

Phase Coherence as a Measure of Acoustic Quality, part two: Perceiving Engagement

David Griesinger

Consultant, 221 Mt Auburn St #107, Cambridge, MA 02138, USA

PACS: 43.55.Fw, 43.55.Mc, 43.66.Ba, 43.66.Hg, 43.66.Jh, 43.66.Qp

ABSTRACT

The first of these three papers described the physics and physiology that enables humans to detect nearly instantly the apparent closeness of a sound source. In this section we describe some of the research that forms the basis for many of the myths that govern the use of early reflections in current hall designs, and the personal experiences that lead the author to question these myths. With the help of many fine artists, musicians and engineers I learned to hear how the apparent closeness of a sound directs the brain to pay close attention – and thus the degree of engagement we have in a performance of drama or music. Together these experiences constitute a plea for acoustic designs that encourage engagement over a wide range of seats. Since engagement and the clarity of both speech and music are closely related, this plea is really asking that we design halls where all the notes can be heard. Composers will thank us.

INTRODUCTION

Current thinking about early reflections is dominated by the work of Haas on the audibility of single reflections, and the work of Barron on the spatial impressions that arise from such single reflections as a function of the horizontal angle of incidence (azimuth).

Roughly speaking, Haas discovered that if a reflection (or the sound from a reinforcing loudspeaker) arrives at a listener within 50ms of the direct sound the reflection reinforces the loudness of the direct sound without reducing the intelligibility. Some of the work of Haas has been misrepresented in the common belief that the “first wavefront” of a sound is always audible, and will dominate the sonic perception. Haas showed this to be true when there is only one reflection, and when that reflection is appropriately delayed. But in concert halls the first wavefront is often masked by multiple reflections. Barron found that when single reflections arrive from the side they tend to induce the perception of space, or “spatial impression”.

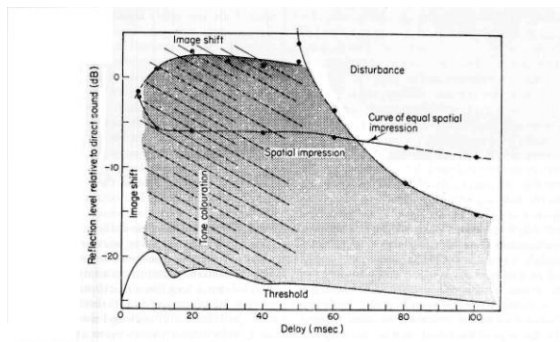


Figure 1: The effect of spatial impression from a single lateral reflection as a function of the time delay and level of the reflection. From Barron [1]. Note that the strength of the

reflection varies from -25dB to +3dB relative to the direct sound. This is a direct to reverberant ratio (D/R) of +25dB to -3dB. These D/R ratios are typical of recordings – and not of concert halls.

Barron’s work has greatly influenced hall designs. It led rapidly through the work of Marshall to the assumption that lateral reflections were to be maximized in good acoustic designs. I did not initially question this assumption, and made a series of lectures and papers extending Barron’s work. [2][3] [4] But all this research, and to some degree the work of Haas, has basic flaws when it is applied to concert halls.

The importance of multiple reflections

Part of the problem is that (as both Haas and Barron noticed) single early reflections are often not audible as such. They blend in with the direct sound, altering some of its properties. If they come from the medial plane (in front, overhead, or directly behind) they may not be audible at all.

The other – and more important problem – is that Barron and others, including myself, studied what happens when low-level reflections are gradually added to a strong direct signal. The D/R in these experiments is almost always greater than one. In actual spaces there are multiple reflections. The ear is sensitive to the *SUM* of the reflected energy, not to the energy of each reflection individually, and the energy in the sum of the reflections is nearly always greater than the energy in the direct sound.

Toole and Olive [5] studied the audibility of individual reflections in rooms, and in [6] Toole implies that if individual reflections are below the level of audibility with respect to the direct sound they can be ignored. This is clearly not the case. In small rooms individual reflections are almost always inaudible individually, but there are a great many of them, and their sum is highly audible.

The energy in the reflections in halls is greater than the energy in the direct sound for an obvious reason: the strength of the direct sound decreases 6dB with every doubling of distance from the source. In a typical concert hall half the seats have a D/R less than -10dB. In almost all seats the D/R is less than -3dB. A major purpose of a hall – besides keeping out the rain – is to capture sound that does not travel directly to a listener, and re-direct it to the listeners. In the best halls the re-directed sound increases loudness without compromising engagement.

