



HOW LOUD IS MY REVERBERATION

David Griesinger

Lexicon

100 Beaver Street

Waltham, MA 02154 dg@lexicon.com

Abstract

It is vital to know the loudness of reverberation in rooms, but this is poorly predicted by current theory and measures. We present the results of a continuing series of experiments on the loudness of running reverberation. We propose a method for determining the reverberant loudness from a measured impulse response, and provide some insight into how reverberation is perceived. Optimum reverberant levels for music depend strongly on the reverberation decay shape and the degree of musical masking. When masking is high small changes in reverberant level can have large perceptual effects.

Introduction

Running Reverberation (RR) - the reverberation one hears while the music is playing - is one of the most important perceptions in musical acoustics. Yet there is no universally accepted method of describing or measuring the amount of RR. Ideally we would like a description or a measure which corresponds in a fundamental way with the biology of human hearing. A good test of such a measure is the ease and reliability with which subjects can make comparisons. We find that subjects can easily make judgments about the relative loudness of reverberation in two different segments of recorded music. In particular, they can match the loudness two different types of artificial reverberation on the same musical source, and they can do this comfortably and with high repeatability. Because the concept of loudness seems natural to the subjects, and because they perform well on matching experiments based on loudness, we have some confidence that reverberation loudness is a fundamental property of hearing.

The lack of a reliable measure of reverberant loudness is especially serious in designing and testing halls and rooms for musical performance. For sound engineers the most critical factor about reverberation is the reverberant level - the position of the reverberation level slider on the mixing board. If the reverberation is not audible enough, we simply increase the level. It is nonsense to think that listeners in the concert hall are



different, and yet acousticians generally ignore the concept of reverberant loudness. They speak only of reverberation time, believing that if the reverberation time is in the range recommended for particular types of music that the sound will be fine. This does not always work.

In room acoustics the time it takes for the sound to decay 10 decibels (multiplied by six to be comparable in magnitude to the reverberation time) is usually used to describe reverberant level. This measure is called the Early Decay Time, or EDT. Reverberation level matching experiments at the Massachusetts Institute of Technology and in our laboratory show that the reverberant profile - especially the amount of pre-delay - and the reverberation time have a large effect on our perception of the loudness of running reverberation. EDT matches the perceived loudness only for a subset of the data. Where the reverberation time is about 2 seconds, the decay is exponential, and the Clarity 80 is about 0db, EDT works reasonably well. When the direct sound is stronger, or if there is a significant pre-delay or plateau in the decay curve, EDT works poorly, and this is often the case in concert halls. While our current work has yet to reveal a comprehensive measure for reverberant loudness, a simple modification to the EDT - where the slope of the first 350 to 380ms of decay is used instead of the slope for the first 10dB of decay - may be a more useful measure in halls.

The level matching experiments show that RR is a specific perception, separate from room impression, intelligibility, sharpness of localization, etc. It arises from the ability of the hearing mechanism to parse sound events into foreground and background. When RR is easily heard this parsing is independent of the spatial properties of the sound field, and is mostly determined by sound energy more than 150ms after the ends of musical notes. It is usually musically optimal for RR to be partially masked by musical material. Thus the optimal level depends on the type of music. When masking is high small changes in reverberant level can have large perceptual consequences.

Musician self support

This work was started through a study of musician self-support on concert stages. A series of binaural recordings using microphones just above the pinnae of the performer was made of solo recorder, French horn, and voice in various spaces. It was noticed while making the recordings that some spaces gave a similar musical impression to the player, even though the rooms were very different in size. For example, similar self support resulted from playing in a room of 200k ft³ with a 2.0s RT60, and in a 2k ft³ space with an RT of 0.5s. In general while playing music it is not possible to hear the details of the impulse response of the room. Although reverberation can be easily heard, the impression is generic. We found we could duplicate these tapes by adding artificial reverberation to anechoic binaural recordings of the same player, and adjusting the reverberator to match the

reverberant profile as measured by recorded handclaps.




Figure 1. Waterfall pictures from handclaps of two rooms with similar running reverberance. Top:



200kft³ 2.0sec RT Bottom: 2kft³ 0.5sec RT. These two curves have quite different C80, RT, EDT, etc.

