Multichannel sound systems and their interaction with the room

David Griesinger

Lexicon
3 Oak Park
Bedford, MA 01730
dg@lexicon.com

Abstract

The ideal number and placement of drivers in small listening rooms has been controversial. Most research has examined only the pressure (loudness) as a function of frequency and source-listener positions. We believe that two additional properties of the soundfield, externalization and envelopment, are important to listener preference. Externalization is solely a low frequency problem, and the interaction between envelopment and loudspeaker position is particularly important at low frequencies. We proposed in a previous paper a mathematical method of quantifying these two perceptual properties given a measured or calculated binaural impulse response. The Average Interaural Time Difference (AITD) is the measure for externalization, and the Diffuse Field Transfer function (DFT) is the measure for envelopment. This paper will present a simple image model for rectangular rooms, and use it to study the interaction between multiple drivers and listening rooms. The image model is used to predict the values of pressure, AITD, and DFT for different combinations of room properties and driver locations. It is found that the low frequency pressure uniformity and the AITD can be increased in the prime listening area by using multiple low frequency drivers – especially at the sides of the listeners. When playing material where the bass energy is primarily monaural, the drivers on the left side of the room should lead or lag the drivers on the right side by a constant phase angle of 90 degrees. With musical signals we find DFT is also increased by multiple woofers. Once again a side placement is optimal.

1. IMAGE MODEL FOR SMALL RECTANGULAR ROOMS

The work of developing measures is greatly aided by an efficient method of checking how they work in typical rooms. Unfortunately it is difficult to find a room which seems typical, and it is exceedingly tedious to measure a large number of binaural impulse responses for each new loudspeaker arrangement. We need an efficient computer model. Image models have the virtue of simplicity and computational speed. Unfortunately the image model assumes that the surfaces involved in the model behave as simple plane reflectors – or that the reflecting surfaces are large compared to the wavelength of the sound being reflected.

This assumption is clearly violated when we study small rooms. However the principal error is due to diffraction at the boundary between two surfaces of different reflectivity. When all the surfaces of the room all have the identical reflectivity the model could give reliable results. In fact, we have tested the model against measurements in two rooms, and have found the model to predict the results surprisingly well. In spite of the absorbing ceiling in one of the rooms the model produces plausible results. Our conclusion – the image model may not be perfect, but for developing measures and concepts it is more than good enough!

2. A SIMPLE MODEL FOR THE HUMAN HEAD

However we are looking for more than the sound pressure at a point in the room. We are looking for the interaural time difference for a human head placed at a given point. ITDs are influenced both by the distance between the ears of a listener and by sound diffraction around the head. Ideally we should sum delayed and attenuated head related transfer functions for each image. While not inconceivable, this procedure would be computationally expensive – and each HRTF would need to be quite long to accurately yield the ITD at low frequencies.

Fortunately published HRTF data suggest that for frequencies below 125Hz the interaural delay can be accurately predicted if we model the head by two omnidirectional receivers, spaced by about 25cm. The model works fairly well above this frequency if the interaural spacing is reduced. This simple head model allows us to find the sound pressure at each ear by summing the pressure contribution from each image. We do not have to worry about the sound direction. This head model is an enormous simplification.

It is valid only for low frequencies, but it makes our image model practical.

3. DETAILS OF THE ROOM/HEAD MODEL

Our image model is written in MATLAB. The code is available from the author by email on request. The image model uses loops and conditionals, so the Matlab C compiler (with real number math) is highly recommended. 49 binaural receiver positions and two or four source positions can be evaluated in a few minutes using compiled versions of critical subroutines, something that takes hours without compilation.

We use a recursive method to calculate the images that result from an arbitrary source position in a rectangular room. We first find the line of images formed by the side walls, out to the selected image order. We then reflect this line of images with the front and back walls to form a plane of images. This plane is then reflected with the floor and ceiling to produce a series of image planes. The strength of each image is found by multiplying the source strength by the reflectivity of each surface encountered.

Once the images are found, the distance from each image to each receiver position is calculated. These distances are combined with the image strength to calculate the binaural impulse response for the particular source/receiver pair. To find the impulse response we use a sampling technique. The sampling process results in timing errors, which can be particularly important when studying ITDs at low frequencies. For the work on externalization we find a sample rate of 44100Hz gives good results, and allows the resulting impulse response to be convolved with recorded source material. The work on envelopment required an initial sample rate of 176400Hz to give consistent results. We found that splitting the pressure (not the energy) from each reflection linearly between adjacent samples gives the best results, and we sum pressures (not energy) from multiple reflections.

The Fourrier transform of the impulse response gives the steady-state pressure response as a function of frequency. One can use the resulting response curve to estimate the number of images needed. As the number of images increases the length of the impulse response increases, as does the sharpness of the individual peaks and notches in the response. In practice one increases the order of the reflections until the pattern stabilizes. Ideally one would like to double the order between each test. Unfortunately the sharpness of the resonances increases approximately linearly with the reflection order, and the computation time increases as the cube of the order. There is a large payoff in using the minimum number of images.

When the room surfaces have a reflectivity of 0.9 we need an order of at least 16 to approximate the room response. Fortunately we have not had to model a room which is that reflective. To model the rooms we have measured, a reflectivity of about 0.8 seems to work, and a reflection order of 11 seems sufficient.

Our model calculates the contribution of each image to the total amplitude and phase at the receiver position. Although the method assumes that the surfaces have no net phase shift with each reflection, such a phase shift could be modeled over a narrow frequency range by simply altering the dimensions of the room.

There have been studies that compare the pressure distribution measured in real rooms to results calculated with an image model. In general the accuracy of the model has been good. Although we have not made careful measurements of many different rooms, at least in two listening rooms the pressure distribution at several frequencies was checked with a sound level meter. The match to the predicted patterns from our image model was good. In another experiment the variation of pressure with frequency over a range of 30Hz to 100Hz was measured. The average absorption was then adjusted in the model to make the best match with the measured response. For the particular room – a 12’x15’x9’ listening room – the best match occurred with an average reflectivity of 0.8 for all the surfaces. Once this was chosen, the model and the measurement agreed within 2dB.

4. DFT IN AN ANECHOIC SPACE

In the previous paper we presented a measure called the Diffuse Field Transfer function (DFT) for low frequency envelopment in small listening rooms. In the examples here we will plot results for the 63Hz octave band, although results at higher frequencies are likely to be equally interesting. We found in the section on calibration of the DFT that an optimum value would be approximately 0.24ms in the 63Hz octave band. This calibration depends on details of the bandwidth and time constants chosen, and has little real physical meaning. However the results we will show here will plot the DFT in dB relative to this value.

As mentioned in the previous paper a single sound source in an anechoic space cannot produce envelopment, and the DFT value is consistently <-40dB. When there are two sound sources (stereo woofers) driven by decorrelated signals in an anechoic space the DFT can be significant. In fact, when we are directly between the two sources, we will get the calibrated value of 0dB. In conventional stereo listening the speakers are in front of us, at an angle of +-30 degrees. Figure 1 shows the DFT for conventional stereo in an anechoic space.

