Speaker placement, externalization, and envelopment in home listening rooms

David Griesinger

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

Abstract

The ideal number and placement of low frequency drivers in small listening rooms has been controversial. Most research has assumed listener satisfaction is determined by the sound pressure as a function of frequency and source-listener position. We believe two additional properties of the soundfield, externalization and envelopment, contribute to listener preference. We propose mathematical methods for quantifying these two perceptual properties given a measured or calculated binaural impulse response. The Average Interaural Time Difference (AITD) is our measure for externalization, and the Diffuse Field Transfer function (DFT) is our measure for envelopment. An image model for small rectangular rooms is used to predict the values of pressure, AITD, and DFT for different room properties and driver locations.   It is found that the low frequency pressure uniformity, the AITD, and the DFT 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. Listening tests confirm the results of the calculations.

1. Introduction

In the best concert halls and opera houses low frequency sounds envelop the listeners. Although one is aware that the attack of the kettledrums come from the stage or the pit, the ring of the drum and the rumble of the bass drum come from all around the hall. The bass viols and the cellos have the same property, particularly when they play pizzicato. One of the joys of an organ concert is hearing the bass swirl around the cathedral when a pedal note is held. When the acoustics produce envelopment music has a living quality that is highly prized by conductors and players. When recorded music is played through loudspeakers, envelopment can often seem adequate at frequencies above 1000Hz, but

poor at lower frequencies. In fact, many recording engineers seem to be unaware that low frequency envelopment is either possible or desirable.

Envelopment at higher frequencies can also play unexpected tricks. Normally the sound image from a conventional stereo system stays fixed between the two loudspeakers. But this is not always the case. Occasionally sounds seem to surround the listener, even in a non reflective room.

Listening rooms also suffer from a perceptual anomaly that has no obvious counterpart in performance spaces. Low frequency instruments in popular music, such as the kick drum and the bass guitar, are almost always perceived as coming from inside the head. This perception does not occur in the concert venue, even when these instruments are amplified. This "in the head" localization is unique to recorded music. It is always perceived as artificial by the author. In this paper we will use the word "externalization" to describe this perceptual property.

Both externalization and envelopment depend strongly on the recording technique, but they appear to be independent of each other. In rooms where low frequency envelopment is perceptible, low frequency instruments in classical music are often perceived as external, while low frequencies in popular music are often "in the head".

Both envelopment and externalization are highly dependent on properties of the room. Years ago we noticed that it is possible to perceive low frequency envelopment in some home listening rooms, and not in others. In [43] we attempted to study envelopment through measurements of localization. We noted that in many listening rooms it was possible to localize low frequencies to a particular loudspeaker, but phantom images were unstable. Panning low frequencies between two loudspeakers did not yield the same positional dependence as is noted at higher frequencies. The phantom image tended to pull to the center of the listener's head – and be judged as closer to the center than it was intended to be.

In [43] we found that the apparent position of the low frequency sound could be brought more into alignment with the high frequency sound if the separation at low frequencies was increased electronically. The circuit, dubbed a "spatial equalizer" has become popular with many engineers. However the primary virtue of the spatial equalizer for these engineers turned out not to be improved localization, but enhanced envelopment. The spatial equalizer works by increasing the left minus right (L-R) component of the sound below 300Hz. In most listening rooms the improvement in localization is subtle, but the improvement in envelopment is obvious.

In the succeeding years we noticed that this circuit is completely inaudible in some rooms. Ironically, in most sound mixing studios the circuit is inaudible because the antiphase component of the low frequency sound cannot be heard. These rooms typically have a high degree of symmetry and carefully control the low frequency reverberation time. The widespread use of such rooms for sound mixing has had at least two undesirable effects. These rooms tend to make professional engineers unaware that low

frequencies can be enveloping. These rooms also encourage the use of microphone techniques that enhance imaging at the expense of envelopment. For example, recording techniques that utilize only closely spaced omnidirectional microphones (such as most binaural techniques) produce excellent imaging with earphones, but poor low frequency envelopment with loudspeakers. If your room does not permit you to hear envelopment, you will not know what you are missing.

The externalization of low frequencies is even more mysterious. Typical home stereo systems often externalize low frequencies, whereas symmetrical listening spaces are the most likely to sound artificial. These rooms are often described by their owners as possessing "tight" low frequency imaging. To my ears the low frequencies are centered, but they are unlike anything one would hear in a concert.

So we have at least three mysteries to untangle. First, why do some rooms support low frequency envelopment, and what can be done to provide it in rooms that do not? Second, why do the kick drum and the bass guitar almost always end up banging away inside your head and what can we do to get them out? Third, why do some recordings sound enveloping even when you listen to them through two front loudspeakers in a relatively non reflective room?

2. Envelopment in concert halls

The study of envelopment in concert halls has been marked by contradictions between common observation and accepted theory. In [44] we outline a theoretical framework that resolves these discrepancies. The framework has the following major parts:

1. Envelopment at low frequencies is perceived when the interaural time delay fluctuates at a rate of between 3Hz and 20Hz. Above 400Hz fluctuations in the interaural intensity difference (IID) and fluctuations in the ITD are both important.

Below 400Hz the interaural time delay is the principle cue for localizing the horizontal direction (azimuth) of low frequency sounds. In the absence of reflections, the ITD determines azimuth with high accuracy – within a few degrees at frequencies of 500Hz and above. Below 500Hz the accuracy is proportional to the frequency, so that localization to +-20 degrees is still possible in the 63Hz octave band.

Lateral reflected energy causes the ITD – and thus the perceived azimuth – to shift. When the sound source is broad-banded, or consists of a musical tone with vibrato, the shift in ITD becomes a fluctuation. For sources of both speech and music the fluctuation is essentially random (or chaotic) in nature. Fluctuations at rates slower than 3Hz are perceived as source motion. Above this frequency they are perceived as envelopment.

2. Where the sound source consists contains rapid attacks – such as the start of a speech phoneme, or the attack of a musical note – the onset of the sound at the listener is uncorrupted by reflections. In this case the ITD during the attack accurately determines the sound direction, and later fluctuations produce envelopment. Thus it

3

is possible to have both good localization (a low apparent source width) and high envelopment at the same time. At higher frequencies the IID performs a similar role, and localization is determined by both IID and ITD.

3. For speech or music that is relatively transparent to reverberation, the fluctuations are maximal during the pauses between phonemes or notes. The loudness of the reverberation during these pauses determines the degree of envelopment. The importance of the reverberation during pauses is increased by the formation of the background sound stream. Where the background stream is easily separated from the foreground stream, envelopment is ~6dB stronger for a given direct to reverberant ratio.

4. For thickly orchestrated music where such pauses are rare, the overall amount of fluctuation determines the degree of envelopment. This is the case when the source is relatively continuous, such as massed strings, or a pink noise test signal.

We wish to add a fifth observation, that came to the author's attention after reference [44] had been published:

5. In carefully controlled loudspeaker experiments by Morimoto's group in Kobe, it has been found that reflections that come from behind the listener are more enveloping than reflections that come from the front. They have proposed measuring envelopment through the front/back ratio.

Statement number 5 indicates that the envelopment perception does not arise solely from interaural fluctuations. Where it is possible to discriminate between front and rear sound, rear sound is perceived as more enveloping. We know from the physics of sound perception that it is possible to discriminate front sound from rear sound in two ways. At low frequencies the two can be discriminated through small movements of the listener's head, that cause predictable changes in the ITDs. Unfortunately head movements produce no shifts in ITDs when the sound field is largely diffuse – this localization cue works only with sound that is well localized. At high frequencies it is possible to discriminate sound that comes from behind by notches in the frequency spectrum at about 5kHz. These notches are present when sound comes from greater than about 150 degrees from the front.

We conclude that the front/back ratio is primarily important for sound that includes a substantial amount of energy above 3kHz.

The statements given above outline the relationship between the physical sound field and the perception of envelopment. Within this framework there are several unavoidable difficulties. For example, it is plain from statements three and four that the perceived degree of envelopment will depend both on the loudness of the sound source, and the type of music being played. This effect is easily observed, and it does not make it easy to develop reliable measures. Statement number 5 implies that the spectrum of the source

and the spectrum of the reflections will make a significant difference to how envelopment is perceived.

For frequencies below 1000Hz the perception of envelopment depends on maximizing the fluctuation in the interaural time delay at the listener's ears. For music that is relatively transparent to reverberation it is the reverberant component of the perceived sound that creates envelopment. For more continuous music, it is the total fluctuation in the ITD that counts.

For frequencies below 1000Hz in small listening rooms the reverberation time is usually sufficiently short that the room itself is unable to develop significant fluctuations in the ITD in the pauses between syllables or notes. Thus creating the perception of envelopment in the background stream depends on how the room acoustics interact with the material (particularly the reverberation) on the recording. The loudspeakers and the room form a transfer system, which ideally can transmit the enveloping properties of the recording to the listener. Our job is to find a way of measuring the effectiveness of this transfer system, and then to find how the transfer can be optimized.

At frequencies above 1000Hz the front/back ratio will have important implications for our choice of loudspeaker positions.

3. Envelopment at high frequencies

In the frequency band between 300Hz and 3000Hz envelopment is determined by a combination of fluctuations in the IID and the ITD. This fact was discovered independently by the author and Blauert. Blauert notes in [7] and [8] that fluctuations in IID have slightly different perceptual properties than fluctuations in ITD. He concludes that different brain structures are involved in the two perceptions. Envelopment at frequencies below 300Hz is determined almost entirely by fluctuations in the ITD. The important virtue of the framework presented in [44] for the perception of envelopment is that the concept of interaural fluctuations takes the study of envelopment out of the realm of psychoacoustics and into the realm of physics. We can model the mechanisms that create these fluctuations, and we can measure them.

Once we have moved the problem into the realm of physics, we can see that the problem of envelopment has two parts – the nature of the envelopment receiver, and the nature of the sound field surrounding the receiver.

The receiver for envelopment is the human head, the outer and inner ears, and the brain. The soundfield is the combination of direct and reflected energy at the listening position. The head/ear system is the antenna for the perception of envelopment. Like other antenna, this system has directional properties. We can study the directional dependence of the head/ear system for single reflections by convolving test signals with published Head Related Transfer Functions (HRTFs) such as the ones put on the web by Bill Gardner. After the convolution we can detect and measure the interaural fluctuations that result using the software detectors that will be described later in this paper.

This work is yet to be done. However several years ago the author did models of the high frequency behavior of interaural fluctuations using a simple head model. We found that the resulting fluctuations depend on the azimuth of the reflection. For 1000 Hz the optimum angle lies in a cone centered on a line drawn through the listener's head. This cone makes an angle of about 60 degrees from the front of the head. At about 1700Hz the cone includes the standard loudspeaker positions at +- 30 degrees from the front. This data agrees with subjective experiments into the angular dependence of ASW and spaciousness reported by Ando.

(Ando and Morimoto use the interaural cross correlation (IACC) as one of their measures of the spatial properties of concert hall sound. Morimoto uses the IACC of the first 80ms of a binaural impulse response as a measure. Beranek and Hidaka call this measure (IACCe). For reasons described extensively in [44], the author finds the measure often misleading, particularly in small rooms. However the IACCe does describe the high frequency angular dependence of envelopment quite well.)

Using interaural fluctuations it can also be shown that if there are two sound sources in an anechoic space, and uncorrelated band-limited noise is played through each source, the sense of envelopment will be maximum when both sources are at the optimal angle for the frequency band in question. For frequencies below 700Hz the optimal angle is 90 degrees – the sources must be at the sides of the listener. As the frequency rises, the optimum position moves toward the medial plane.

The angular dependence of envelopment at high frequencies has important consequences for sound reproduction systems. A few years ago the author was playing a stereo tape of applause through a standard two channel system in a relatively dead room. He was very surprised to hear the applause coming from all around him, even though the loudspeakers were clearly in front. Adding a bandpass filter quickly showed that the perception was a simple result of the angular dependence of high frequency envelopment. The 1700Hz band produced a surrounding perception, even though the loudspeakers were at +-30 degrees. Because the applause had significant energy in this band, the result was enveloping. We conclude that stereophonic reproduction works in part because envelopment in a particular frequency band can give an impression of envelopment in all bands. Even with a relatively narrow front loudspeaker pair, envelopment at some frequencies will be high.