If we duplicate Barron's work with multiple reflections, and with D/R ratios less than -3dB, we find that far from being benign, when the sum of reflections that arrive in the first 100ms after the direct sound exceeds a critical value, the ear is no longer able to detect the direct sound as separate from the reverberation. The clarity, timbre, and localizability of the sound changes in profoundly. Surprisingly, the apparent loudness of the reverberation also decreases, as does the perception of envelopment.

FORMATIVE EXPERIENCES WITH OPERA

The author has had the good fortune to work with sound and reverberation in many capacities. My earliest experiences came as a sound engineer of classical music, work I still enjoy. I recognized the critical lack of natural-sounding artificial reverberation, and started a career designing these devices. This led to extensive work in recording, sound playback, and electronic architecture - in spaces as large as Chicago's Grant Park, and as small as automobiles.

Electronic manipulation of acoustics has enormous advantages as a research tool. It is possible to quietly wander through a hall during a rehearsal or a performance, and vary the acoustics with a remote control. Suddenly valid A/B comparisons can be made – and in the company of a skilled conductor or director the real – as opposed to imagined – effects of small modifications can be heard, and remembered.

Coming from the field of sound recording, and as a designer of reverberation equipment, I was all in favour of more reverberation than some halls provided. When I had the chance to work in opera houses I was rapidly disabused of this notion. My first settings for the reverberation enhancement in the Berlin Staatsoper were deemed way too strong by Mr. Barenboim. He insisted that the clarity of the singers be in no way reduced by reverberation.



Figure 1: Deutsches Staatsoper, Berlin

The solution was to increase the reverberation time and the reverberation level at frequencies below 500Hz – but to reduce the level and RT at upper frequencies. The result was that the extremely high engagement of the Berlin Staatsoper

was retained, while the orchestra, and the fundamental frequencies of the singers, gained the richness they had lacked. This opera house is still my favourite in Europe. Critics were very happy with the change – but oblivious to its acoustic origin. “Barenboim has managed to get the Staatscapella to sound like a real orchestra” was the comment.

Peter Lockhart, the assistant conductor of the Amsterdam Muziektheater, and I were in the Muziektheater stalls with a remote control during a rehearsal of “Siegfried”. We raised the strength of the reverberation above 1000Hz by 1dB. “Stop – that's too much” said Peter. “Why”, I asked. “Because the singer just moved away from us by 3 meters” he said. And he was right. In time I was able to hear this myself. It is difficult to do, because the visual impression of distance is so dominant. The effect of reflections and reverberation on sonic distance seems to be all or nothing. Either the singers are clear – and engaging – or they seem far, and less interesting. But you have to close your eyes to hear the difference as in increase in distance.

Hartmut Haenchen, the music director in Amsterdam, was conducting in the pit when I chose to raise the reverberation level by 1/2dB. He immediately waved to me from the podium, and told me to put it back. He had easily heard the increase, even standing right in front of the pit orchestra. I had similar experiences with Michael Schönwandt in Copenhagen.

All these experiences convinced me of the vital importance of immediately hearing the results of acoustic adjustments. I am convinced that without the ability to rapidly compare the small changes we were making we would not arrive at the settings we eventually used. We would probably have used more reverberation, and compromised engagement.

In fact, Haenchen had come to Amsterdam from the position of music director of the Dresden Semperoper. The Semperoper is far more reverberant than Amsterdam. From my seat in the front of the first balcony (see figure 4) the singers seemed far away. Balance between the singers and the orchestra was poor and the singers were not engaging. Haenchen initially wanted me to duplicate the sound of the Semperoper – just as Barenboim wanted something like the Festspielhaus in Bayreuth. But when they were able to A/B the difference in clarity such levels of reverberation created, they wanted no part of it.

DOCUMENTATION WITH BINAURAL RECORDINGS

There are other ways of performing instant A/Bs besides electronic architecture. I was working with a well-known acoustician when I noticed that he was recording all the performances he attended with microphones in his ears, and a small DAT tape recorder. I was impressed. He said did not think you could learn anything about hall acoustics without documenting what you thought you had heard.