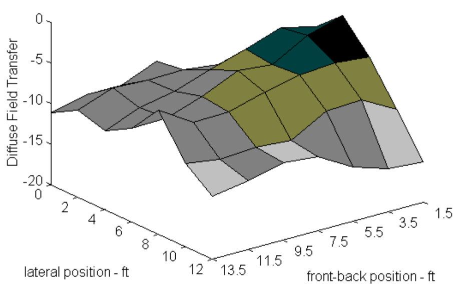

a:

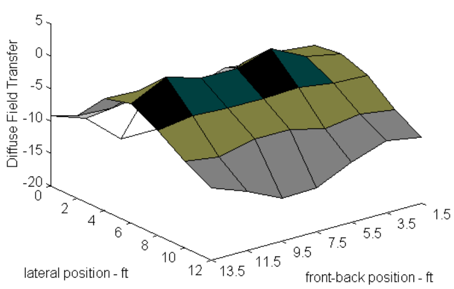

b:

Figure 1: Diffuse Field Transfer function (DFT) for a 12'x15' anechoic space. 63Hz octave band. A: two uncorrelated speakers in the front, +-4' from the center. The value of 0dB is optimal. Note in the listening area the DFT is reduced by the +- 30 degree angle between the loudspeakers and the listener. B: The same space with the speakers at the side, at 11.2' from the front. Note that the DFT is optimal through the listening area. This corresponds to stereo subwoofers at the sides of the listening area.

In figure 1 the listening plane is at 4' from the (immaginary) floor, and the speakers are at 1.5' from the floor. Figure 2 shows the value of DFT along the center line of the room.

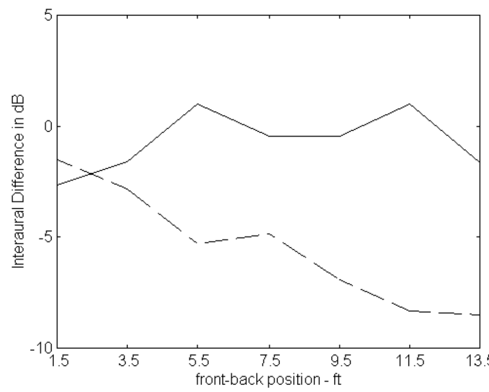

Figure 2: DFT along the center line for the two configurations in figure 1. ___ = two drivers at the sides of the listener __ __ __ = two drivers at the front. Note the approximately 5dB difference in envelopment in the listening area.

Note that with the speakers in the front the envelopment is reduced by 5dB compared to having the speakers at the sides. Figure 2 also shows the approximately +-2dB accuracy of the DFT calculation. The ideal value, assumed to be ~.24ms, is achieved in the middle of the listening area when the drivers are at the side.

5. DFT IN A REFLECTIVE ROOM FOR NOISE SIGNALS

When we octave band noise at 63Hz as a test signal and we add reflective room surfaces it is possible to have significant values of DFT with a single loudspeaker. This result was surprising to the author, who was expecting the results to agree with his experience with musical signals. Some quick listening tests revealed that the DFT was accurate. We will show later that musical signals give quite a different result.

Figure 3 shows the DFT for a 12'x15'x9' room, with a surface reflectivity of 0.8. Figure 3a shows what happens with a single driver, at 4' to the left of the center line in the front. DFT is surprisingly high throughout the listening area, even with a single source.

Figure 3b shows what happens when there are two sources. DFT increases somewhat, and the uniformity of the DFT also increases (although this may be an artifact of the noisy DFT measurement.) Figures 6a and 6b show that at a surface reflectivity of 0.8 there is little advantage in envelopment to having two drivers, a result that reflects the use of a broadband test signal. However it is clear that as the reflectivity goes down having two drivers will become much more important.

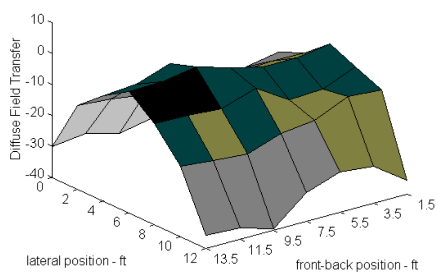

3a:

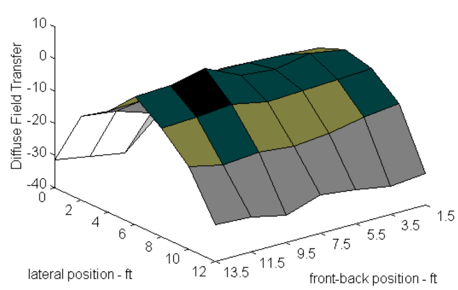

3b:

Figure 3: DFT for a 12'x15x9' space with surface reflectivity of 0.8. 63Hz octave band. 3a has a single driver in the front of the room on the left side. Figure 3b has two drivers separated by +-4' in the front, each with uncorrelated noise. Note the somewhat higher DFT in the second case, with slightly improved uniformity.

Figure 4 shows the DFT along the center line of the same room for three cases. The highest curve is the DFT when all surface reflectivities of 0.8, as in figure 6b. The large dashes show the DFT when the lateral reflectivity is reduced to 0.6, with two uncorrelated drivers. Note that there is very little decrease in the DFT. The curve with small dashes shows the DFT with only one driver, with a lateral reflectivity of 0.6. Note that the envelopment is significantly reduced. This corresponds to the case of a single subwoofer.

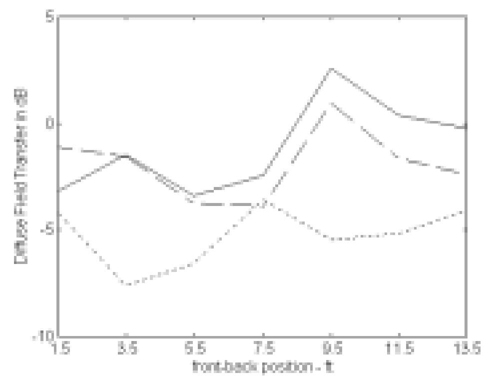

Figure 4: The DFT along the center line of a 12'x15'x9' room. 63Hz octave band. Speakers in the front of the room at +-4'. ____ = Two speakers, room reflectivity of 0.8 on all surfaces. __ __ __ = same, but with a lateral reflectivity of 0.6. - - - - = a single loudspeaker with a lateral reflectivity of 0.6. Note the reduced envelopment when there is a single sound source and a low lateral reflectivity.

Using the DFT modeling tool with other room dimensions gives similar results. In general, when the lateral reflectivity is high the monaural DFT determines the overall envelopment of the room. When the lateral reflectivity drops below 0.6 the envelopment drops dramatically unless there are multiple low frequency drivers. When the lateral reflectivity drops below 0.5 there is a large advantage to locating the low frequency drivers at the sides of the listeners.