Morimoto's front/back ratio becomes important when the reflected sound contains significant energy above 3000Hz. In the author's experience with concert halls and opera houses it is rare that there is significant energy at those frequencies coming from the rear. Sound from the rear of a hall is often spatially diffused, and has been reflected off many surfaces. Typically each surface takes a toll on the high frequency content. In halls where an electronic enhancement system has been installed it is possible to experiment with the frequency spectrum of the reflected energy, and invariably the sound is better if frequencies above 3000Hz are attenuated. Thus the author is not convinced that the front/back ratio is an important measure for concert hall design.

However Morimoto has shown conclusively that the front/back ratio is important when listening with loudspeakers, and the author can confirm the observation. This issue becomes quite important in the design of surround sound playback systems. Several of the best sound mixers for surround sound have noted that they prefer the rear loudspeakers to be placed +- 150 degrees from the front, and not the more typical +-110 degrees. At 150 degrees these loudspeakers are capable of producing the frequency notches in the HRTFs that indicate rear sound, and speakers at 110 or 120 degrees cannot. The author is convinced that loudspeakers at 150 degrees produce a more exciting surround sound picture. A loud sound effect from 150 degrees from the front is much more exciting than one from 120 degrees.

In addition, one of the most effective uses of a surround sound system for popular music is the presentation of a live performance, where the applause and noise from the audience draws the listener into the experience. Applause has substantial high frequency content, and the front/back ratio becomes important.

However, what is good at high frequencies is not necessarily good at low frequencies. As we will see, for optimal envelopment at low frequencies the surround loudspeakers should be at the sides of the listening area. We might conclude that a single pair of surround loudspeakers is not sufficient – although if we are clever there may be ways around this problem.

4. Envelopment at frequencies below 1000Hz in small rooms – the DFT

Envelopment below 1000Hz depends on interaural fluctuations, and on the ability of human perception to separate sound into a foreground and a background stream. Where this separation is possible, it is the spatial properties of the background stream – specifically the amount of interaural fluctuations – that determine the envelopment we perceive. In a small room there is almost never sufficient late reflected energy to contribute to the background perception. The late energy must be supplied by reverberation in the original recording. To predict the degree of envelopment we perceive we must be able to predict the strength of the interaural fluctuations during the reverberant segments of the recording.

Recording engineers know that for best results the reverberation in a two channel recording should be uncorrelated – completely different in the left and right channels. It is the job of the loudspeaker/room system to cause the listener to have adequate interaural fluctuations when this condition occurs. The loudspeaker/room system is acting as a transfer system, transferring the decorrelation in the recording to the listener's ears. We need a measure for how effectively this transfer works.

Finding a test signal for measuring envelopment:

This problem turned out to be much more difficult than expected. Much of the difficulty lies is choosing a test signal that will adequately represent music. As mentioned in the

section on concert hall acoustics, the perception of envelopment is highly influenced by the presence of gaps in the music that allow reverberation to be heard. In this study we will assume that such a gap has already occurred – we are trying to model the transfer function of the reverberation within that gap.

The correlation time of musical signals:

However music has another very interesting property that is highly relevant to this study. Music generally consists of notes – segments of sound that have a recognizable pitch. Although the notes may be rich in harmonics, the fundamental frequency is often steady.

If we autocorrelate a musical signal with itself we will see that over a certain length of time the autocorrelation function is strongly non-zero. That is, as long as a note continues there will be a strong fluctuation in the autocorrelation function. Ando has published material that found that the average length of the non-zero region of the autocorrelation function was related to the type of music being played, with modern serial music having a short correlation time, and romantic symphonies having a long correlation time. Such a result would be expected from the average length of the notes in the various musical styles, as well as the presence in many types of modern music of percussive sounds with no well established pitch. Ando discovered that as the correlation length of a musical example increased, the ideal reverberation time of a performance space also increased. We will see that this result can be predicted from the properties of the DFT.

When we started working for a measure of envelopment in rooms we used band limited pink noise as a test signal. Reference [44] contains several experiments and observations of the spatial properties of band limited pink noise. The correlation length of band limited noise depends on the bandwidth chosen (and to some degree on the filter type.) For our first work we choose to use a bandwidth approximately equal to the width of a critical band in the human ear. The results showed that this was not a good choice. Although the results appeared to be accurately reflect the envelopment of noise signals in small rooms, they did not predict the envelopment of musical signals.

As an experiment, we measured the frequency width and phase fluctuation in the reverberation from musical notes of various lengths in a real concert hall (Boston Symphony Hall.) The resulting reverberation had high coherence. We were able to use the reverberation in Boston as a test signal in calculating the DFT, and we found that to achieve a similar results with noise we had to use a filter bandwidth of 1-2Hz. A disadvantage of such a narrow frequency is that the DFT in a small room is often highly frequency dependent, and to find an average value of envelopment one must separately calculate the DFT at many different frequencies. Unfortunately there appears to be no short-cut. We must calculate our measure using a narrow band test signal.

The reason we need a narrow band test signal is obvious in hindsight. Small rooms have relatively short reverberation times, often less than 0.5 seconds. The time constant of such a space (the time it takes the sound to decay by 1/e) is the reverberation time divided by 7, or 70 milliseconds. If the music (or the reverberation from the music) has a

correlation length that is significantly longer than this time constant the room does not generate interaural fluctuations directly. This means that a single driver in the room will not produce interaural fluctuations at the ears of a listener. The room can however detract from the transfer of fluctuations from multiple drivers, and this is the effect we wish to measure.

A detector for envelopment at the listening position

A detector for envelopment is also a difficult problem. We had hoped to be able to use a measure developed for concert halls, the interaural difference, or IAD. To make a long story short – it doesn't work. We were in fact unable to find a proxy measure that could even duplicate the perceived envelopment in an anechoic space, let alone the envelopment in a reflective space. For example, it is obvious that a single sound source in an anechoic space is incapable of producing envelopment. Our measure must show that this is the case. By similar reasoning, two sources driven by highly correlated material should also give near-zero envelopment in an anechoic space, and the measure should reflect this. The IAD fails on both counts.

(As luck would have it, the IACC measure does distinguish between echoic and anechoic spaces. However – again making a long story short – the IACC fails to predict observed results at frequencies below 200Hz.)

We found it was necessary to go back to first principles. The hypothesis in [44] predicts that the perception of envelopment at low frequencies depends on fluctuations in the ITD. There is no shortcut – to measure envelopment we must convolve the binaural impulse response from each sound source with an independent signal, and then measure the fluctuations in the ITD that result. In practice this is difficult. Evolution has had millions of years to perfect methods of extracting ITD information from noisy signals at the eardrums. We had to develop an algorithm in software that make this extraction. Our current measure still needs some work, but it seems to give useful results. Using this algorithm and a narrow band noise signal as a test probe we developed our measure for envelopment in small rooms. We call it the DFT, or diffuse field transfer function.

The process of finding the diffuse field transfer function can be summarized:

1. Calculate (or measure) separate binaural impulse responses for each loudspeaker position to a particular listener position. A high sample rate must be chosen to maintain timing accuracy. In our experiments 176400Hz is an adequate sample rate.

2. Low-pass filter each impulse response and resample at 11025Hz, and then do it again, ending with a sample rate of 2756Hz. This sample rate is adequate for the frequencies of interest, and low enough that the convolutions do not take too much time.

3. Create a test signals from independent filtered noise signals. Various frequencies and bandwidths can be tried, depending on the correlation time of the musical signal of interest.

4. Convolve each binaural impulse response with a different band filtered noise signal, and sum the resulting convolutions to derive the pressure at each ear.

5. Extract the ITD from the two ear signals by comparing the positive zero-crossing time of each cycle.

6. Average the ITDs thus extracted to find the running average ITD. The averaging process weights each ITD by the instantaneous pressure amplitude. In other words, ITDs where the amplitudes at the two ears is high count more strongly in the average than ITDs where the amplitude is low.

7. Sum the running average ITD and divide by the length to find the average ITD and the apparent azimuth of the sound source.

8. Subtract the average value from the running average ITD to extract the interaural fluctuations.

9. Filter the result with a 3Hz to 17Hz bandpass filter to find the fluctuations that produce envelopment.

10. Measure the strength of these fluctuations by finding the average absolute value of the fluctuations. The number which results is the Diffuse Field Transfer function, or DFT.

11. Measure the DFT as a function of the receiver position in the room under test.

The most difficult part of this process is building the ITD detector in software. The detector must be robust. The signals at the ears are noise signals – in many places the amplitude is low, and the zero crossings can be highly confused. Our detector should use very simple elements – just timers and filters – to do the job. It must be very difficult to confuse. The design of this detector is beyond the scope of this paper. Persons interested in its design, or in the Matlab code for the whole DFT measurement apparatus, should contact the author. In the current version of the code, it takes about 15 seconds to find the DFT at a single receiver position, so a 7x7 array of positions can be calculated in about 12 minutes (on a laptop with a 150MHz Pentium.)

The current code uses noise signals with 10000 samples each. This is about 3.7 seconds at the sample rate of 2756Hz. We are interested in random fluctuations in a signal in the 3Hz to 17Hz band – with typical maximum energy around 5Hz. Needless to say, the accuracy of such a measurement given a 3.5 second duration is not high. The DFT we get from this measurement has a semi-random variation of about 2dB, with occasional

values as many as 3dB different from what was expected. Thus the DFT we present here is somewhat noisy. We could improve the accuracy by increasing the length of the noise signals – gaining the usual improvement by a factor of 1.4 for each doubling of the signal length.

5. Design of a system for calculating the DFT

With a single driver in an anechoic space the DFT should be zero. In practice, in spite of our limited length of noise, the values we get are at least 40dB less than the maximum values with two drivers. Thus our detector passes this test.

5a. Bandwidth of the test signal for noise:

There are two major adjustments to the detector that we must set as best as we can to emulate the properties of human hearing. The most important of these is the bandwidth we choose for the noise signal. When we want to study the envelopment of noise signals in a room we would like a noise signal with the bandwidth and filter shape of a single critical band on the basilar membrane. In the Matlab code we use a sixth order elliptical bandpass filter, similar to the ones in a sound level meter. We need to choose the bandwidth.

In an effort to calibrate the DFT detector, a series of experiments on the envelopment of low frequency noise signals was performed. It is possible to probe the properties of the human envelopment detector through experiments with single lateral reflections, and with multiple lateral reflections. We are interested in how the envelopment impression depends on the delay of single reflections, or the combination of delays in multiple reflections. The apparatus described in [44] was used, with continuous band filtered pink noise as a source. The results were highly interesting.

First, (with a single subject) we found that the envelopment from a single lateral reflection depends on the delay of the reflection. There is an interference effect. For frequencies in the 63Hz octave band a single lateral reflection at 5.5ms delay produces a very wide and enveloping sound field, with relatively low sound pressure. A delay of 13ms produces a nearly monaural impression, with little or no envelopment at all. As the frequency rises the envelopment goes through one more cycle, becoming first super wide, and then somewhat less wide. Beyond 20ms all delays sound about the same.

This interference behavior arises from easily calculated cancellation between the direct sound and the reflection. Such interference is not possible when the delays are greater than the coherence time of the noise signal, and this depends on its bandwidth. Thus the properties of the basilar membrane filters can be studied through the interference effect. We find that at least at 63Hz the basilar membrane can be modeled by an elliptical filter of one octave width. If the basilar membrane was significantly sharper than one octave at 63Hz, we would expect the interference effect to extend to greater delays. If you use a ½

11

octave filter in the DFT detector you find that the interference effect with a single reflection does extend to higher delays. A one octave filter at 63Hz seems about right.

As an aside – we also found to our surprise that when multiple reflections are used the envelopment and the DFT depend strongly on the combination of delays chosen. This observation has pronounced implications for concert hall design. It appears that when there are multiple lateral reflections the delay times of the reflections relative to each other matter a lot. In fact, once a pattern has been set, the delay of this pattern relative to the direct sound can be varied with no change in envelopment. Some patterns are highly enveloping (greater than a diffuse field) and some are not enveloping at all. It seems that it is not just the total energy in the early reflections that is spatially important!