Since then I have been attending operas and concerts all over the world, and recording many of them with tiny microphones stuck to my eyeglasses above my pinna. For the last three years I have been doing this with probe microphones which sit gently on my eardrums. The probe microphone system, when headphones are equalized with the same microphones, is capable of astounding realism, at least for me and about 50% of the people who listen to it. [7]. With this system you can A/B two different seats in the same hall, or two different halls.



Figure 2: Binaural system for recording the sound pressure at the eardrum. The microphones are behind the pinna, with a hard tube into the ear canal, equipped with a very soft silicone tip. When high quality on-ear headphones are equalized with the same microphones and a mathematical inverse filter exceptional realism results – without the need of head tracking.

This system has been invaluable. With it I have been able to convincingly record and play back the differences in engagement between one seat and another seat only a few rows back. The recordings made with microphones above the pinna are not as realistic, but they give a good general idea of the clarity of the sound, and the balance between soloists and orchestra. They also play back reasonably well through loudspeakers in a non-reverberant room. Any time I think I am on the wrong track I can go back and listen to some of the recordings. They are reassuring.

A major advantage of all these recordings is that they can be used as input to the measurement system described in part one. Work on measuring halls this way is in progress.

DRAMA EXPERIMENTS IN COPENHAGEN

Due to the success of the enhancement system in the old Royal Theatre in Copenhagen, Steve Barbar and I were asked to improve the speech intelligibility in a drama theatre – the “new stage” across the street. We had previously installed 64 loudspeakers in this theatre for use in opera, but the system was not intended for speech. We used a pair of line-array microphones to pick up the sound from the actors on stage, and I designed a sophisticated electronic gating circuit to remove as much as possible of the reflections and reverberation from the direct sound. The resulting signal was distributed through the speakers with appropriate sound delays.



Figure 3: Copenhagen New Stage

Note that the New Stage is a shoebox. Not the best shape for a drama theatre, as the average listener is too far from the

actors for good engagement. (This comment also applies to recital halls for chamber music.)

I turned the system on and off every 10 minutes during a live performance of Chekov in Danish with a full audience. Five of the major drama directors in Copenhagen were in the audience. At the intermission the directors were unanimous. “The system works – the actors are louder and more intelligible” they said. “We don’t like the system – turn it off.” “All right”, I said, “tell me why.”

At first they could not tell me. But I kept at it. Finally one said “The system makes the actors seem further away. I would rather the actors be unintelligible than sound further away. If the audience can’t understand them, the audience will listen more intently. This is just what I want.” The other directors agreed. They decided the solution to the intelligibility problem was better training of the actors.

Schönwandt was conducting “Tristan” next door. When he finished he wanted to know all about the experiment. He fully concurred with the results – and found the conclusion about training the actors highly amusing.

RECENT OPERA HOUSE EXPERIENCES

Moscow



Figure 4: The Dresden Semperoper from my seat in the front of the first balcony.

Much of the work described in these papers started when I was working in Moscow at a new opera theatre built next to the old Bolshoi. The new theatre was intended as a replacement for the old theatre which was scheduled to be rebuilt. The new theatre was modelled after the Semperoper in Dresden, which was rebuilt in 1983 after being destroyed in the second world war. The redesigned Semperoper eliminated the layers of fabric and other absorptive surfaces that were – and are – typical of European opera houses. I measured the fully occupied reverberation time at 1.6 seconds at 1000Hz – quite long for a venue with only about 1200 seats. I have a recording of “Arabella” from the front of the first balcony. The singers seem far away, and the balance between singers and orchestra is poor. The sound is reverberant, and the orchestra sounds good. But the singers are not engaging.



Figure 5: the new Bolshoi Theatre

The new Bolshoi is smaller, and more reverberant than the Semperoper. In addition there are strong focused early reflections from the curved side walls into the stalls. Singers sounded even further away than in the Semperoper. I had been asked to add reverberation to the opera house – which was not what it needed. I eventually got a conference with the general manager – who listened intently, and indicated he agreed. He had come from the Mariinsky theatre – and knew what drama needed. He indicated that more absorption would be on his to-do list, but that it might not be politically feasible. Everything is political in Russia.