6. DFT IN A REFLECTIVE ROOM FOR MUSIC SIGNALS

In the previous section we found that a single loudspeaker reproducing a noise signal was capable of producing substantial envelopment in a small room if the reflectivity of the surfaces was over 0.65. Music can have much narrower bandwidth. We are interested in the transfer of the reverberant component of a recording to a listener. If we imagine a bass instrument – such as a string bass or organ pedal – that produces a tone and then stops, the bandwidth of the resulting reverberation can be quite narrow. As we reduce the bandwidth of the test noise signal, the DFT from a single driver becomes much lower, while the DFT from a pair of drivers with independent signals stays about the same. Figure 5 shows the DFT along the center line of a 12’x15’x9’ room with the speakers either at the front in the narrow end of the room, or at the sides of the listening area. The filter frequencies chosen were 62Hz to 65Hz, for a 3Hz bandwidth.

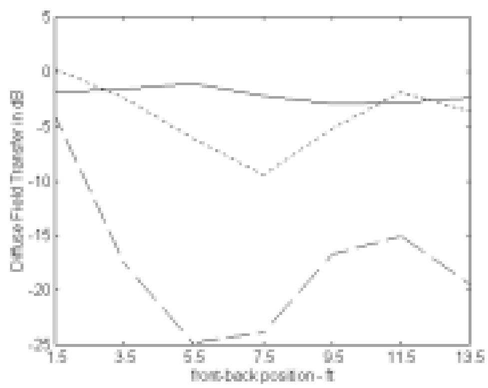

Figure 5: The DFT along the center line of a 12’x15’x9’ room with surface reflectivity of 0.8.
____ = Two uncorrelated loudspeakers at the sides of the listeners, at 11’ from the front.
 - - - = Two uncorrelated loudspeakers at the front of the room, 4’ apart.
__ __ __ = A single driver in the front of the room, 4’ to the left of the center line.
Bandwidth of source signal is 3Hz, from 62Hz to 65Hz. Compare this figure to figure 4, where the bandwidth of the test signal is 45Hz.

We conclude that for many types of bass instruments there is a substantial advantage to stereo low frequency loudspeakers, even in reflective listening rooms. The DFT with a narrow band test signal – or with actual reverberation from a musical source – can be used to quantify the difference, and to find optimal loudspeaker positions. Once again, there appears to be an advantage to placing the low frequency drivers at the sides of the listening area.

7. AITD AND PRESSURE FROM A SINGLE DRIVER IN A REFLECTIVE ROOM

In a previous paper we presented a measure for externalization of a sound source. The measure was called the Average Interaural Time Delay or AITD. When the AITD was developed it was intended as a measure for both externalization and envelopment. Although it was ultimately not useful for envelopment, it is clearly closely related to the DFT in the way it varies with room shape and speaker placement. AITD is much simpler to calculate.

In the previous paper we showed some curves of how the AITD behaves in an anechoic space. The anechoic case is a good test of the theory, but not common in practice. In a real room standing waves reduce the total AITD, making the sound source more difficult to localize, and making the “in the head” perception more likely.

For example, Figure 6, 7, and 8 show the lateral, medial, and total AITD for the 63Hz octave band in a 17’x23’x9’ room, with wall reflectivity of 0.8. The driver is in the upper left corner, at position x=1’, y=1’, and z=1.5’. The receiver plane is at z=4’.

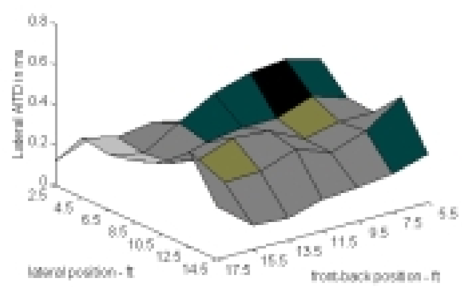

Figure 6: Lateral AITD in the 63Hz octave band for a 17’x23’x9’ room, surface reflectivity 0.8. Note the value is lower than in the free field.

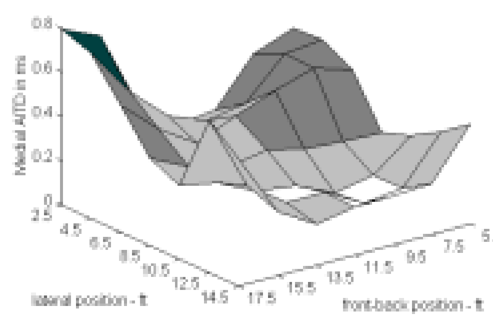

Figure 7: Medial AITD for the same room. Note the value is larger than figure 5, but still

lower than free field. This drawing represents the lateral AITD if the primary listening axis is parallel to the short wall, rather than parallel to the long wall.

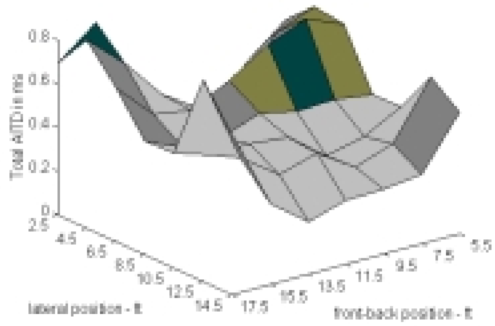

Figure 8: The Total AITD for the same room. In the listening area the AITD is roughly half the value of the free field. Sounds are somewhat externalized in such a field.

Note that all the AITDs are lower than for an anechoic room. The total AITD is minimal in just the area of the room you are most likely to choose for critical listening.   Figure 9 shows the total AITD along the center line of the room. It is reasonably constant at about 0.4ms. Experience has shown that low frequencies in this room are weakly externalized when a single driver is used.

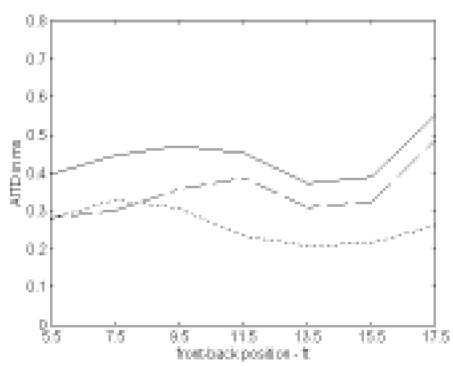

Figure 9: ___ = Total AITD along the center line of the same room. - - - = Lateral AITD, __ __ __ = Medial AITD. Receiver height at 4'. Some externalization is possible in this sound field, but it is not as easily externalized as a free field.

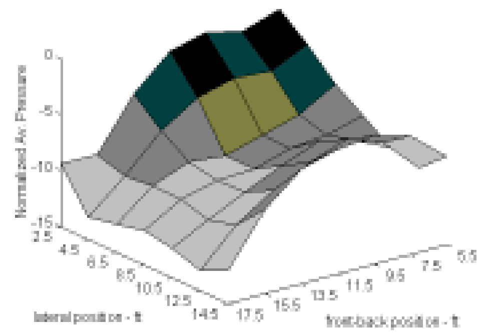

Figure 10: Normalized Average Pressure in the 63Hz octave band from a single driver of unit strength in the upper left corner of a 17'x23'x9' room. The wall reflectivity is 0.8. Note the pressure in the listening area is low. Equalization can raise the pressure, but it cannot change the AITD.