Bill Gardener (a Ph.D. candidate at the Massachusetts Institute of Technology) and I performed a series of preliminary experiments where a subject varied the direct to reverberant ratio to match the apparent loudness of different reverberation profiles. When the two reverberations had a very different RT subjects complained that the sounds were obviously different, but even with this observation they found there was a perception similar enough that it made sense to match levels in this way. The experiment gives similar results with the same signals in both earphones (mono), indicating that the perception of reverberance depends on the way the brain parses information, and not on the spatial properties of the field. Bill is conducting a series of controlled listener experiments on the phenomena we uncovered (1). In analyzing the data from the preliminary experiments we noticed that plots of the 40ms widow integrated impulse responses of two spaces with similar running reverberance tend to cross at about 180 to 200ms. We hypothesized that RR could be predicted by integrating the impulse response over a rectangular window centered on this time, and comparing this value to the value from a similar window integration beginning with the direct sound.

Experiments with single pitches

Here are results from experiments with a musical signal of single repeated notes of variable length.

1. Loudness of a musical note of constant level repeated once per second increases with note duration rapidly in the time range of 10 to 50ms, and then gradually (about 6dB more) as the duration increases to 200ms. These times appear to be independent of the frequency of the note. This is consistent with psychoacoustic data from Zwicker. (3).
2. When presented with a stimulus consisting of a single frequency tone which alternates between two different levels with various timings and duty cycles, (typically ~160ms in the high state) we perceive primarily the loudness of the LOWER level tone. The perception is that of a constant tone being disturbed by a louder tone. This ability of the hearing mechanism to focus on the level of the lower level tone is quite remarkable.
3. If we think of the higher level tone as a series of musical notes of the same pitch, and the lower level tone as the reverberation of these notes, the audibility of this reverberation depends critically on the spacing between notes. When this spacing is shorter than 30ms the ability to hear the loudness of the reverberation disappears rapidly. When the gap width increases from 50ms to 200ms the loudness of the lower level tone rises about 12dB.
4. When the low level tone is replaced by exponentially decaying but randomly varying reverberation, the perception remains that there is a tone of constant level interrupted by louder tone bursts. This is so even though the level of the reverberation in the gaps is not at all constant, and can be quite different from gap to gap.

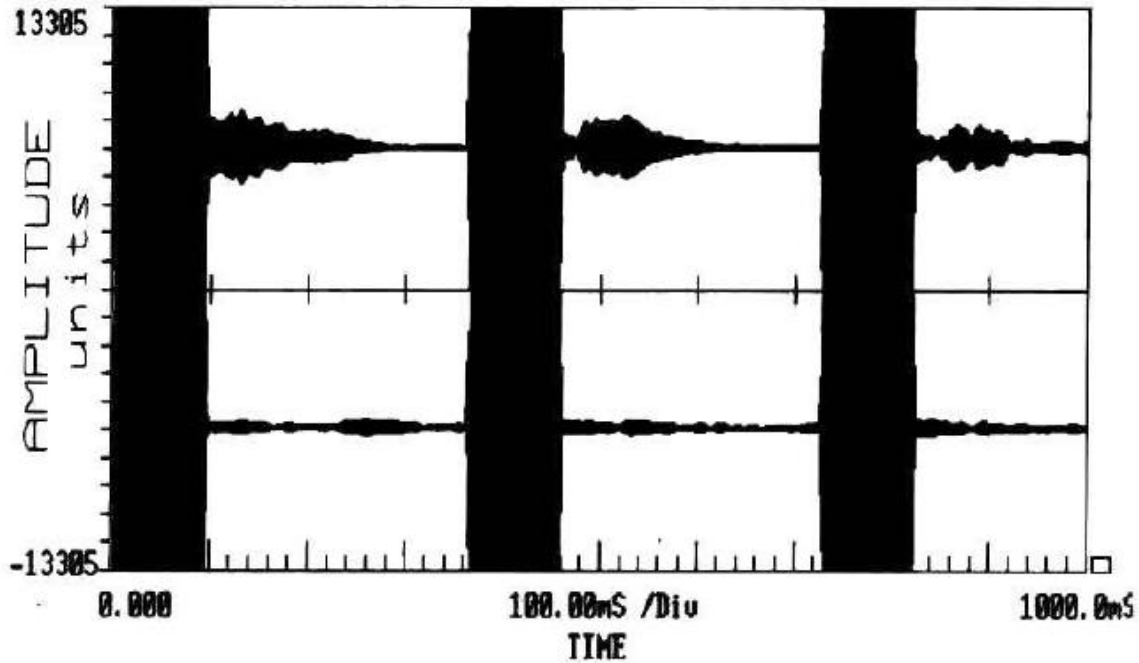


Figure 2: Reverberation from 100ms tone bursts of equal loudness. Top 0.5s RT, bottom 2.0s RT. Vertical axis increased a factor of two to make the reverberation more visible.

