between panels and the real ceiling, but a wall mounted absorber needs a minimum depth of 6 inches for fiberglass (703) board, and a foot for glass wool.

Don't use thin treatment. Carpeting and drapes absorb 2 to 14% of bass while soaking out 60 to 70% of the top end, yielding a room with no presence and extreme boominess. Bad for playing, worse for recording. At 70% efficiency, they also require an excessive amount of treatment and wall area. Interestingly enough, both products cost far more than proper acoustical materials, and are not necessarily more attractive. By and large, fiberglass in one form and another is probably the most practical treatment available, and it can be covered in any number of handsome fabrics or in Tectum if a durable wall is needed.

With the type and location of studio treatment in hand, we can finally address the question of how much absorption is needed.

The following data are not hypothetical. The Dc figures were discovered during extensive reality testing of a newly written acoustical design computer program. The test method consisted of retro engineering a number of recording studios, in each of which the writer had done some hundreds of sessions. The majority of the studios were acceptable, a few were marginal, two were bad, and two superb. The object of the exercise was to find a common parameter that related to actual studio performance, and the voice Dc proved to be a figure of merit for isolation in properly treated rooms. Other correlations became evident over several years of repeated computer runs on these and other studios in an acoustics course taught by the author at a local college.

Designing for isolation is both simpler and more difficult than it first appears. The simple part is very simple indeed, as voice Dc (and therefore isolation) turns out to be function of the amount of absorption in a room regardless of room size. The absorption required for 28 dB of acoustical loss at 20 feet a voice Dc of 11- 1/2 feet is about 2700 Sabins. Sounds easy.

If the practice were as straightforward as the theory, one could stuff 2700 square feet of fiberglass in a studio and open for business without further ado. Unfortunately, what's wanted is 2700 Sabins of absorption, and the actual amount of treatment for that figure can vary from less than 1500 to just over 2500 depending on the size of the room.

That's the first bear in the woods.

The reason for a difference between actual treatment and effective absorption is that the standard Sabin formula is linear, and absorption in highly treated rooms is not. In fact, when 80% of the wall surface absorbs at 1 Sabin per square foot, the effective absorption of the treatment is doubled. There is a formula for this effect, (Norris Eyring) which is reasonably accurate, but since it involves the use of natural logarithms, it is tedious to use.

Second bear.

The third bear is the well-populated acoustical forest is the difficulty of accurately assigning absorption values to various materials already in the room. Most standard materials can be looked up in tables printed for the purpose, but there are always a few things that aren't on a table. Additionally, it is very easy to mistake one kind of acoustical material for another and come up with significant errors in calculations.

Calculations are a pain anyway, so it's best to circumvent the bears by measuring the room.

There are several thoroughly scientific ways to do this, and any number of manufacturers eager to sell equipment for the purpose, but as a practical matter such measurements are of little or no use to the studio owner. Cheap equipment yields cheap results, and the data gleaned from upscale equipment require expert (and costly) interpretation. In the first instance, the figures aren't completely trustworthy, and in the second routinely repeating the tests will cost a fortune.

Following the KISS (keep it simple, stupid) rule, the writer prefers to measure a room by determining it's voice Dc. The equipment costs nothing, it takes about two minutes, and the results are more than accurate enough for real world use. Better still, being a simple-minded test, it reports simple-minded figures.

No interpretation, no ambiguity. Best of all, a Dc check makes it's measurement at about 100 Hz, where improper treatments cause a the majority of isolation problems.

Measuring a voice Dc is child's play provided one keeps in mind that the purpose is to determine the global characteristics of the space. Toward that end, it is essential to make the measurement in the acoustical center of the room. Given normal treatment, that will be in the physical center. In cases where the absorption is considerably greater on one wall than another, the acoustical center will have to be found.