With the help of a scalper I attended an opera performance and a ballet in the old Bolshoi theatre. Both were magnificent. The hall has 3000 seats – more real seats than La Scalla – and an occupied reverberation time of 1.2 seconds. The singers have enormous emotional power throughout the hall, and balance between them and the orchestra is overwhelmingly in favour of the singers. The ballet was surprisingly reverberant. The orchestra was playing the reverberation, and the very wide pit made the sound enveloping.

This is what an opera house should be! Too bad they are taking it down. I hope they don't rebuild it like the Semperoper.

Two recently completed opera houses

I attended performances in two recently completed European opera houses, sitting in two seats in the first one, and three in the second. I recorded what I heard with my probe system.

The first was a performance of “Siegfried” in a brand-new house. The sound in the back of the top balcony was distant, but reasonably clear. I needed – and did not have – binoculars to see the action clearly. Excess early reflections were blocked by the balconies, and there was a very prompt – and thus not problematic - reflection from the ceiling, which was only a few feet overhead. Balance between the singers and the orchestra was poor, as the approximately 1.6 second reverberation time of hall favours the orchestra over the singers. For the last act I managed to sneak into the stalls. There was a strong reflection bouncing from the surface below the first balcony to the side wall and back into the stalls. The orchestra appeared to come from the side wall instead of the pit. The singers struggled to be heard. I played the tape to one of my colleagues. “I can't understand the words,” he said, “and I'm German.”

In the second house my first position was in standing room at the back of the first balcony. The balcony shielded me from many of the early reflections and reverberation, and the sound was clear and engaging. I liked it. The next seat was in

the front of the second balcony, the third in the front of the third (top) balcony. I played these recordings to about 12 of my acoustician colleagues. All were impressed by the realism, and said what they heard in the recordings corresponded very well to what they had heard in similar seats. Nearly everyone preferred the sound in the front of the second balcony. The exceptions were me and the two who understood Italian. “This is not opera,” they said “the drama is lost if you can't hear the words.” I don't understand Italian – but I agree.

(Sound clips from this house can be heard in reference [3] in the third of these three preprints.)

EXPERIENCES IN LARGE CONCERT HALLS

In 2004 Leo Beranek asked me to join him for a talk at the 50th anniversary of the ASA. He wanted me to talk about concert halls – and I thought I knew too little about the subject. So I compared my recordings of violin concertos from three well known halls – Boston Symphony Hall, Avery Fisher Hall in New York, and the Kennedy Centre in Washington DC.

All these halls have similar cubic volume, rectangular shape, and similar numbers of seats. But they sound very different, with Boston clearly the best. I concluded at the time that the major differences were in the stage houses – but I did not test seats more than half-way back, and I was not convinced the stage house was the only reason they sounded different. These differences are discussed further in the third part of this talk.

A few years later I was asked to write a short mathematical article on concert acoustics for the IEEE.[8] Once again I felt I knew too little about the subject, so I did some experiments. I modelled the Boston Hall and the Concertgebouw in Amsterdam, and did binaural convolutions of the result with my own HRTFs. The models sounded very good, and plausibly like the real halls. But the two sounded quite different.

I swapped the late reverberation from one hall with the late reverberation of the other – and found it made no difference. The shape of the build-up of reverberation in the two halls was similar - but there was an additional time delay in Amsterdam of about 10ms. When I shortened this time delay to match the delay of the Boston hall the two models sounded identical. Amsterdam is almost square in plan. The greater delay of the side reflections, and the slightly closer average distance between the musicians and the listeners increases the clarity compared to Boston, but both halls create good engagement over a wide range of seats.

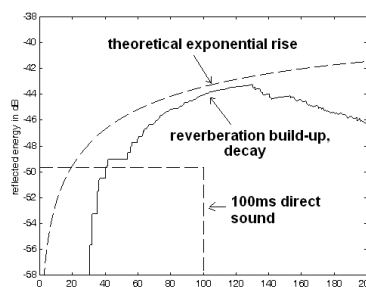


Figure 6: Reverberation build-up and decay from a 100ms excitation in a model of the Amsterdam Concertgebouw. The seat was chosen so that the D/R would be -10dB for a continuous excitation. Note there is more than 35ms time delay before the reflected sound pressure equals the direct sound pressure. This allows the brain to perceive the direct sound as

separate from the reverberation. Good azimuth perception and high engagement results. The value of *LOC* is +6dB.