5b. Bandwidth of the test signal for music:

When we wish to calculate the DFT for music the test signal must have a longer correlation length than for noise. As mentioned above, we can use real reverberation as a test signal, or we can emulate the reverberation with a noise signal of 1-2Hz bandwidth. We will show some results based on narrow band noise.

5c. Other considerations in the DFT measurement system:

Another physiological variable in the DFT detector is the time constant used in the running ITD filter. Without this filter the ITD detector is not very accurate, so there is a considerable reason to include it. Here we simply guess. A time constant of about 50ms seems to work well. In practice the filter is implemented with a variable time constant. The TC is 50ms for strong signals, and rises linearly as the signal amplitude falls. This amplitude dependence keeps zero crossings at low amplitudes (which tend to be very noisy) from affecting the running average very much.

The 3Hz to 17Hz bandwidth used for the fluctuations is also a bit of a guess. It is based on measurements made with amplitude modulated and phase modulated pure tones. The bottom line is that this bandwidth gives quite reasonable results, so it seems a good choice for now.

The output of the DFT measurement is a number that represents the average of the absolute value of the interaural fluctuations. It is expressed in milliseconds. What is the meaning of this number? How large should it be when the envelopment is "just right", and is it possible for it to be too high?

5.2 Calibration of the DFT measurement system

We might think we could calibrate the DFT by using impulse responses measured in concert halls. This method has some advantages, but is probably not what we want to do. The DFT depends ultimately on the relationship between the strength of the lateral sound field in the region of the listener and the medial sound field. In a true diffuse field the medial signal will dominate the lateral signal, since the medial field includes both the

front/back direction and the up/down direction. In a concert hall this situation is altered by the floor reflection. At 63Hz the floor reflection enhances the lateral direction by canceling the vertical sound waves. At 128Hz the opposite happens, with the vertical sound being enhanced by the reflection. Which situation is optimal? In any case, why should we expect that a concert hall – even a very good one – would be optimal?

Once again experiments were performed to measure the envelopment using the apparatus of [44]. We found that below about 200Hz the most pleasing overall sensation of envelopment occurs when two uncorrelated noise sources are on opposite sides of a listener in an anechoic space. (Above this frequency the envelopment from such an array seems too wide, and a more diffuse soundfield is preferred.) This implies that below 200Hz the optimum value of the DFT is similar to the values we would measure in a concert hall at 63Hz, where the floor reflection enhances the lateral component of the sound. Our conclusion is that a diffuse field is not optimal at low frequencies. We can calibrate our DFT detector by measuring its value when it is exactly between two noise sources in a simulated anechoic space. The value we get is frequency dependent, but when the test signals has a bandwidth of one octave at 63Hz the anechoic DFT is about 0.24ms. In the work that follows we will use this method of calibrating the DFT, expressing our results in decibels relative to this anechoic value.

Externalization

6. ITDs and Externalization

The "in the head" perception also arises from the behavior of the ITD. When we are outdoors or in a large space, a sound source at the side produces ITDs of about 0.75ms at the listener's head. We easily perceive from this ITD that the source is external and at the side. When the sound source is in the medial plane (directly in front, overhead, or behind) the ITD is much lower. We perceive the sound as centered. We can usually tell if the source is directly in front or behind us. When this is possible, we perceive the sound as external.

It is well known that our ability to discriminate between rear sound and frontal sound at low frequencies depends on our ability to move the head. In natural hearing small movements of the head produce predictable changes in the ITD. Small head movements localize and externalize a source in the medial plane.

As we will see, in a small room the situation is quite different. For medial sources (or for a phantom image of a center source) a listener will experience a low ITD regardless of how the head is rotated. It is not possible to determine if the sound is from the front or the back.   The perception becomes internal and artificial – a perception peculiar to recorded music.

It is not necessary that the sound source be localized to a particular direction for externalization to occur. If the ITD changes randomly – not in synchrony with head movements – externalization is still successful. The resulting sound is perceived as external and enveloping. Our job in this paper is to find a measure that describes the degree to which low frequencies are externalized in a particular sound field – ideally from a modeled or measured binaural impulse response. We can then use the measure to find how externalization can be optimized.

7. Binaural impulse response and ITD

The Fourrier transform of the impulse response yields both pressure and phase as a function of frequency. To study the case where there is more than one driver active at the same time we find it convenient to first find the impulse response from each source separately. The phase of the transform from each source is adjusted if necessary, and then the transforms are added to find the total pressure and phase. The transform of a binaural impulse response is found by separately calculating the transforms for each ear. At low frequencies we assume the head can be approximated by two omnidirectional receivers, separated laterally by 25cm. We do not attempt to model diffraction around the head, which is minimal at the frequencies of interest.

The ITD as a function of frequency can be found by finding the phase difference between the transforms of the left and right binaural responses. The phase difference is then converted to a time difference by dividing by frequency.

It is useful before going further to examine the perceptual meaning of the impulse response and its Fourrier transform. Both are mathematical concepts. It is not obvious that either should have any relationship to what we hear with music. The impulse response is the response of the room to a pistol shot or small explosion. These are fortunately rare in everyday life. Nonetheless there is a considerable body of literature that attempts to relate musical perceptions to direct features of the impulse response.

The transform of the impulse response describes the steady state response of the room to single frequency sinusoids. These are plausibly musical. In a small rooms after the time it takes for the sound to travel across the room a few times the pressure is nearly at the steady state value. Most musical notes are longer than this, and often the attacks of low frequency instruments are slow enough that the room can be modeled by the steady state condition. However each frequency in the transform represents only a single musical note. We would like to know the response of the room over a range of notes. We need a measure for average pressure and average phase, where the average is over frequency.

8. Normalized Average Pressure (NAP), and Average ITD (AITD)

To find the average pressure we could simply sum the pressure squared over the range of interest to find the total pressure in the frequency band. The summation assumes that each frequency was present in the steady state – which is not particularly likely – but this measure of average pressure is useful nevertheless. To make the measure closer to what

we hear we weight the sum by -3dB per octave, so each 1/3 octave band gives an equal contribution to the sum.

We start with a binaural impulse response, either measured or calculated. That is, a separate impulse response for the left ear and the right ear of a dummy head.

IL(t) = impulse response for the left ear
IR(t) = impulse response for the right ear
Sum = sum of the frequency bins over the frequencies of interest

2. FFTL(f) = fft(IL(t))
3. FFTR(f) = fft(IR(t))

6. Average Pressure = sqrt(sum((magnitude FFTL(f)^2)/f)

This average pressure contains a frequency factor, which we would like to eliminate when we plot pressure as a function of frequency. We can define a normalized average pressure:

7. Normalized Average Pressure = Average Pressure / (sum (1/sqrt(f))

For low frequencies we assume that the magnitude of FFTL is approximately equal to the magnitude of FFTR.

We suggested above that the ITD at a particular frequency could be found by comparing the phase of the left and right parts of a binaural response. The problem is that this phase difference is independent of the pressure at that frequency. It is possible – in fact it is nearly always the case – that the frequencies where the phase difference is the largest are the ones for which the pressure is the lowest. We are not likely to take much notice of a large ITD if the pressure is low. Unlike a lateral figure of eight microphone, which gives a maximum output at the nulls of a lateral standing wave, we cannot determine the ITD of a sound we cannot hear. If we stand at such a null, we hear nothing.

It would be possible to develop a model that determines ITD the same way the ear does, by first separating the sound into critical bands, and then finding the timing of zero crossings for each ear. See [44.] As we will see later, such a model is quite expensive computationally. Since we want to calculate the AITD at many frequencies and at many points in a room, there is a strong incentive to develop an efficient model.

To make an average ITD that is closer to what we actually hear we weight the ITD with the pressure. ITDs with low pressure are given low weighting, and ITDs with high pressure are weighted more strongly. We then normalize the resulting sum by dividing by the average pressure over the same frequency range.

Given this description the method of finding the AITD follows directly:

5. Phase_angle(f) = angle(FFTL)-angle(FFTR)
6. ITD(f) = Phase_angle(f) /(2*pi*f)

To find the AITD we multiply the absolute value of ITD(f) by the absolute value of the pressure, and sum over frequency. We also apply a –3dB per octave weighting to give equal weight to all bands. The sum is then normalized by dividing by the weighted sum of the absolute pressure over the same frequency range.

6. AITD = sum(magnitude(FFTL(f))*abs(ITD(f))/(sqrt(f))/(average pressure)

It is also possible to define a NET ITD (NITD) – one which preserves the sign of the ITD in the sum. The NITD is always less than the AITD, but it indicates our ability to determine the direction of a sound source, not just its externalization.

7. NITD = sum(magnitude(FFTL(f))*ITD(f)/(sqrt(f))/(average pressure)

Note that for simplicity we use only the magnitude of FFTL(f) to represent the pressure at both ears. We could average the magnitudes of FFTL and FFTR, but there seems to be no reason to do so in practice.

9. AITD and NITD – what do they mean?

The NITD is a measure that expresses the apparent direction of a sound source in a given frequency band. Not every frequency within the band will localize to the same degree. In fact, due to standing waves some frequencies may localize to the wrong direction. The NITD averages out these differences, giving us a lower value than we would get using individual frequencies. It does reflect the average localization of a narrow band noise source. We are interested in externalization of the sound, and externalization can take place even in the absence of correct localization. Thus for externalization the average absolute value of the ITD, the AITD is a better measure.

At low frequencies in a free field – with no reflections – AITD and NITD are equal, and independent of frequency. Their value depends on the angle between the listener and the source, and can be easily calculated.

If source_angle is the angle between the source and a line drawn between the listener's ears (listener facing forward):

Lateral AITD = 0.75ms*cos(source_angle) See figure 1

We can also define a Medial AITD – by simply rotating the listener's head by 90 degrees in azimuth and calculating the AITD again. In a free field – assuming the source angle has not changed:

Medial AITD = 0.75ms*sin(source_angle) See figure 2

It makes sense to define a Total AITD, which is simply the RMS sum of the lateral and the medial AITD. Since cos^2 + sin^2 is one, in a free field

total AITD = 0.75ms. See Figure 3

The total AITD is a measure of how easily a sound source can be localized. In a free field it has the constant value of 0.75, which means the source can be localized to its true direction with full accuracy. In a reflective room the total AITD is almost always less than this value, as standing waves tend to reduce the ITD at the listener. Where the total AITD is significantly less than 0.75 the sound source will be "in the head".

At present it is not clear how close the AITD must be to 0.75ms for complete externalization. In our informal experience a value of 0.3ms or lower will almost always sound internal. Values of 0.4ms or so can sometimes be internal and sometimes be external. Clarifying this question will require further experiments. However the AITD is useful as a measure. We want to find strategies of speaker placement and recording technique that bring the AITD as close as possible to the free field value. Since in most cases the listener does not move the head very much, at low frequencies we are mostly interested in the lateral AITD.

10. 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!

11. 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.

12. 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.

Results

13. DFT in an anechoic space for noise signals

Earlier 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 with noise signals 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 no physical meaning at this time. The results we will show here will plot the DFT in dB relative to this value.

As mentioned earlier 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 4 shows the DFT for conventional stereo in an anechoic

space. The listening plane is at 4' from the (imaginary) floor, and the speakers are below, at 1.5'. Figure 5 shows the value of DFT along the center line of the room. Note that with the speakers in the front the envelopment is reduced by 5dB compared to having the speakers at the sides. Figure 5 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.

14. 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 6 shows the DFT for a 12'x15'x9' room, with a surface reflectivity of 0.8. Figure 6a 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 6b 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.

Figure 7 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.

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.

14. DFT in a reflective room for music signals

In the previous 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 8 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.

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.

15. 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 9, 10, and 11 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'. 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 12 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 13 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 14, 15, and 16 show similar data for a smaller room, 12'x15'x9'. The unusual shape of the Lateral AITD surface in figure 15 is due to a strong standing wave at about

70Hz. Most of the other frequencies are well represented by figure 16. Once again we see that pressure is generally lowest just where we would like to listen, and so is the lateral AITD.