Figure 10 shows the normalized pressure in the 63Hz octave band for the same room. Note that pressure is not uniform. There is a concentration of pressure near the driver, and the minimum pressure is in the preferred listening area. This is true even though we are averaging over an entire octave.

Figures 11, 12, and 13 show similar data for a smaller room, 12'x15'x9'. The unusual shape of the Lateral AITD surface in figure 11 is due to a strong standing wave at about 70Hz. Most of the other frequencies are well represented by figure 12. Once again we see that pressure is generally lowest just where we would like to listen, and so is the lateral AITD.

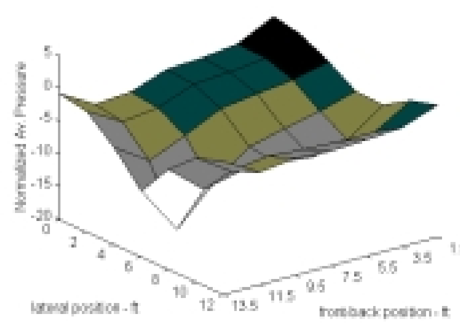

Figure 11: The Normalized Average Pressure in a 12'x15'x9' room in the 63Hz octave band from a single driver in the top left corner.

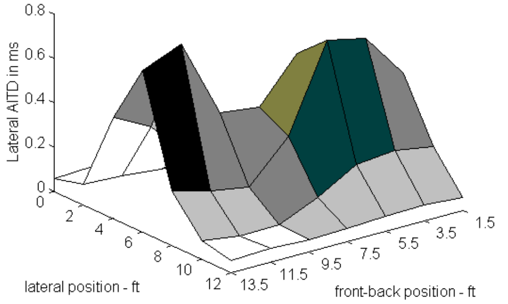

Figure 12: Lateral AITD in a 12’x15’x9’ room, surface reflectivity 0.8, single driver in the top left corner, 63Hz octave band. This band has the largest AITD for this speaker position in this room. The next figure is more typical.

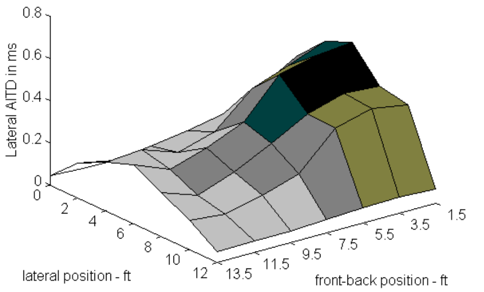

Figure 13: The Lateral AITD over a range of 20Hz to 90Hz in a 12’x15’x9’ room from a single driver in the front left corner. Surface reflectivity is 0.8.

The medial AITD for the 12’x15’ rooms is plotted in figure 14. Unlike the room of figure 7, this room shows a substantial forward localization. The difference is significant. The medial AITD represents the lateral AITD for a listener who is facing the long wall of the room. If we decided to set our stereo system along the long wall rather than along the short wall, the lateral AITD would be much higher. The difference shown here has some historical significance. The work in this paper was prompted in part by the author’s observation that in his 12’x15’ listening room the sound was much more pleasant when the system was oriented so the listener faced the long wall. There are probably several reasons this orientation is preferred in this room, but the high values of lateral AITD are likely to contribute significantly. Typically one is using two full range loudspeakers in such a room, not one. In this case the meaning of the high value of lateral AITD is that when there is a strong low frequency signal in one of the two stereo channels (and not the other) the low frequencies will be external and localized to the side. For music where there is substantially random phase between the two channels, the sound will be both external and enveloping.

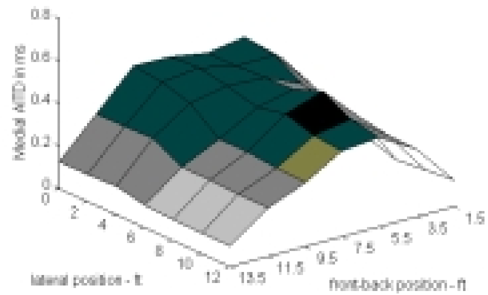

Figure 14: Medial AITD for the 12’x15’x9’ room over the 20Hz to 90Hz range. Single driver is in the upper left corner. This picture represents the lateral AITD if the room is set up with the primary listening axis parallel to the short walls, rather than parallel to the long walls. Notice that in this room setting the axis parallel to the short walls gives a much larger lateral AITD than setting it parallel to the long walls. The difference is highly audible.

11. PRESSURE AND AITD FROM A SINGLE DRIVER THAT IS NOT IN THE CORNER

In audio as in life there is no free lunch, but it is possible that by moving the driver to the side of the room we could increase the lateral AITD at the expense of the medial AITD. Figures 15 – 19 show that this works rather well. Putting the LF driver to the side causes much the same type of increase we saw in the 12’x15’ room when the listener faced the long wall. The low frequencies become external, and tend to localize in the direction of the driver. In practice this means the low frequencies shift from inside the head to the side of the room. Whether this perception will be preferred depends on your expectations. In practice, the sense of externalization is much stronger than the sense that the low frequencies are coming from the side. One is not particularly aware of where the low frequencies are coming from, but at least they are external.

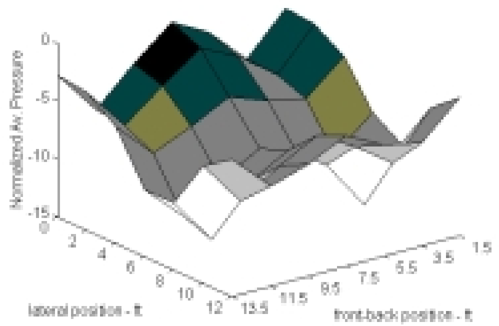

Figure 15: Normalized Average pressure in the same room, but the driver is now at the side of the room, at y=7.5', x=0.1', z=1/5' Although the driver is not in the corner of the room, the pressure in the listening area is not significantly reduced. See figure 11.

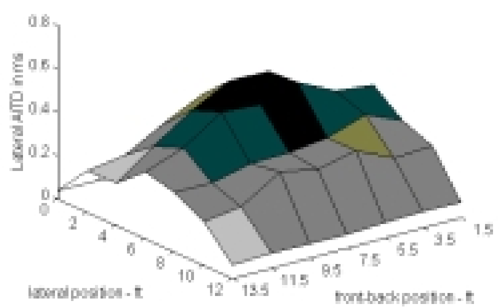

Figure 16: The Lateral AITD over a range of 20Hz to 180Hz in a 12'x15'x9' room with the driver at the side. Note the significantly higher values than for figure 13.

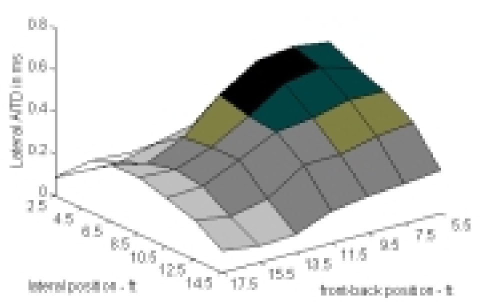

Figure 17: The Lateral AITD over a range of 20Hz to 90Hz in a 17'x23'x9' room with the driver in the upper left corner. See the pressure response in figure 10.