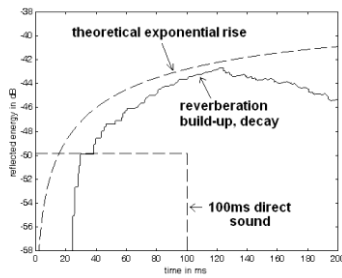


Figure 7: Reverberation build-up and decay from a 100ms excitation in a model of Boston Symphony Hall. The seat was chosen so that the D/R would be -10dB for a continuous excitation. Note that the initial delay is less than in Amsterdam. The sound is slightly less clear, but still engaging. The value of *LOC* is +4.2dB.

Clearly the rate at which reverberant energy builds up, and the level of this reverberation compared to the direct sound is an extremely important – and neglected – aspect of hall design.

As a further experiment, I took the model of the Boston hall and reduced its dimensions by a factor of two. The sound went from clear and reverberant to muddy and unpleasant. If too many reflections come too soon, the result is disastrous.

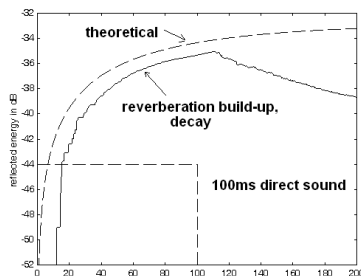


Figure 8: Reverberation build-up and decay from a 100ms excitation in a model of the Boston Symphony Hall with the dimensions all reduced by a factor of two. The seat was chosen so that the D/R would be -10dB for a continuous excitation. The value of *LOC* is +0.5dB. This hall has a reverberation time of under one second, but the sound is muddy and there is no engagement. Since the hall is smaller than the original, and the reverberation builds up more quickly, there is far more reverberant energy in the 100ms window that follows the onset of the direct sound. The direct sound is completely masked by the reflections.

As we will see in part three, the shape of the stage and the delay of the first lateral reflections are only part of the reason that Boston and Amsterdam are held in such high regard. I now believe an essential element of their success is the frequency dependence of the scattering elements on the walls and ceiling. These features increase the direct to reverberant ratio at high frequencies by directing high frequency reverberant energy downward to the front of the hall, where it is absorbed before it can travel to the back of the hall. The result is similar to the electronic enhancement in the Berlin Staatsoper. Reverberation and engagement can co-exist.

EXPERIENCES IN SMALL CONCERT HALLS



Figure 6: The stage of a small concert hall that lacks clarity and engagement.

A little more than a year ago I had the opportunity to work in a 300 seat chamber music hall that had problems with muddiness and lack of engagement. If you sat close enough to the stage to hear the sound clearly the sound was too loud. If you sat further back it was difficult to localize and separate instruments. The occupied reverberation time was about one second – so by conventional thinking the sound should have been clear. But it was not. In the opinion of several of my acoustician friends the lack of clarity could be solved by adding a small shell behind the musicians. My work with models suggested the opposite. I felt the problem was that the first reflections came too soon, and were too strong. The solution would be to add absorption to the back wall and side walls of the stage.



Figure 7: The same hall with some absorption added to the back of the stage.

We tried both solutions. The small shell made the muddiness worse throughout the hall. Adding a small amount of absorption – just panels at the bottom of the back wall - made an enormous improvement. The dean of the music school that owned the hall said he had never heard what a difference a seemingly small acoustic modification could make to the power of the music. The piano sounded like a newer and far superior instrument. There are plans to permanently install some absorptive panels at the back of this stage.

In a new European hall with about 350 seats I recorded a violin-piano concert in several different rows by discretely hopping across rows between pieces. (The performance was sparsely attended.) Clarity was good in rows 1 to 3, poor in rows 4 through 13, and for some reason was good again in row 14, the last row. But to my ears most of the seats in this hall were muddy and unclear.

The hall was specified to be a shoebox with a 1.7 second or greater reverberation time, and the specification was met by the acousticians that designed it. But a small shoebox hall with a long reverberation time is almost guaranteed to have too much early reflected energy. The hall must have very little absorption to meet such a criterion, and consequently the reverberant level will be high compared to the direct sound. The shoebox shape brings the side walls close to the audience, decreasing the delay of the first reflections. The shoebox shape also increases the average distance from the musicians to the listeners, further decreasing the D/R ratio. To make things worse, in this hall the stage house is large, has a flat back wall, and is completely devoid of absorption. The reflection from the stage back wall is prompt and strong. It adds to the prompt and strong early reflections from the side walls and ceiling. Poor engagement is the inevitable result.