The medial AITD for the 12'x15' rooms is plotted in figure 17. Unlike the room of figure 10, 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 a good place to start. 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.

16. 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 18 – 22 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.

17. 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 23 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 that 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.

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 24 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 25 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.

18. 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?

18a: 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.

24

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 26 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 27 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.

18b 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.

Figure 28 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 29-32 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.

19. Pressure and AITD from four drivers.

There is an additional improvement when four drivers are used. Figures 33-37 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.

20. Listening tests

This paper introduces two new methods of evaluating the sound of a room, the AITD 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 AITD, 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. Above this frequency for most of the rooms we have modeled the AITD becomes large enough to provide externalization regardless of the placement of the loudspeakers. 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 DFT has not been specifically tested as a measure. However, the major predictions that we have been able to make are well supported by the author's experience. First, the DFT predicts that multiple low frequency drivers driven by independent reverberant signals will be preferable to a single low frequency driver. Second, it predicts that low frequency drivers located to the sides of the listening area will produce more envelopment than drivers at the front. This has also been observed in our listening rooms.

20. Conclusions

AITD modeling and DFT modeling both show 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. Both externalization and DFT depend to some degree on the reflectivity of the room. Where the room reflectivity is high the room can become more enveloping, particularly with broadband signals. In general externalization can be improved by shifting the phase of the low frequencies by 90 degrees for the left and right sides of the listener.

Conventional surround systems that utilize a combination of two full range loudspeakers and a number of smaller satellite loudspeakers would most likely produce a more enjoyable sound if the full range speakers were located at the sides of the room, and the smaller speakers in front. Although such a loudspeaker arrangement is unusual, it has sonic advantages. In particular cases it may also be more convenient to place the larger speakers alongside the listeners rather than in the front of the room.

References

1.   Ando, Y. and Singh, P.K. and Kurihara, Y., Subjective diffuseness of sound field as a function of the horizontal reflection angle to listeners. Preprint received by the author from Dr. Ando
2.   Ando, Y. and Kurihara, K., Nonlinear response in evaluating the subjective diffuseness of sound fields. J. Acoust. Soc. Am. 80 [1986], 3 pp 833-836
3.   Barron, M., Spatial Impression due to Early Lateral Reflections in Concert Halls: The Derivation of a Physical Measure. J. Sound and Vibration 77(2) [1981] pp 211, 232
4.   Beranek, L., Music, Acoustics and Architecture. John Wiley, 1962
5.   Beranek, L., Concert Hall Acoustics - 1992. J. Acoust. Soc. Am. 92, [1992]
6.   Beranek, L., Concert and Opera Halls How They Sound.   Acoustical Society of America, 1996
7.   Blauert, J., Zur Tragheit des Richtungshorens bei Laufzeit- und Intensitatsstereophonie. Acustica 23 [1970] p287-293
8.   Blauert, J., On the Lag of Lateralization Caused by Interaural Time and Intensity Differences. Audiology 11 [1972] pp265-270
9.   Blauert, J., Raumliches Horen. S. Hirzel Verlag, Stuttgart, 1974
10.   Blauert, J., Spatial Hearing. MIT Press, Cambridge MA 1983
11.   Bradley, J.S., Contemporary Approaches to Evaluation of Auditorium Acoustics. Proc. 8th AES Conference Wash. DC May 1990 pp 59-69
12.   Bradley, J.S., Contemporary approaches to evaluating Auditorium Acoustics. Proceedings of the Sabine Conference, MIT June 1994.
13.   Bradley, J.S. and Soulodre G.A., Spaciousness judgments of binaurally reproduced sound fields. Ibid. p 1-13
14.   Bradley, J.S., Comparisons of IACC and LF Measurements in Halls. 125th meeting of the Acoustical Society of America, Ottawa, Canada, May 1993

15. Bradley, J.S., Pilot Study of Simulated Spaciousness. Meeting of the Acoustical Society of America, May 1993

16. Bradley, J.S. and Souloudre, G.A., Objective measures of Listener Envelopment. J. Acoust Soc. Am. 98 [1995] pp 2590-2597

17. Bradley, J.S. and Souloudre, G.A., Listener envelopment: An essential part of good concert hall acoustics. JASA 99 [1996] p 22

18. Gardner, W. and Griesinger, D., Reverberation Level Matching Experiments' Proceedings of the Sabine Conference, MIT June 1994. p263

19. Gold, M.A., Subjective evaluation of spatial impression: the importance of lateralization. ibid. p97

20. Griesinger, D. Measures of Spatial Impression and Reverberance based on the Physiology of Human Hearing. Proceedings of the 11th International AES Conference May 1992 p114 - 145

21. Griesinger, D., IALF - Binaural Measures of Spatial Impression and Running Reverberance. Presented at the 92nd Convention of the AES March 1992 Preprint #3292

22. Griesinger, D., Room Impression, Reverberance, and Warmth in Rooms and Halls. Presented at the 93rd Audio Eng. Soc. convention in San Francisco, Nov. 1992. AES preprint #3383

23. Griesinger, D., Progress in electronically variable acoustics. Proceedings of the Sabine Conference, MIT June 1994.

24. Griesinger, D., Subjective loudness of running reverberation in halls and stages. Proceedings of the Sabine Conference, MIT June 1994. p89

25. Griesinger, D., Quantifying Musical Acoustics through Audibility. Knudsen Memorial Lecture, Denver ASA meeting, 1993.

26. Griesinger, D., Further investigation into the loudness of running reverberation. Proceedings of the Institute of Acoustics (UK) conference, Feb 10-12 1995.

27. Griesinger, D., How loud is my reverberation. Audio Engineering Conference, Paris, March 1995. Preprint # 3943

28. Griesinger, D., Design and performance of multichannel time variant reverberation enhancement systems. Proceedings of the Active 95 Conference, Newport Beach CA, June 1995

29. Griesinger, D., 'Optimum reverberant level in halls' proceedings of the International Congress on Acoustics, Trondheim, Norway June 1995

30. Hartman, W., Localization of a sound source in a room. Proceedings of the 8th international conference of the Audio Engineering Society, May 3-6 1990

31. Hidaka, T. and Beranek, L. and Okano, T., Interaural cross correlation (IACC), lateral fraction (LF), and sound energy level (G) as partial measures of acoustical quality in concert halls. J. Acoust Soc. Am. 98 [1995] pp 988-1007

32. Jullien, J.P. and Kahle, E. and Winsberg, S. and Warusfel, O., Some Results on the Objective Characterization of Room Acoustical Quality in Both Laboratory and Real Environments Proc. Inst. of Acoustics Vol. 14 part 2 (1992). presented at the Institute of Acoustics conference in Birmingham, England, May 1992

33. Kahle, E., Validation d'un modele objectif de la perception de la qualite acoustique dans un ensemble de salles de concerts et d'operas' Doctorate Thesis, IRCAM June 1995

34.   deV. Keet, W., The Influence of Early Lateral Reflections on the Spatial Impression. The 6th International Congress on Acoustics, Tokyo, Japan, Aug 21-28 1968. pp E53 to E56

35.   Kreiger, A., Nachhallzeitverlangerung in der Deutschen Staatsoper Berlin. Tonmeister Informationen (TMI) Heft 3/4 Marz April 1996.

36.   Morimoto, M. and Maekawa, Z., Effects of Low Frequency Components on Auditory Spaciousness Acustica 66 [1988] pp 190-196

37.   Morimoto, M. and Posselt, C., Contribution of Reverberation to Auditory Spaciousness in Concert Halls. J. Acoust. Soc. Jpn (E)10 [1989] 2

38.   Morimoto, M. and Maekawa, Z., Auditory Spaciousness and Envelopment. 13th ICA, Yugoslavia 1989

39.   Morimoto, M., The relation between spatial and cross-correlation measures 15th ICA Norway 1995

40.   Schroeder, M.R. and Gottlieb, D. and Siebrasse, K.F., Comparative study of European concert halls: Correlation of subjective preference with geometric and acoustic parameters. J. Acoust. Soc. Am. 56 [1974] pp 1195ff

41.   Schubert, P., Die Wahrnehmbarkeit von Einzelruckwurfen bei Musik. PhD thesis, TU Dresden 1967

42.   Schultz, T.J., Acoustics of the Concert Hall" IEEE Spectrum (June 1965) pp 60-62

43.   Griesinger, D. "Spaciousness and Localization in Listening Rooms and their Effects on the Recording Technique" JAES v34 #4 p255-268 April, 1986

44.   Griesinger, D. "The Psychoacoustics of Apparent Source Width, Spaciousness and Envelopment in Performance Spaces" Acta Acustica Vol. 83 (1997) 721-731
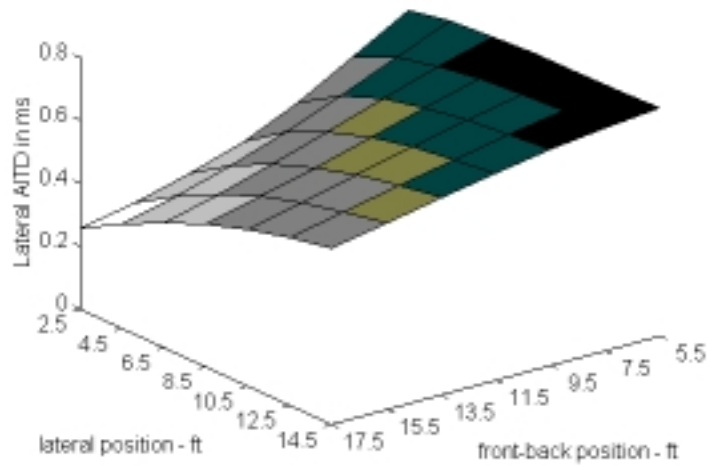
Figure 1: The Lateral AITD in the center of a 17'x23' anechoic space with a driver in the upper left corner (at .1',.1'). Where the value is high, the sound source will be localized to the side of the listener.
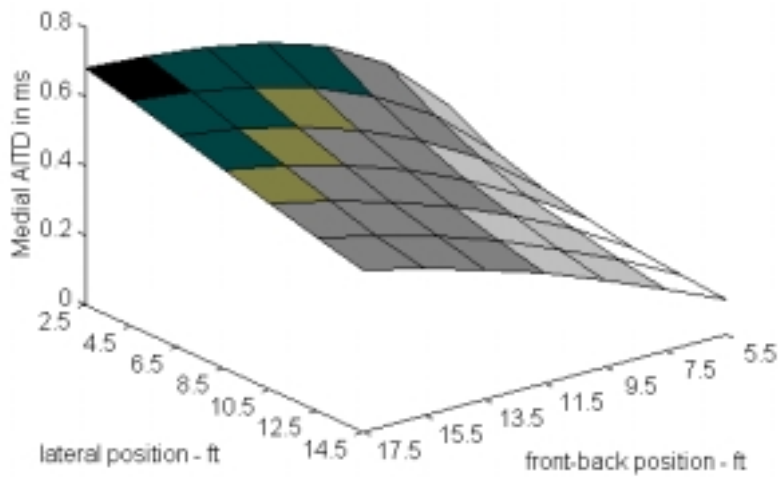
Figure 2: The Medial AITD in the center of a 17'x23' anechoic space with a driver in the upper left corner at .1', .1' Note that where the lateral AITD is low, the medial AITD is high.
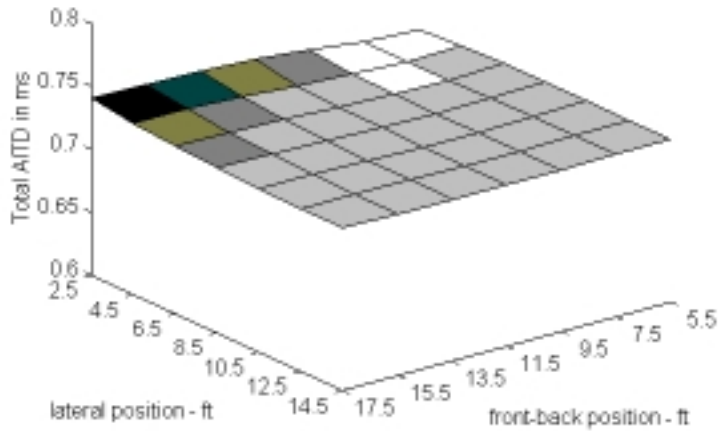
Figure 3: The Total AITD (the RMS sum of the lateral and medial AITDs) in the same space as figures 2 and 3. Note that the driver can be accurately localized everywhere, as we would expect in an anechoic environment. The differences from 0.75ms are due to sampling errors in the model.