5. When we compare the loudness of RR from a 0.5s exponential decay to the loudness from a 2.0s decay we find:

a. When the gap is short - under 50ms - the RR loudness is simply proportional to the total energy in the gap. A 0.5 second decay and a 2.0 second decay of equal energy in the first ~160ms are equally loud.

b. When the gap is over 200ms in length the total energy of the 0.5 second decay must be about 10dB higher for equal RR loudness. This result is quite extraordinary when these two signals are seen on an oscilloscope. It would seem that the hearing mechanism inhibits the detection of most of the level from the 0.5 second decay, or assigns it to the note itself. Only after a time delay of 160ms or more is the energy assigned to the RR perception.

6. When the pre-delay of a 0.5 sec RT is increased from zero we see:

a. With simple stimuli of constant note length as long as the pre-delay is less than the length of the note the energy in the gap increases smoothly as more of the reverberant energy appears in the space between notes. This is a physical effect - instead of being masked by the continuing direct sound the reverberant energy is moved into the gap where it can be heard.

b. When the gap between notes is on the order of 200ms the loudness increases with pre-delay somewhat less than one would expect if the level at 180ms or so after the end of the note was all that mattered.

Solo instruments as sources



Musician self support appears to be identical to the perception of running reverberance. We ran the level matching experiments with solo recorder (Bach Partita II BWV 1004, Allemande), French horn (first movement Mozart K447), and tenor voice as sources. All these selections gave results similar to each other and to the experiments reported in (1) using anechoic clarinet and sax. For most individuals data from these matching experiments is well predicted by the ratio between the direct sound and the reverberant level at about 180ms, particularly if the experiment is limited to exponential decays in the range of 0.5s RT to 2.0s. This result is surprising and encouraging. When pre-delay is added to a 0.5s RT some individuals report the expected 6dB increase in level for each 50ms of pre-delay. However on the average the increase in level with pre-delay is less - 3-4dB for every 50ms (1). In addition very little increase in loudness was noticed for RT > 2 seconds.

Measures Based on the Impulse Response

Current measures of musical acoustics tend to be based either on the impulse response itself (C80) or on the backwards integrated impulse response (RT60 and EDT). The impulse response represents the sound of the room when it is excited by a pistol shot. These are rare in music. The backwards integrated impulse response represents the decay of a very long note which abruptly stops. Such long notes are also rare in music, although they are not as rare as pistol shots. More commonly music excites rooms with notes of finite length. Fortunately it is easy to smooth the impulse response in such a way as to represent the response of the room to a note of finite length, and we propose to search for measures of reverberance with curves smoothed in this way. Such a smoothing has nothing to do with the physiology of hearing - we have simply made a decay curve which more closely approximates the response of the room to typical music.

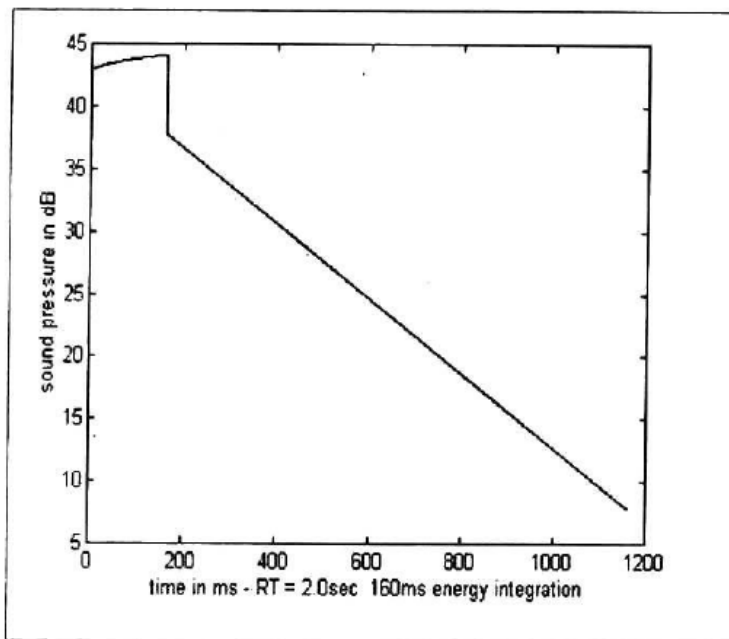


Figure 3: 160ms window integrated curve

In a simple analogy to the Schroeder integration, if we sum the power in an impulse response which lies inside a sliding window of the length of an average note, we get a curve like Figure 3. Note that the amplitude increases as the note plays, and then decays after the note ends.