A well-known acoustician who was with me agreed with my assessment of the sound, and the binaural recordings clearly showed the differences in engagement between the rows. But the shoebox shape is firmly enshrined as the ideal shape for a concert hall regardless of size, and more small halls of this shape are being built every year. The average seat in these halls does not encourage a listener to make live chamber music concerts a frequent part of their life.

I used the probe microphones to record the sound from a string quartet in two seats in London's renowned Wigmore hall. The hall is a long, narrow shoebox. I was convinced I could localize at least some of the instruments in the quartet from a seat half-way back into the hall, and could occasionally localize them from a seat two-thirds of the way back. But when I listened later without the visual image I was surprised by the poor localization, and due to a strong prompt reflection from the side wall half of the time the sound was actually coming from my right side. Engagement in the recording was poor. I sat much closer in a previous visit to the hall – and loved the sound. This hall – like many others – has a deservedly good reputation, but only for the best seats. I have not yet attempted to measure the engagement in these recordings. I do not expect it to be high.

A medium sized (1000 seats) shoebox hall

The problems that accompany a long reverberation time and prompt early reflections do not go away as a hall gets larger. I recently attended several concerts in a new medium sized shoebox hall, and recorded a string quartet in two different seats in the stalls. Once again the reverberation time was specified as at least 1.7 seconds. A seat in row F provided clear localization and engagement for all the instruments in the quartet. The sound was beautiful and exciting. In row K the clarity was lost – all the instruments blended into the centre of the sonic image, and the inner voices were often inaudible. The music was nice, beautiful in a way, but it lacked the excitement of row F. Row K was less than half-way back into the hall – which does not bode well for the rest of the seats.

I played these recordings for one of the designers of the hall, who clearly heard the difference – but claimed that it was unimportant. “The hall looks and sounds like a concert hall. If a listener wants a clear sound he can sit in row F. Otherwise he can choose row K.” But the sound in row K was not more reverberant than row F – if anything it was less reverberant. Why would anyone choose it? Perhaps you could call the sound “well blended” but I prefer Beranek's description of such seats; “you can sell them to tourists.”

As we will see in part three of these talks, large, medium, and small halls exist that really do provide exciting, reverberant, sound to a large majority of seats. Some of these are old and very much admired. We can build new halls that sound as good – but we will have to give up some of our favourite myths. We should be more concerned about whether we can hear all the music, not just some of it. Our convictions about the ideal shape of a concert hall need adjusting.

ACKNOWLEDGEMENTS

The experiences described in this paper would not have been possible without the participation, encouragement, and support of friends and colleagues too numerous to name. But I would especially like to thank Steve Barbar of Lares Associates for his passionate commitment to electronic acoustics, fine ears, and his unlimited appetite for hard work. I also thank Leo Beranek for his continually generous encouragement and support.

REFERENCES

- 1 M. Barron: “Auditorium Acoustics and Architectural Design” E&FN SPON 1993.
- 2 W. Gardner and D.H. Griesinger, “Reverberation Level Matching Experiments” Proceedings of the Wallace Clement Sabine Centennial Symposium, Cambridge MA 5-7 June 1994. (a copy of this paper is on the author's website.)
- 3 D.H. Griesinger, “The effective loudness of running reverberation in halls and stages” – *ibid*. A similar paper “How Loud is my Reverberation” is on the author's web page.
- 4 D.H. Griesinger, “The psychoacoustics of apparent source width, spaciousness & envelopment in performance spaces” Acta Acustica Vol. 83 (1997) 721-731. (this paper is on the author's web page)
- 5 S. Olive and F. Toole, “The Detection of Reflections in Typical Rooms” Preprint 2719 from the Audio Engineering Society Nov. 1988. Available from the AES.
- 6 F. Toole “Sound Reproduction, Loudspeakers and Rooms” Book available from Amazon.com
- 7 D.H. Griesinger “Frequency response adaptation in binaural hearing” A powerpoint presentation given at the May 2009 meeting of the ASA in Portland OR. Available on the author's web-page.
- 8 D.H. Griesinger, “Concert Hall Acoustics and Audience Perception” IEEE Signal Processing Magazine [1] March 2007