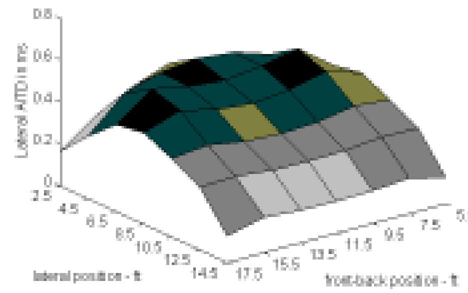

Figure 18: The Lateral AITD over a range of 20Hz to 90Hz in a 17'x23'x9' room with the driver at the side of the room in position x=0.1', y=11.5', z=1.5' Note the substantially higher values than for figure 17.

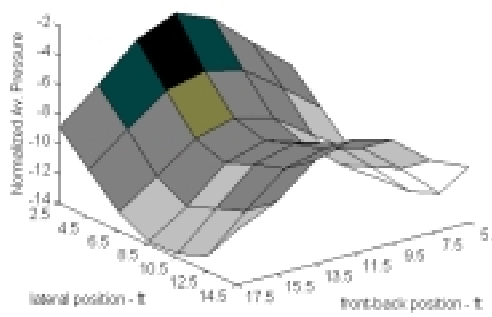

Figure 19: The Normalized Average pressure over the same room as figure 16. Note the pressure is not significantly lower than in figure 10, even though the driver is not in the corner of the room.

12. LATERAL ITD FROM TWO DRIVERS – APPARENT POSITION OF PHANTOM IMAGES AT LOW FREQUENCIES

If we have two low frequency drivers in the room there will in general be interference between the pressure produced by each driver. As mentioned before, if the signals to the two drivers are not correlated, this interference will be minimal. However, by long tradition almost all popular music is recorded so the low frequencies are highly correlated in the two stereo channels. The reasons are various. In FM broadcast when there is little correlation too much energy goes into the subcarrier, and in LP records the cutting stylus tends to lift out of the groove. Besides these technical reasons, usually the bass is louder if it is in phase, and most engineers think that louder is better in popular music.

There is another long tradition in stereo music recording, the phantom image. Recording engineers have long controlled the perceived azimuth of a sound source by adjusting the relative level of the two drivers. The most common pan law assumes that the apparent position of a sound image can be smoothly moved between the two loudspeakers by controlling the relative amplitude of the two speakers with a sine/cosine pan.

If p is a pan angle varying from 0 to 90 degrees, and A is a music signal, then

Left speaker = A*cos(p)
Right speaker = A*sin(p)

Reference [43] in the previous paper cites a considerable literature on the validity of this pan law, and demonstrates that at low frequencies the movement is not what is expected. One of the virtues of our room-head model is that we can investigate these pan laws.

We are interested in investigating how two sound sources respond when they are driven with various phases and amplitudes. Lets start with the two sources in phase, and investigate the effects of varying the amplitudes. What ITDs (and thus what perceived azimuth) are generated? We could answer this question for a number of points in the room, but for this paper we will do so only for the "ideal" position – at the vertex of an equilateral triangle which includes the two loudspeakers.

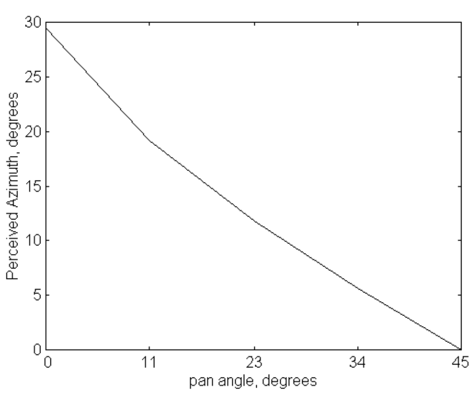

Figure 20: Perceived angle, as calculated from the ITD, for a listener at the ideal listening position, with loudspeakers +-30 degrees, as a sound is panned from left to center. Anechoic environment

Figure 20 shows the Net lateral AITD at the prime listening position in an anechoic room, as a pan law varies from p = 0 (full left) to p = 45 degrees (center). As expected, when p=0 the AITD has the value of sin(30)*0.75, which we plot at a perceived angle of about 30 degrees. As the sound pans the ITD decreases, and the sound appears to move smoothly to the center. There is a slight tendency for the perceived position to lie closer to the center than one would expect from the angle p, but the match is pretty good. (The match of this figure to the measured laws in reference [43] is extremely good.)

As a real sound source moves from left to center in an anechoic space, the medial AITD increases from sin(60degrees)*0.75ms, to 0.75ms. This is not the case as a phantom source moves. Our model shows that as a phantom source moves in an anechoic space, the medial AITD is constant, holding the value for full left pan. The symmetry of the loudspeaker layout enforces this non-intuitive result.

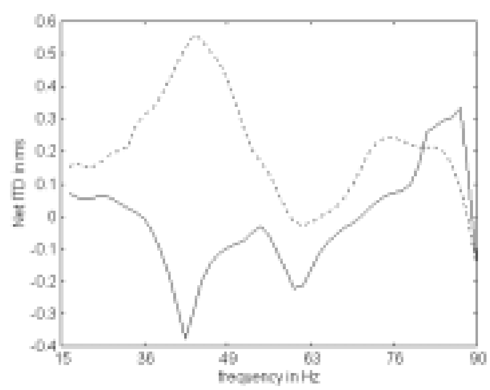

Figure 21: The frequency dependence of the ITD as a function of frequency for a single driver at x=.5', y=.1', z=1.5' for a listener at x=6', y=9.5', z=4' in a 12'x15'x9' room with reflectivitiy 0.8. ___ = lateral ITD. - - - = medial ITD. Note that for frequencies below 70Hz the lateral ITD is negative, and the speaker appears to be in the opposite corner.

When you perform the experiment in a reflective room, the result is drastically different. First of all, when p=0 (full left pan) the AITD is not necessarily equal to sin(30)*0.75. The room conspires to make the net ITD a strong function of frequency. Figure 21 shows the frequency dependence of the ITD for p=0 in a 12'x15'x9' room. Note that for frequencies below 70Hz the ITD is negative, which means the speaker is localized to the wrong side. The average absolute ITD – the AITD – is positive over the range of 20Hz to 90Hz, but the net AITD over the same range is near zero. The medial AITD is also highly frequency dependent. It seems that at 37 Hz it is possible to localize the sound to the front, but not at other frequencies.

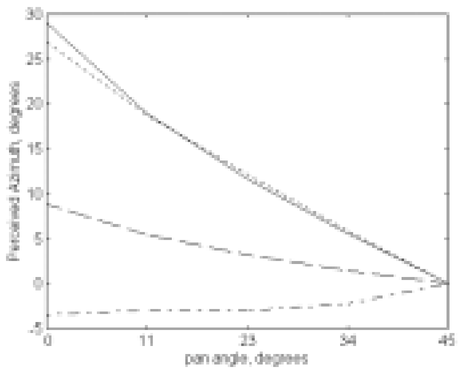

Figure 22: Pan law for four different room reflectivities. ___ = anechoic, - - - = 0.5, __ __ __ = 0.65, - . - . - . = 0.8. Note the ability to localize horizontally goes down rapidly as the reflectivity goes up. Same room and positions as figure 21.