A: B:
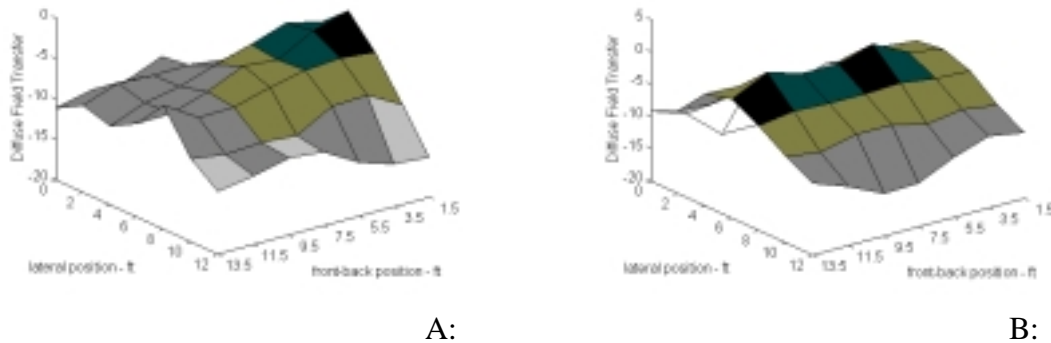
Figure 4: Diffuse Field Transfer function (DFT) for a 12'x15' anechoic space. 63Hz octave band. A: two uncorrelated speakers in the front, +-4' from the center. The value of 0dB is optimal. Note in the listening area the DFT is reduced by the +- 30 degree angle between the loudspeakers and the listener. 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.
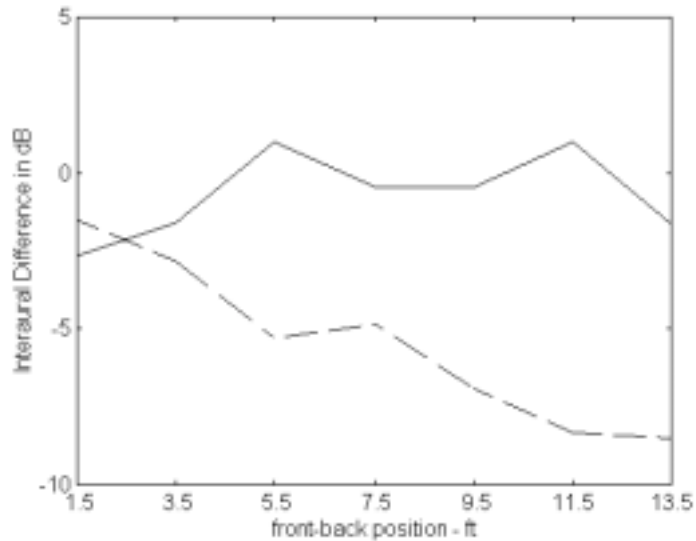
Figure 5: DFT along the center line for the two configurations in figure 4. ___ = two drivers at the sides of the listener __ __ __ = two drivers at the front. Note the approximately 5dB difference in envelopment in the listening area.
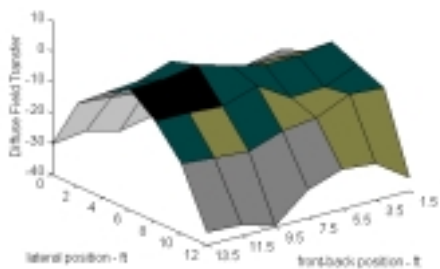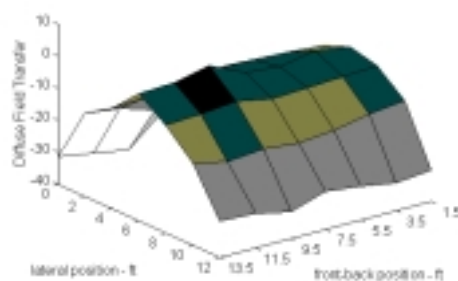
3a   3b

Figure 6: DFT for a 12'x15x9' space with surface reflectivity of 0.8. 63Hz octave band. 6a has a single driver in the front of the room on the left side. Figure 6b 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.
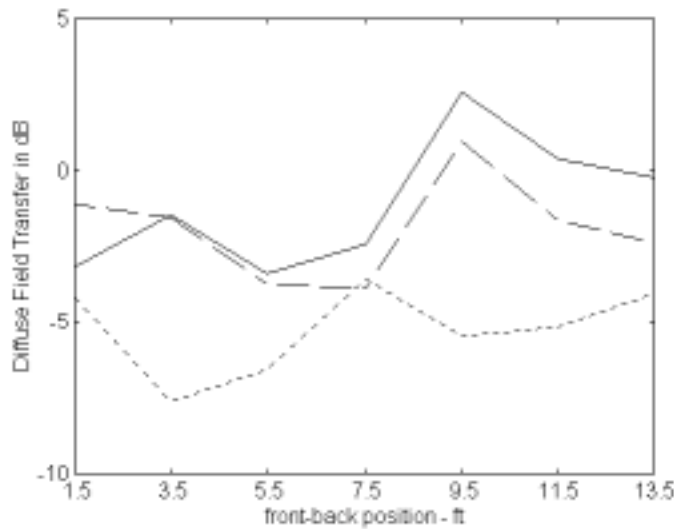
Figure 7: The DFT along the center line of a 12'x15'x9' room. 63Hz octave band. Speakers in the front of the room at +-4'. _____ = Two speakers, room reflectivity of 0.8 on all surfaces. __ __ __ = same, but with a lateral reflectivity of 0.6. - - - - = a single loudspeaker with a lateral reflectivity of 0.6. Note the reduced envelopment when there is a single sound source and a low lateral reflectivity.
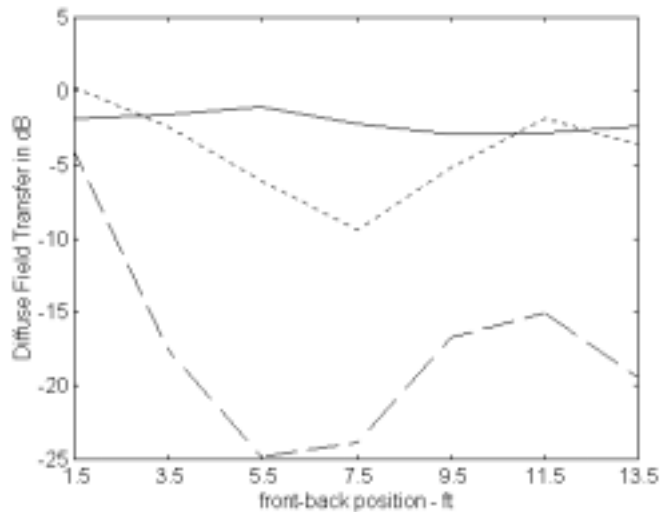
Figure 8: The DFT along the center line of a 12'x15'x9' room with surface reflectivity of 0.8. _____ = Two uncorrelated loudspeakers at the sides of the listeners, at 11' from the front. - - - - = Two uncorrelated loudspeakers at the front of the room, 4' apart. __ __ __ = A single driver in the front 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 7, where the bandwidth of the test signal is 45Hz.
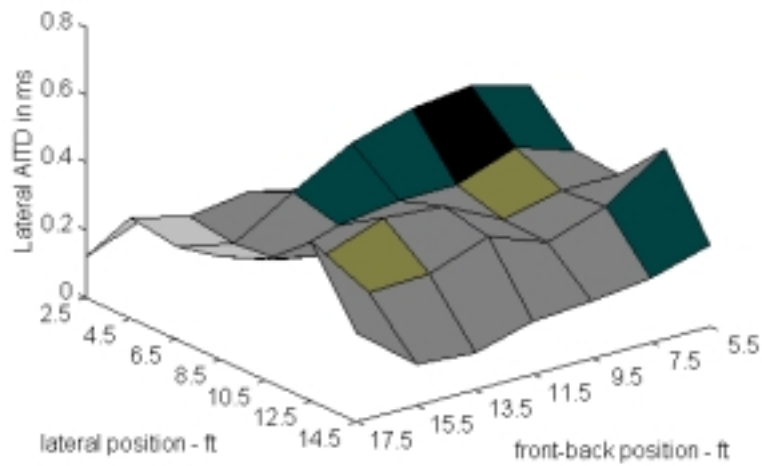
Figure 9: Lateral AITD in the 63Hz octave band for a 17'x23'x9' room, surface reflectivity 0.8. Single driver in the upper left corner. Note the value is lower than in the free field. (Typically 0.7ms.) See figures 1,2, and 3.
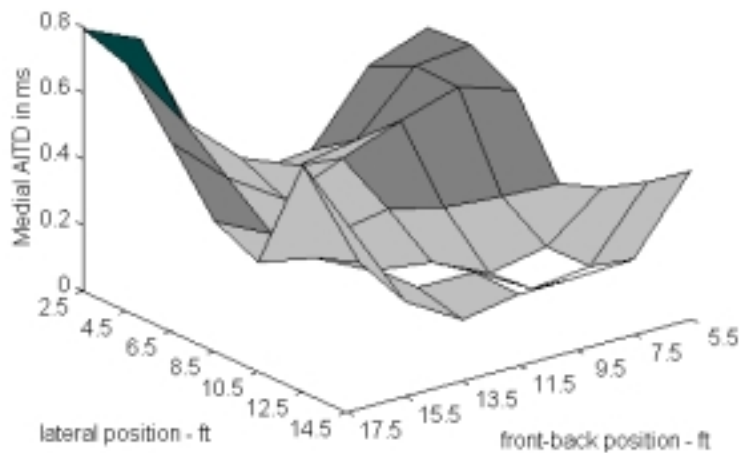
Figure 10: Medial AITD for the same room. Note the value is larger than figure 9, but still lower than free field. This drawing represents the lateral AITD if the primary listening axis is parallel to the short wall, rather than parallel to the long wall.