Again, dead easy. Using the incredibly sensitive instruments found on either side of the human head, one sidesteps away from one wall toward another until the reflected sound from the two are equal in each ear. If the reader has not done this in past, he may find it useful to calibrate his ears to wall sound by stepping up to a live wall and varying his wall to head distance from a couple of feet to a couple of inches until the wall sound is firmly fixed in mind. It is usually perceived as a kind of pressure on the ear, and will very reliably inform the listener of his position in a space. No sound other than the room's random noise is needed, and once the listener knows the sound of a close wall he will find that he can walk to within a foot or so of any live wall with his eyes closed, This is simply a case of practicing a normal human ability into a skill. The blind do it all the time. So do the rest of us, but unconsciously. The writer once deadened one wall of a hallway, and sighted people veered into it to the point of wearing out the treatment.

Having determined the acoustical center of the room, Dc is measured by two people more or less astride the room's center starting at a distance of 15 to 20 feet. One of them walks toward the other droning one, one, one as the other waits for the sound of the talker's voice to suddenly get louder. The process works both ways, with the talker's voice abruptly going constant volume as he retreats, but the writer's experience with several hundred students indicates that toward is easier to hear than away, particularly in the learning stage. It is also easier to hear if the talker walks briskly at first. He can slow down for greater accuracy once the listener has the sound of the transition in mind. The pair can also check room's frequency response by measuring the Dc using the word six, leaning on the s and x and suppressing the vowel, so that most of the sound is at 3 to 5 Khz. This is a pretty rough test, but if the Dc's are wildly disparate, they indicate a room with more absorption in the midrange than at the low end.

While Dc is a square root function of a room's global characteristics and there- fore a rather short ruler, the breakover is sufficiently abrupt to make measurements to within a couple of inches quick, easy, and repeatable by any number of talker/listener pairs. The only conditions under which it doesn't work properly are rooms in which the Dc is greater than the wall spacing (rare) and huge rooms which appear to divide themselves into several acoustical areas due to extreme losses between one wall and another. In the first instance the room will be too small and dead to be of any practical use, and in the second the room volume will be well in excess of a million cubic feet. The writer knows of one at 6 million that acts funny, but it's in no danger of being used for studio work.

Once the Dc of a room has been measured, some acoustical modifications may seem in order. If so, a few cautionary notes should be kept in mind.

First, the Dc varies as the square root of the room absorption, so doubling the effective treatment and thereby halving the reverberation time will extend the Dc to only 1.4 times it's previous figure. This presents no problem in a medium to large room, but good isolation in a 30x20x10 foot studio would require some 1550 square feet of fuzz scattered over only 2400 square feet of surface area. Even with the floor thickly carpeted, leaving a 20x10 foot control room wall reflective would require a 75% treatment of the other walls, and result in a reverberation time of just over one tenth second.

Some rooms are simply too small to treat for live studio work, as they get too dead. The 6000 cubic foot case in point is probably the workable minimum.

Second, a big studio is rarely allowed more that about one second of reverberation time, which results in a voice Dc approaching 20 feet. Obviously, such a room needs no help in isolation, and is best left alone. It is a general rule in acoustics that big rooms are easy. It's the little ones that give you fits.

Third, professional engineers commonly do good work in bad conditions. The writer has done any number of sessions in studios with 7 to 8 foot Dc's which turned out well enough to sell bags of records. It's not impossible to record in a room with poor isolation, it's just damn hard work. The point of proper treatment is that it allows one to get decent sound with any reasonable setup, and it eliminates time lost in fooling around trying to correct the room's faults.

Fourth, none of the figures given are engraved in stone. A twelve foot Dc is better than eight, and less good than sixteen, but acoustics are inherently inexact, and there is no sharp point at which rooms switch from bad to good; they just glide from exasperating to no problem, with the latter occurring and something around a 12 foot voice Dc for the bulk of studio work.

In summary, a few minutes spent in measuring the real world characteristics of a recording studio may reveal unnecessarily poor isolation, and/or some of the treatment methods suggested herein may improve it's general usefulness. Since the measurement involves no expense and the treatment is designed to make experimentation easy, these techniques offer a practical way for a studio to confirm or optimize it's recording rooms.