Figure 22 shows the horizontal localization in this room for four different wall reflectivities. The net AITD is calculated over the frequency range of 20Hz to 90Hz. As can be seen, the ability to localize the low frequencies depends strongly on the reflectivity of the walls. Note that at a reflectivity of 0.8, the net ITD is low and to the opposite side.

It is not clear what these pan law diagrams mean. We would like to treat lateral ITD as the only determinate of azimuth. (It is simpler to ignore the medial ITD.) Our wishes are aided by the fact that the medial ITD is usually much lower than in the anechoic case. With our current understanding of perception this would indicate an “in the head” localization, and not necessarily a smooth shift in azimuth. In reference [43] of the previous paper small head movements were allowed – and the results suggested that sources tended to cluster toward the center as the sound panned across the room. It is possible that the subjects in [43] confused “in the head” localization with “in the center”. We conclude that the problem of pan laws at low frequencies is clarified by the NITD and AITD, but needs further research.

13. PRESSURE AND AITD FROM TWO DRIVERS AS A FUNCTION OF RELATIVE PHASE

In a classical recording with a lot of hall sound, or which was made with spaced omnidirectional microphones, the low frequencies are not in phase. The phase relationship will depend strongly on frequency, or be a semi-random function of time. Our calculation of the AITD when there are multiple drivers depends on knowing the phase relationship between the sources, and is thus not well suited to studying this case. The DFT is a better measure.

However, even with popular music where the low frequencies are almost always in phase, we can use electronic phase shift networks to give the drivers an arbitrary phase relationship. What is the effect of such a shift on pressure and AITD?

13a: Pressure and AITD from two drivers in an anechoic space

When there is a single driver to the side of a listener, the total and the lateral AITD in an anechoic space will be constant, at 0.75ms. We can add a second driver on the opposite side of the listener, using the same amplitude, but with variable phase. Now instead of having a running wave moving across the listener, we have created a standing wave. If we move the listener laterally between the two sources we will measure peaks and valleys in the pressure, and as we vary the phase of the drivers, the positions of these peaks and valleys will shift. When the drivers are in phase we will have a peak at the center. When the phase is reversed, a null will appear in the center. Intermediate phases will give intermediate positions for the peaks and valleys, but will not eliminate them.

The AITD will also vary with phase. When the two drivers are in phase in a symmetric room, a listener at the center will perceive an AITD of zero for all frequencies. It is equally clear that when the drivers are out of phase the ITD will be high, although they may not be audible because of the lack of pressure.

However we need not choose to have the drivers either in phase or out of phase. For example, we expect that a 90 degree phase shift will reduce the center-line pressure by 3dB compared to the in-phase case. What happens to the AITD?

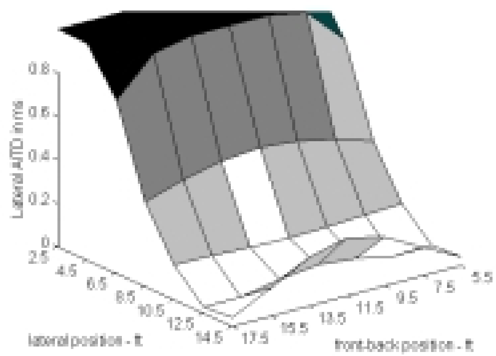

Figure 23: Lateral AITD in an anechoic space, two drivers on opposite sides of the listener, separated by +-11'. Left driver leads the right driver in phase by 90 degrees. 22.5Hz. Note the peak in AITD at -4' from the center, and the minimum at +4'. These values correspond to the minimum and maximum in the pressure response. The value of AITD along the center axis (8.5') depends on the spacing of the minimum and maximum.

Figure 23 shows surface plots of the AITD for a case where there are two drivers on opposite sides of an anechoic space, separated by 22'. An area in the center of +-6' is plotted. The frequency chosen is 22.5 Hz. Note the peak in AITD at the position of the pressure null in the room, and a minimum value of AITD at the position of the pressure maximum. The value at the center of the room – which is intermediate between the positions of these peaks and nulls, must lie between these two values. Because the positional dependence is not linear, the center value depends on how closely the maximum and minimum are spaced from each other.

Figure 24 shows the dependence on frequency of the AITD in the center of the space. Note that at the frequencies of most interest to us, the value is 0.2ms or below. This value is much better than the AITD we would get without the 90 degree phase shift, but is rather low for the purposes of externalization. Fortunately when we look at real rooms the improvement with the phase shift is larger, particularly if we integrate over a wide frequency range. We find that using a phase shift greater than 90 degrees will increase the AITD, at the cost of pressure. Higher shifts than 90 degrees also increase the risk that at some frequencies the ITD will be higher than natural values.

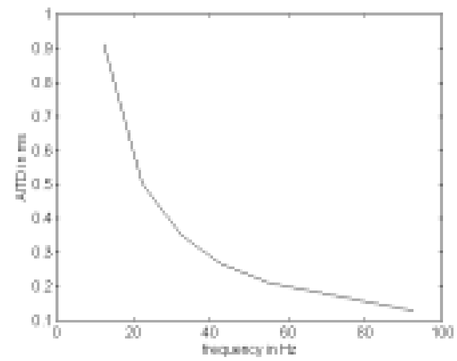

Figure 24: The dependence of the AITD in the center of figure 22 with frequency. Note the typical values of 0.2ms or so above 60Hz. Adding surface reflections can increase this value.

13b Pressure and AITD from two drivers in reflective spaces

First, it is obvious from symmetry that just as in the anechoic case driving two loudspeakers in phase in a symmetric room will cause the lateral AITD to be zero along the center line. Since the medial AITD is also likely to be low, "in the head" localization is almost guaranteed. One way of understanding this result is to realize that all asymmetric lateral room modes must be suppressed, and the asymmetric lateral modes are the only ones capable of producing a lateral ITD along the centerline of the room.

The situation is not improved by driving the low frequencies out of phase. The drivers now will excite only the asymmetric lateral modes. All symmetric lateral modes, all front/back modes, and all up/down modes will be suppressed. The AITD will be higher than the AITD in natural hearing, producing a perception some people refer to as "phasyness". Phasyness can make some people distinctly uncomfortable, and in extreme cases, even nauseous. Clearly we don't want it.

If we run the drivers with a constant 90 degree phase shift: 1. All up/down and front/back modes will be allowed, but their amplitudes will be reduced by 3dB compared to the in-phase case. 2. Asymmetric lateral modes will be allowed. 3. Symmetric lateral modes will also be allowed. 4. If constant phase is a reality, there will be no nulls in the pressure response, because where a symmetric lateral mode has a pressure minimum, the asymmetric mode will have a maximum, and the two will not interfere, because there is a 90 degree phase shift between them!