34
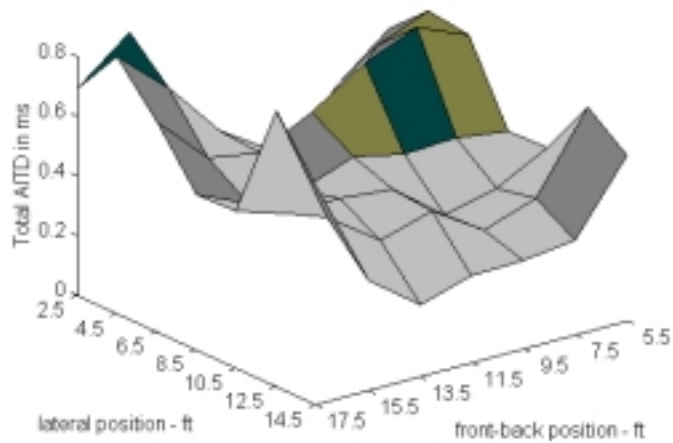
Figure 11: The Total AITD for the same room. In the listening area the AITD is roughly half the value of the free field. Sounds are somewhat externalized in such a field.
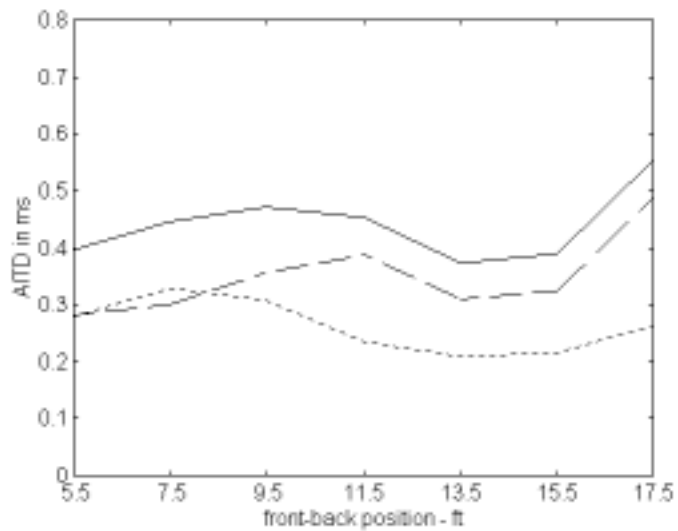
Figure 12: ___ = Total AITD along the center line of the same room. - - - = Lateral AITD, __ __ __ = Medial AITD. Receiver height at 4'. Some externalization is possible in this sound field, but it is not as easily externalized as a free field.
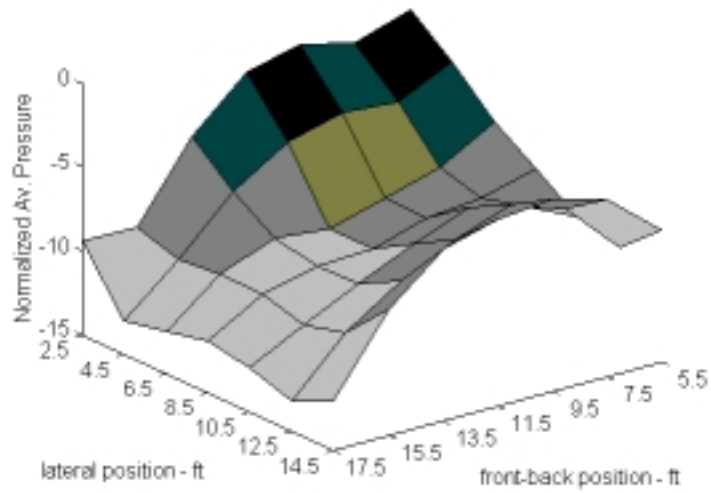
Figure 13: Normalized Average Pressure in the 63Hz octave band from a single driver of unit strength in the upper left corner of a 17'x23'x9' room. The wall reflectivity is 0.8. Note the pressure in the listening area is low. Equalization can raise the pressure, but it cannot change the AITD.
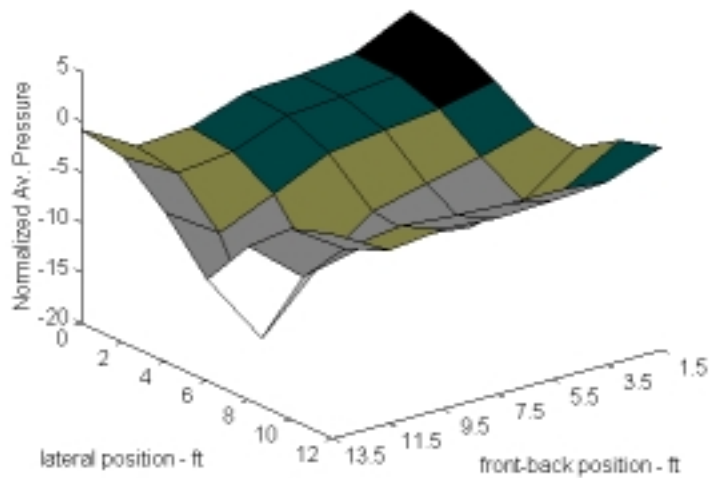
Figure 14: The Normalized Average Pressure in a 12'x15'x9' room in the 63Hz octave band from a single driver in the top left corner.
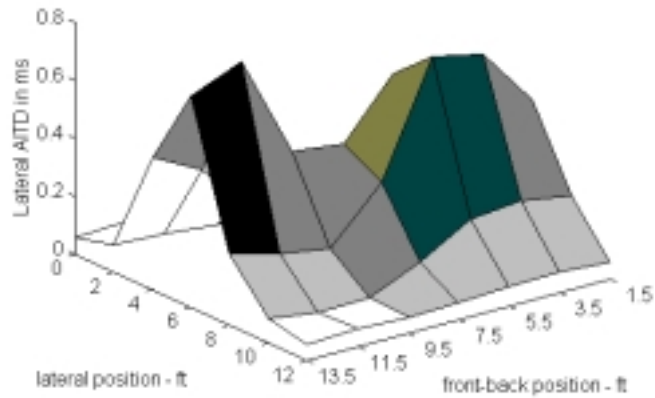
Figure 15: Lateral AITD in a 12'x15'x9' room, surface reflectivity 0.8, single driver in the top left corner, 63Hz octave band. This band has the largest AITD for this speaker position in this room. The next figure is more typical.
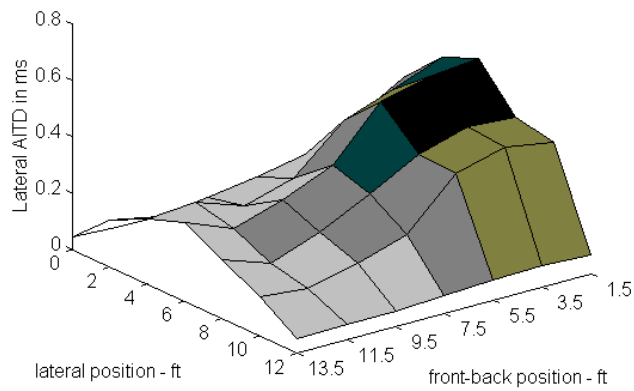
Figure 16: The Lateral AITD over a range of 20Hz to 90Hz in a 12'x15'x9' room from a single driver in the front left corner. Surface reflectivity is 0.8.
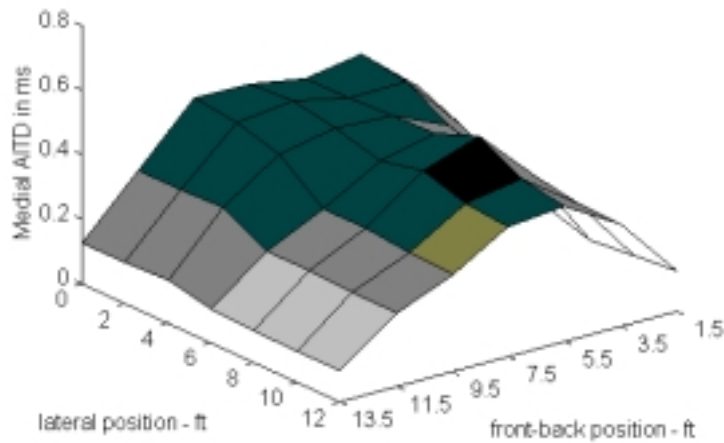
Figure 17: Medial AITD for the 12'x15'x9' room over the 20Hz to 90Hz range. Single driver is in the upper left corner. This picture represents the lateral AITD if the room is set up with the primary listening axis parallel to the short walls, rather than parallel to the long walls. Notice that in this 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.
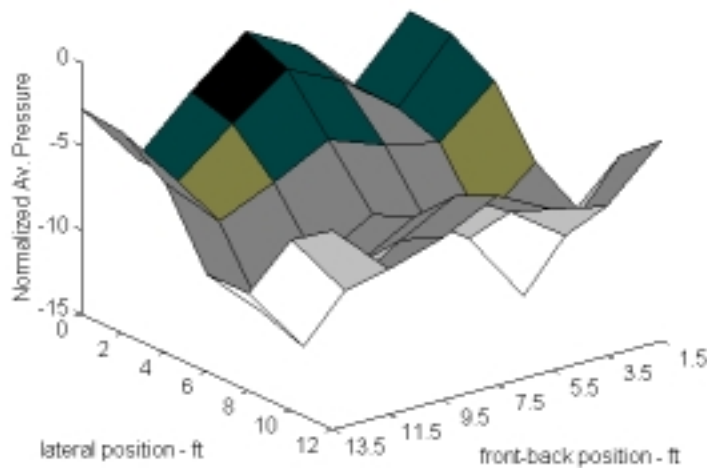
Figure 18: Normalized Average pressure in the same room, but the driver is now at the side of the room, at y=7.5', x=0.1', z=1/5' Although the driver is not in the corner of the room, the pressure in the listening area is not significantly reduced. See figure 14.
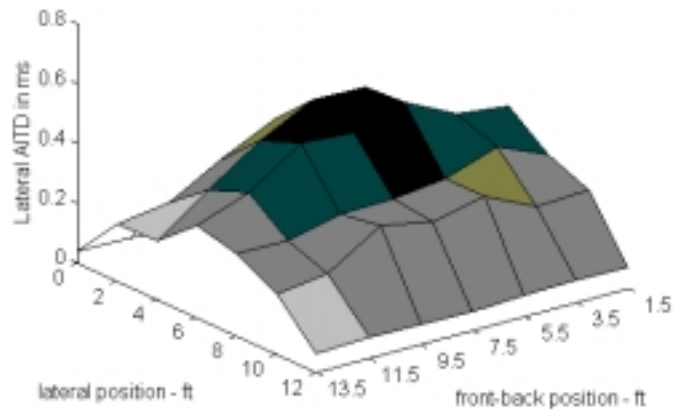
Figure 19: The Lateral AITD over a range of 20Hz to 180Hz in a 12'x15'x9' room with the driver at the side. Note the significantly higher values than for figure 16.
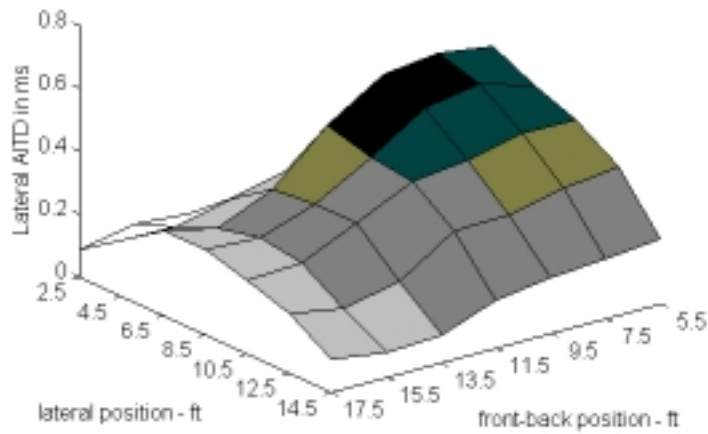
Figure 20: The Lateral AITD over a range of 20Hz to 90Hz in a 17'x23'x9' room with the driver in the upper left corner. See the pressure response in figure 13.
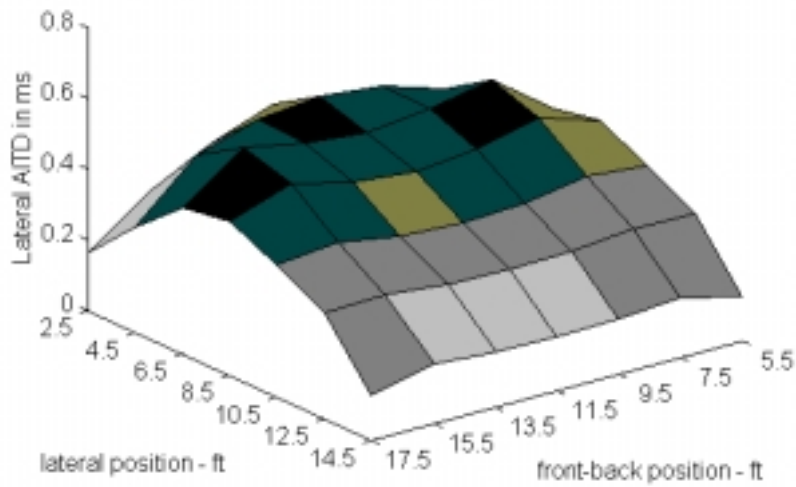
Figure 21: The Lateral AITD over a range of 20Hz to 90Hz in a 17'x23'x9' room with the driver at the side of the room in position x=0.1', y=11.5', z=1.5' Note the substantially higher values than for figure 20.
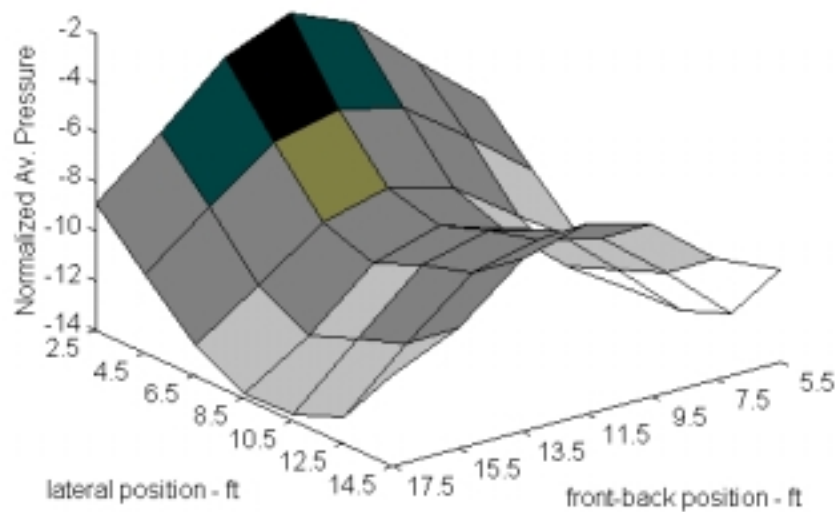
Figure 22: The Normalized Average pressure over the same room as figure 20. Note the pressure is not significantly lower than in figure 13, even though the driver is not in the corner of the room.
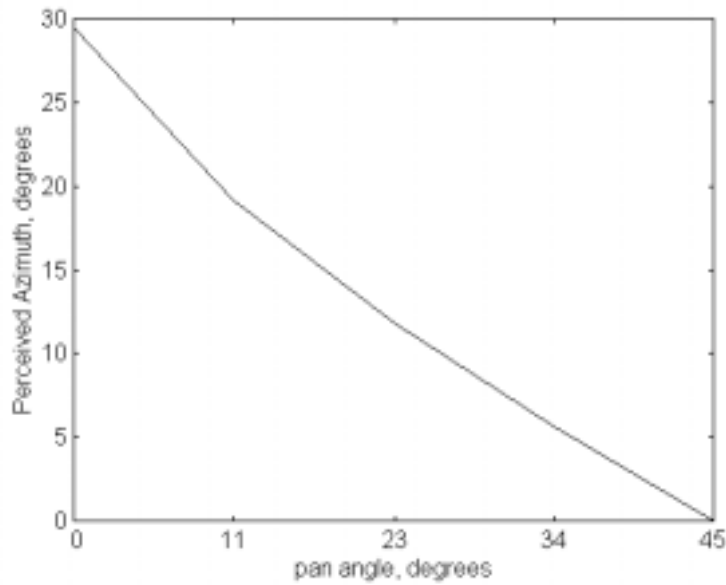
Figure 23: Perceived angle, as calculated from the ITD, for a listener at the ideal listening position, with loudspeakers +-30 degrees, as a sound is panned from left to center. Anechoic environment
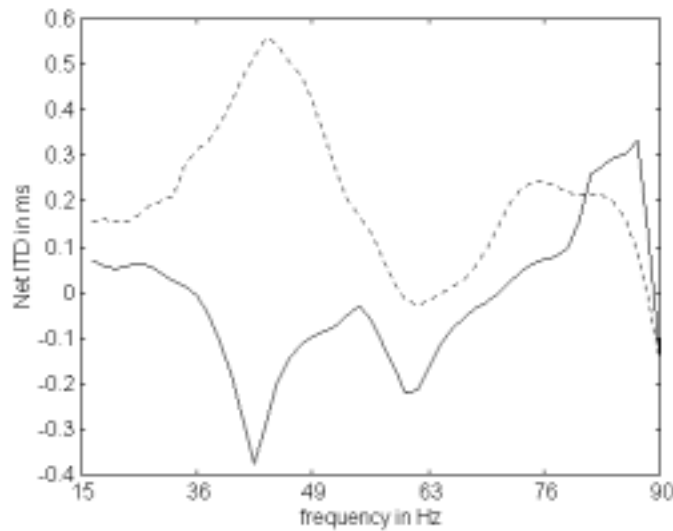
Figure 24: The frequency dependence of the ITD as a function of frequency for a single driver at x=.5', y=.1', z=1.5' for a listener at x=6', y=9.5', z=4' in a 12'x15'x9' room with reflectivitiy 0.8. ___ = lateral ITD. - - - = medial ITD. Note that for frequencies below 70Hz the lateral ITD is negative, and the speaker appears to be in the opposite corner.
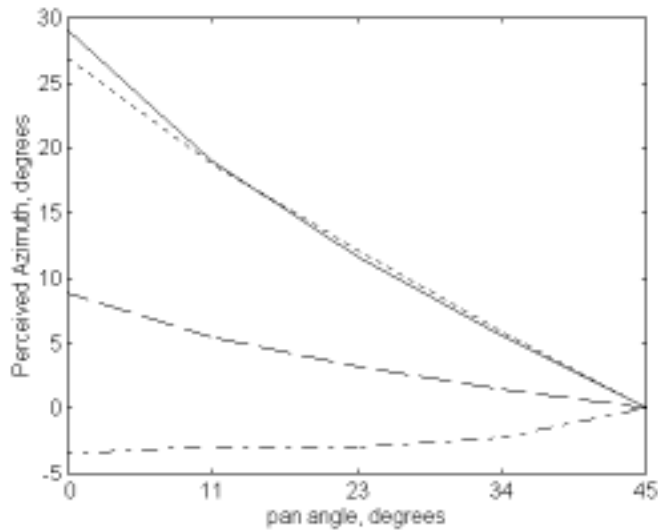
Figure 25: Pan law for four different room reflectivities. ___ = anechoic, - - - = 0.5, __ __ __ = 0.65, - . - . - . = 0.8. Note the ability to localize horizontally goes down rapidly as the reflectivity goes up. Same room and positions as figure 24.
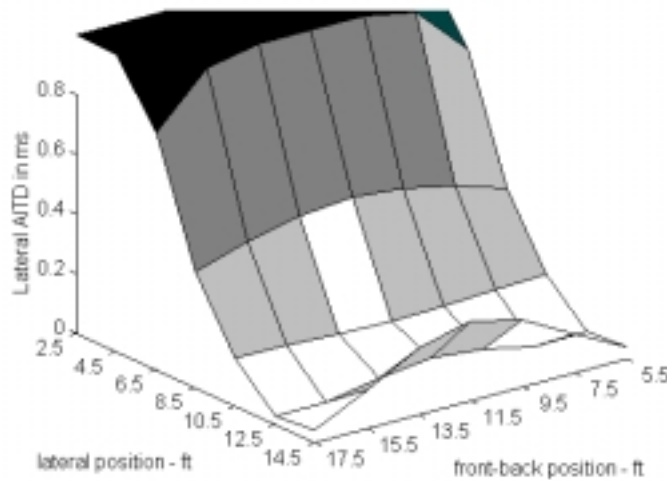
Figure 26: Lateral AITD in an anechoic space, two drivers on opposite sides of the listener, separated by +-11'. Left driver leads the right driver in phase by 90 degrees. 22.5Hz. Note the peak in AITD at -4' from the center, and the minimum at +4'. These values correspond to the minimum and maximum in the pressure response. The value of AITD along the center axis (8.5') depends on the spacing of the minimum and maximum.
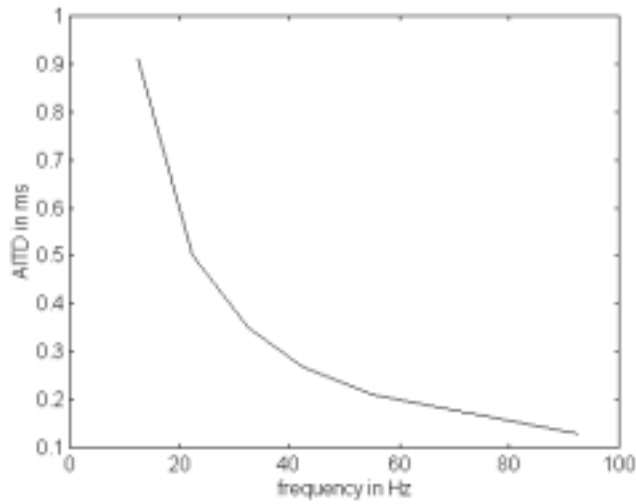
Figure 27: The dependence of the AITD in the center of figure 26 with frequency. Note the typical values of 0.2ms or so above 60Hz. Adding surface reflections can increase this value.

a.

b.
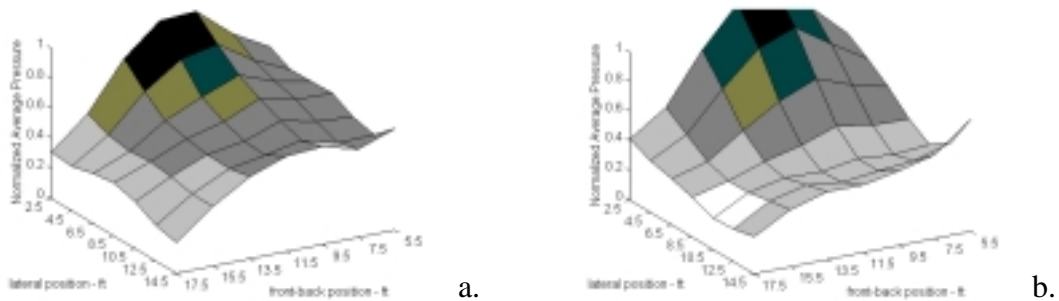
Figure 28: Normalized Average Pressure in a 17'x23'x9' room in the 63Hz octave band, from two drivers along the front wall. Graph a. shows a 30degree phase shift, graph b. shows a 90 degree phase shift. Notice that a dip in the pressure occurs to the right of the central axis.

a.

b.