The prospect of no nulls in the lateral standing waves seems too good to be true, and it is. As we will see, when the wall reflectivities are in the range of 0.8, a 90 degree phase shift leading to the right causes a dip in the pressure on the right side of the room. The reduction in pressure is only 3dB along the central axis, as expected. However there is an improvement in the AITD in the listening area, and this improvement is audible.

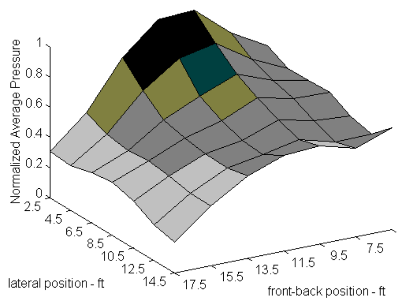

a:

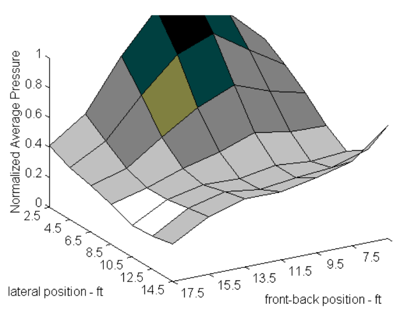

b:

Figure 25: Normalized Average Pressure in a 17’x23’x9’ room in the 63Hz octave band, from two drivers along the front wall. Graph a. shows a 30degree phase shift, graph b. shows a 90 degree phase shift. Notice that a dip in the pressure occurs to the right of the central axis.

Figure 25 shows the pressure in the center of a 17’x23’x9’ room from two drivers symmetrically placed in the front of the room, at 8.5’+-5.5’.

Figures 26-29 show different aspects of the AITD in this room, with various speaker placements. They tell their own story. In general we can say that using the 90 degree phase shift produces significant increases in the lateral AITD in the listening area. There is an additional improvement in both average pressure and in AITD when the low frequency drivers are moved to the sides of the listening area.

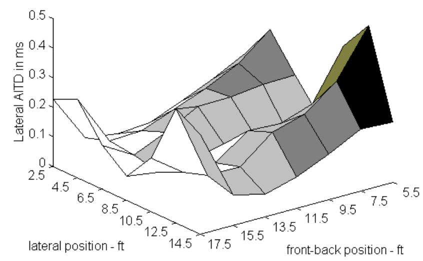

a:

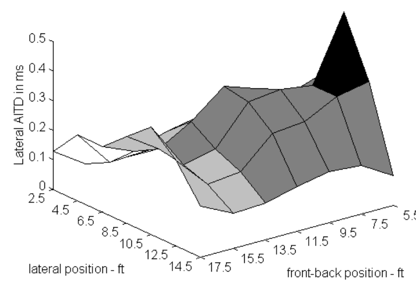

b:

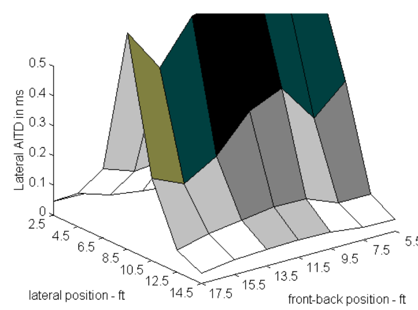

c:

Figure 26: Absolute AITD in the 63Hz octave band for the same room as figure 24. Graph a. is 30 degree phase shift, graph b. has a 90 degree phase shift, and graph c. has a 150 degree phase shift. Note the very low AITD in graph a., a moderate AITD in graph b., and an unnatural (phasey) AITD in graph c.

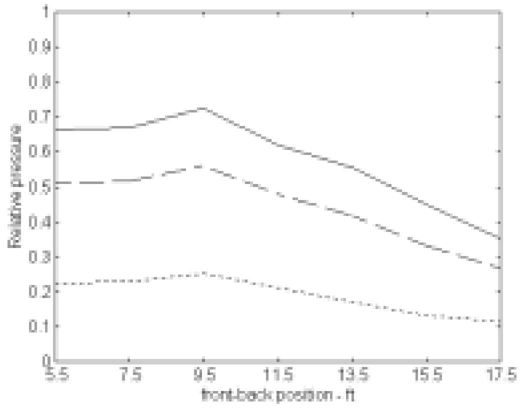

Figure 27: The Normalized Average Pressure along the center axis of the room from figure 25. ____ = 30 degree shift, __ __ __ = 90 degree shift, - - - - =150 degree shift. Note the ~3dB reduction in pressure in the 90 degree curve.

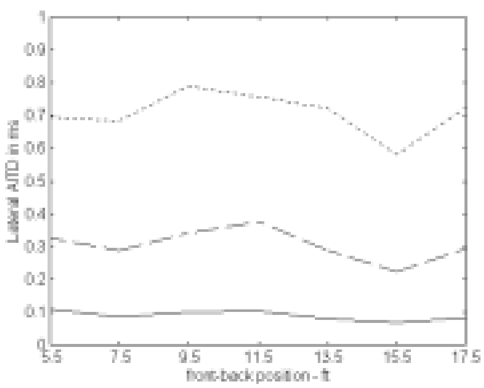

Figure 28: The lateral AITD in the 63Hz octave band along the center axis of the same room. ___ = 30 degrees, __ __ __ = 90 degrees, - - - =150 degrees. Note the very low AITD for the 30 degree case, and the moderate AITD for the 90 degree case. The 150 degree case is unnaturally large and phasey.

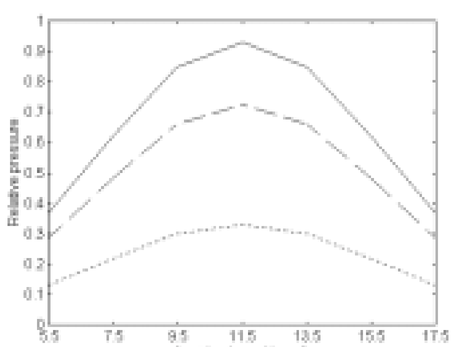

Figure 29: The Normalized Average Pressure in the 63Hz octave band in the same room as figure 24, but with the drivers at the sides of the room at y=11.5'. Receiver is along the center axis. ___ = 30 degree phase shift, __ __ __ = 90 degree phase shift, - - - =150 degree phase shift. Note the pressure in the listening area is higher than with the drivers in the front of the room.

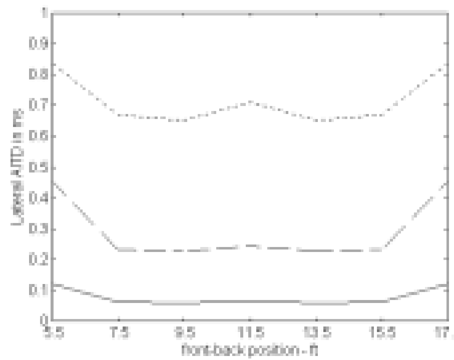

Figure 30: The Lateral AITD in the 63Hz octave band for the same configuration as figure 28. The AITD for the 90 degree shift is again moderate, the 30 degree case is much too low, and the 150 degree case is too high. Note that for this frequency band the AITD for 90 degree phase is approaching the anechoic value. The value below 50Hz and above 70Hz is ~0.4ms.