c.
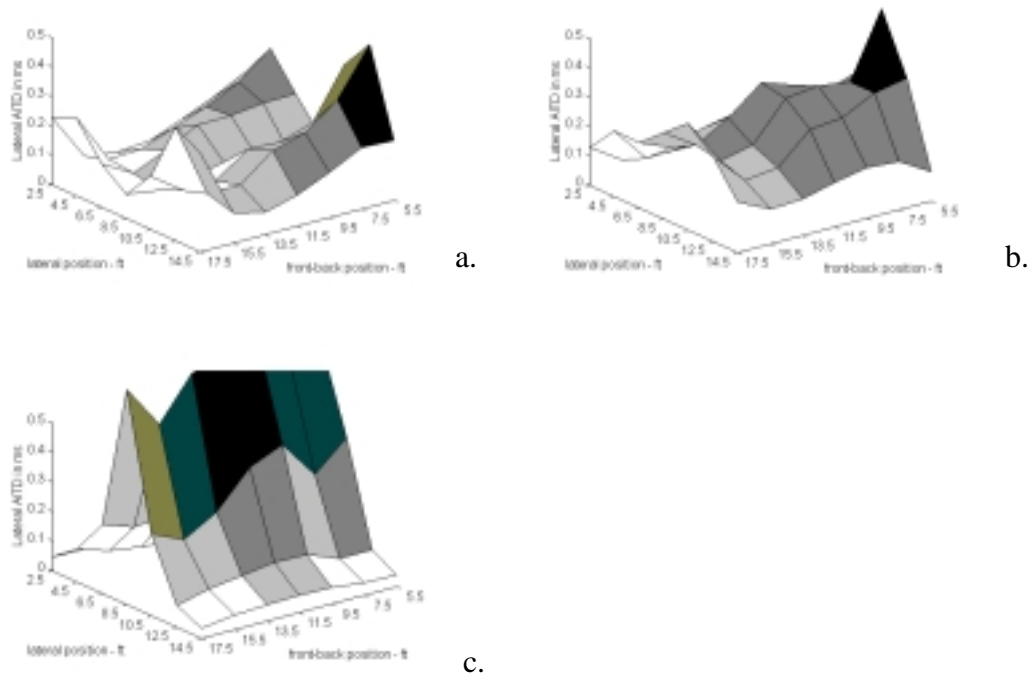
Figure 29: Absolute AITD in the 63Hz octave band for the same room as figure 28. Graph a. is 30 degree phase shift, graph b. has a 90 degree phase shift, and graph c. has a 150 degree phase shift. Note the very low AITD in graph a., a moderate AITD in graph b., and an unnatural (phasey) AITD in graph c.
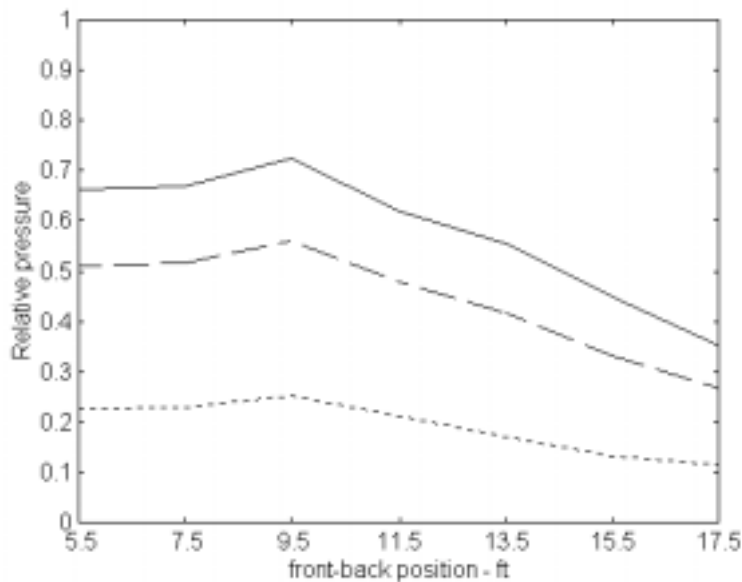
Figure 30: The Normalized Average Pressure along the center axis of the room from figure 28. _____ = 30 degree shift, __ __ __ = 90 degree shift, - - - - =150 degree shift. Note the ~3dB reduction in pressure in the 90 degree curve.
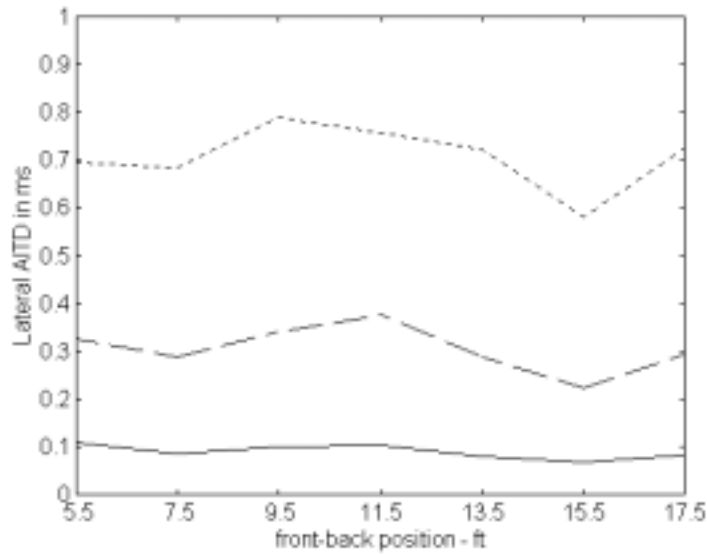
Figure 31: The lateral AITD in the 63Hz octave band along the center axis of the same room. ___ = 30 degrees, __ __ __ = 90 degrees, - - - =150 degrees. Note the very low AITD for the 30 degree case, and the moderate AITD for the 90 degree case. The 150 degree case is unnaturally large and phasey.
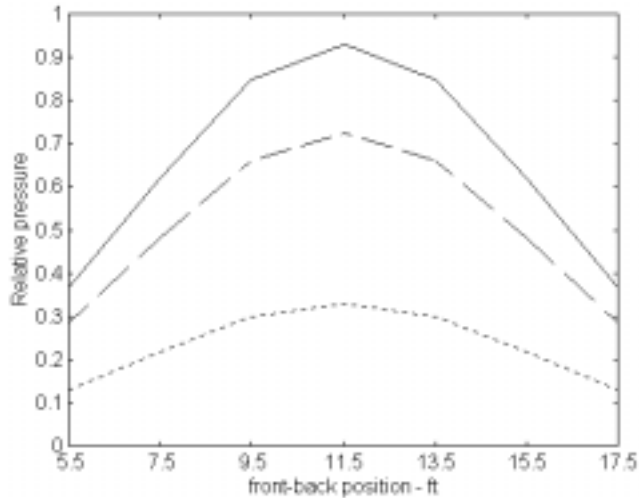
Figure 32: The Normalized Average Pressure in the 63Hz octave band in the same room as figure 28, but with the drivers at the sides of the room at y=11.5'. Receiver is along the center axis. ___ = 30 degree phase shift, __ __ __ = 90 degree phase shift, - - - =150 degree phase shift. Note the pressure in the listening area is higher than with the drivers in the front of the room.
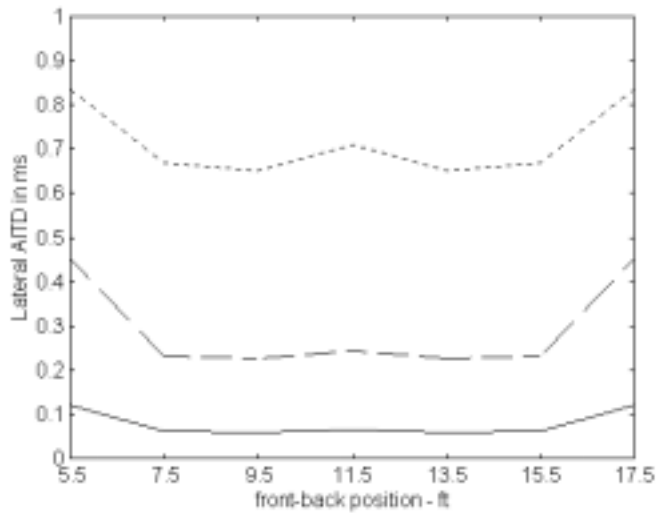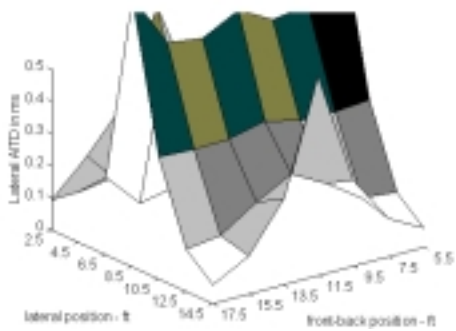
Figure 33: The Lateral AITD in the 63Hz octave band for the same configuration as figure 32. The AITD for the 90 degree shift is again moderate, the 30 degree case is much too low, and the 150 degree case is too high. Note that for this frequency band the AITD for 90 degree phase is approaching the anechoic value. The value below 50Hz and above 70Hz is ~0.4ms.
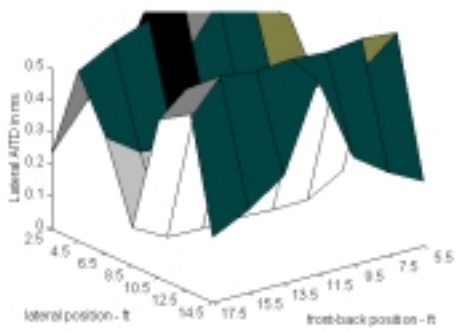
a.
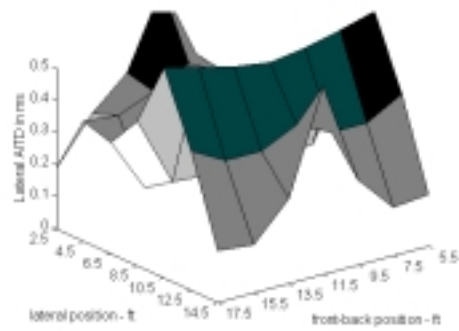
b.

c.

Figure 34: 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.
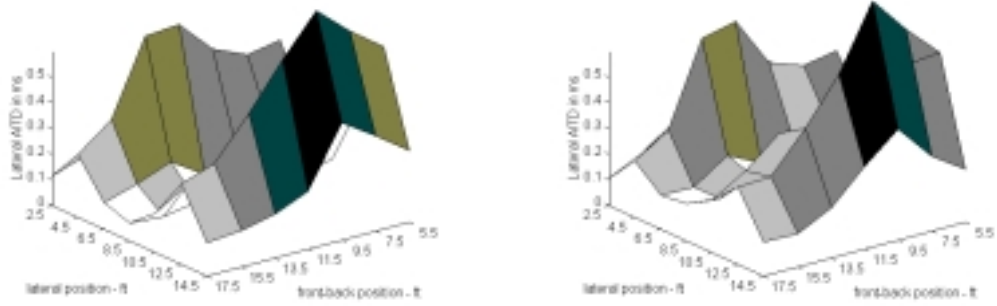
Figure 35: Lateral AITD's in the listening area from 4 drivers, two in the front at +-5.5', and two at the sides at +- 8'. All reflectivities are 0.8. 63Hz octave band. Picture a. is for 30 degree shift, picture b. is for 60 degree shift.
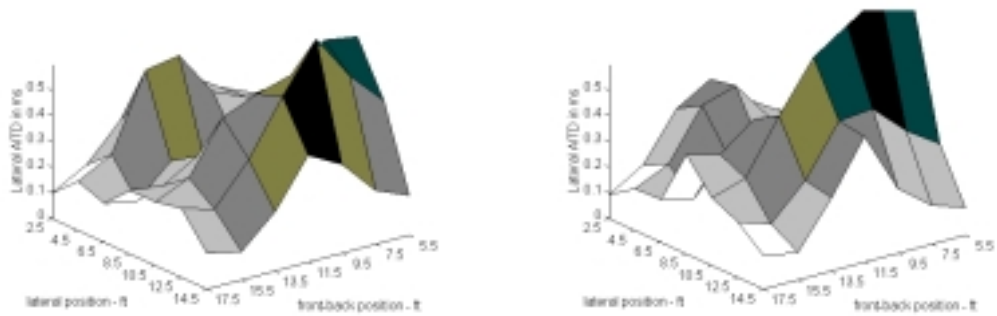

Figure 36: Same configuration as figure 35, but a. is 90 degrees, and b. is 120 degrees.
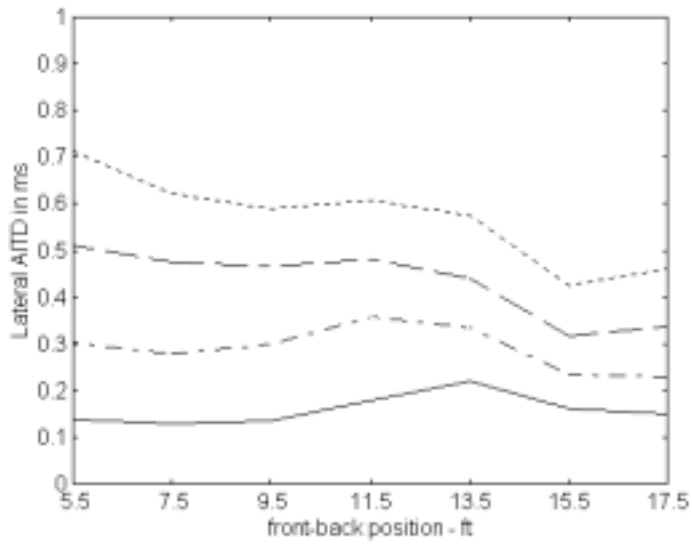
Figure 37: AITD in the range of 30-90Hz along the center line of a 17'x23'x9' room
with four drivers, the right side of the room lagging the left side by a variable phase shift.
____ = 30 degrees, __ - __ = 60 degrees, __ __ __ = 90 degrees, and - - - = 120 degrees.
The increase in externalization at 90 degrees is highly audible.