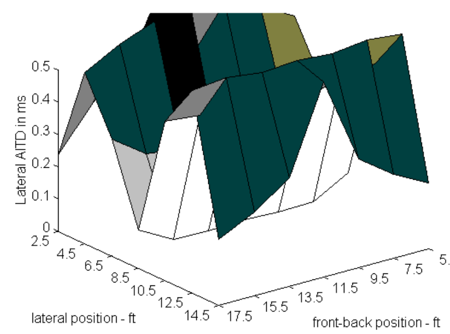

a:

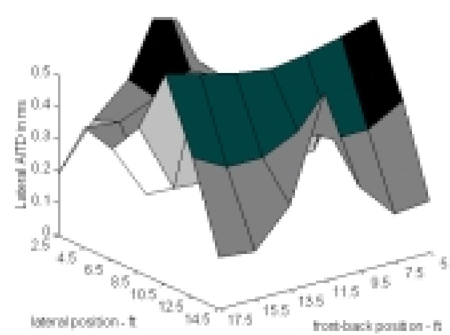

b:

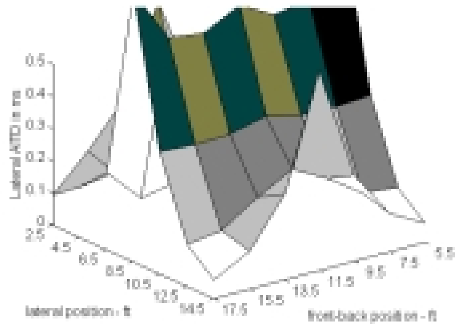

c:

Figure 31: Lateral AITDs in the listening area from two drivers at the sides of the room. Graph a. is for 30 degrees shift, graph b. for 90 degrees, and graph c. is for 150. Note there is a minimum slightly to the left of the center line in graph b.

14. PRESSURE AND AITD FROM FOUR DRIVERS.

There is an additional improvement when four drivers are used. Figures 32-34 show the same 17'x23' room, but with low frequency drivers both in front of the listening area, and at the sides. This configuration corresponds to a 4 or 5 channel surround system with full range drivers. The figures show that we have good results when the low frequencies in the front drivers and the corresponding side drivers are in phase, with a 90 degree phase shift between the left and the right sides of the room.

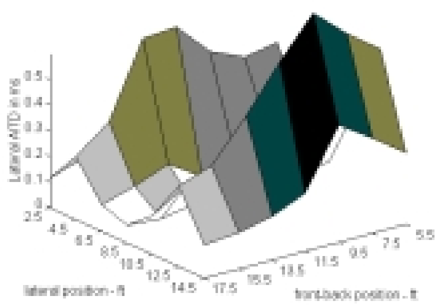

a:

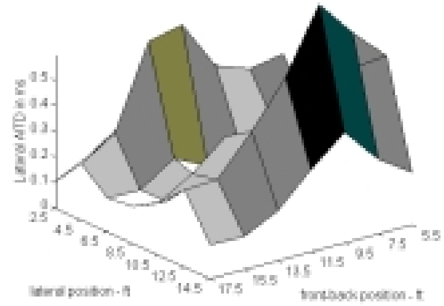

b:

Figure 32: Lateral AITD's in the listening area from 4 drivers, two in the front at +-5.5', and two at the sides at +- 8'. All reflectivities are 0.8. 63Hz octave band. Picture a. is for 30 degree shift, picture b. is for 60 degree shift.

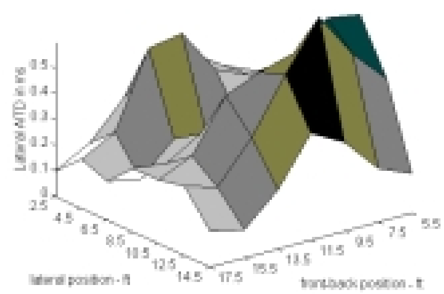

a:

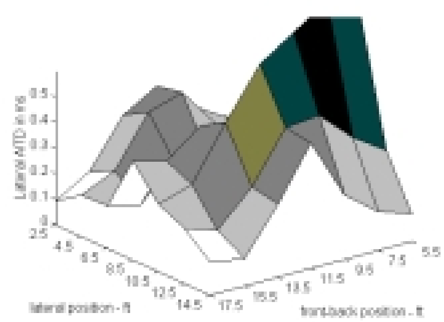

b:

Figure 33: Same configuration as figure 32, but a. is 90 degrees, and b. is 120 degrees.

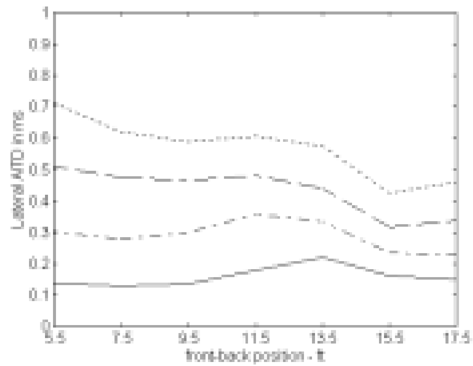

Figure 34: AITD in the range of 30-90Hz along the center line of a 17’x23’x9’ room with four drivers, the right side of the room lagging the left side by a variable phase shift.
____ = 30 degrees, __ - __ = 60 degrees, __ __ __ = 90 degrees, and - - - = 120 degrees.
The increase in externalization at 90 degrees is highly audible.

15. CONCLUSIONS

This paper shows the use of two new methods of evaluating the sound of a room, the AIDT and the DFT. Both measures are sufficiently new and untested that it is difficult to make firm conclusions about what they seem to show. In our limited experience with the AIDT, the measure seems chiefly useful below 100Hz, as a predictor of the degree to which a particular room and loudspeaker configuration will cause sound to be localized outside the head of the listener. Although the model has not been used in many rooms, the electronic circuit based on the model – the 90 degree phase shifter for frequencies below 120Hz – has been tried in several rooms. The improved externalization is highly audible.

The AITD modeling shows that in general it is advantageous to use more than a single low frequency driver, and it is useful to locate these drivers to the side of the listener.

The same conclusion comes from the work on envelopment using the DFT measure. In this case however the improvement achieved from multiple drivers is highly dependent on the reflectivity of the room. When noise is used as a signal source and the lateral reflectivity is 0.8 or more, the overall envelopment depends primarily on the room, and not on the decorrelation in the recording. When the lateral reflectivity is below 0.65, there can be a large improvement in envelopment with two uncorrelated drivers, particularly if they are at the sides of the listeners.

With music as a sound source we find that for many types of bass instruments there is a substantial advantage to stereo low frequency loudspeakers, even in reflective listening rooms. The DFT with a narrow band test signal – or with actual reverberation from a musical source – can be used to quantify the difference, and to find optimal loudspeaker positions. Once again, there appears to be an advantage to placing the low frequency drivers at the sides of the listening area.