$$RR160 = \frac{\int_{160ms}^{t+160ms} p(t)^2 dt}{\int_0^{160ms} p(t)^2 dt}$$

The observation that curves made with a 160ms integration time tend to cross at about 160ms if two spaces have similar musician self support leads to proposal that musician self support can be measured with RR160:



The 160ms window for integration was chosen for two reasons - it is a passable match to the loudness integration time for the direct sound, and it matches the note length used in several of the experiments. Other windows (100ms, 250ms, etc.) also give reasonable results, at least to the precision in our data.

Level matching with quartet or orchestral music

A similar series of experiments was performed with stereo quartet and orchestral music. It was found that much higher reverberant level had to be used to make the reverberation sufficiently audible to get reliable results. For example, the total energy of a 2.0sec exponential decay must be at least -8dB to be sufficiently audible, and even at this level the scatter in the data is large. With anechoic orchestral material listeners preferred a level of -2dB to -4dB musically. Such a high level makes it not possible to obtain equal reverberance with a 0.5s exponential decay. It is possible to compare 1.0s exponential decays and longer. A set of reverberation curves was prepared which sounded "equally reverberant" and measures of these curves were compared. RR160 does not give a particularly good match. Several other measures are more promising, although no single measure appears to match the data precisely. Perhaps the best match to the data when the reverberant energy is very high (C80 -2dB to -6dB) is similar to EDT, but based on TIME and not on a certain number of dB of decay. For example, for a series of curves of equal reverberance where the 2.0s RT has a C80 \approx 0dB, the Schroeder integrals of these curves tend to be equal in level at about 350ms. This happens to correspond to about -10dB from the peak value. Curves where the 2.0s RT curve has a C80 of about -6dB also tend to cross at about 350ms, but this corresponds to about -6dB of decay.




Figure 4. - Equal loudness reverb curves - 1sec RT + spread vs 4sec RT

It appears that notes in orchestral music are frequently longer than 160ms, and the Schroeder integral - which represents the decay of the room when excited by a very long note - is more appropriate than the 160ms power



averaged curves which seem to work for solo music. Early Decay Time is usually measured by using a least-squares method (linear regression) to fit a straight line to the points in the Schroeder integral between the peak value and a point 10dB below peak. There are problems with this method, in that it does not weight the peak value enough when there is significant direct sound. We propose a simpler method for determining the running reverberance of orchestral music. We start with the Schroeder integral, and simply determine the slope of a line which connects the peak value and a point 350ms later. If we express this as an equivalent reverberation time,

if $S(t)$ is the Schroeder integral in decibels vs time, with the peak value at $S(0)$, then

$$RR(orchestral) = \frac{60 * 350ms}{(S(0) - S(350)) * 1000ms / sec}$$

Level matching with speech

With speech as a source and a reference of 2.0s RT at -20dB total energy the reverberation is highly and continuously audible. It is possible to match it with very little scatter in the data. Bill Gardner reports results similar to the experiments with solo music, with some differences in behavior with pre-delay. He makes the observation that with pre-delays of 100ms or more with the 0.5s RT, and for the reference reverberation of -20dB level and 2.0s RT, the reverberation separates into a continuous perception the level of which can be matched quite precisely.

Masking of reverberance by music

The level matching experiments indicate that given a type of music, the audibility of different reverberant profiles can be predicted in a simple way. When we analyzed the anechoic alto recorder tape which started these experiments we found each note was typically about 180ms long, and there was a 50ms to 70ms gap between each. With this tape a 2.0sec RT at C40 = 20dB was highly audible - it could be heard nearly all the time while playing. We then tried a synthesized ascending and descending scale with 70ms gap between the notes. With this scale the reverberance was much harder to hear than with live Bach, and level matching experiments gave results consistent with 5a above.

Further experiments revealed that although scales and short gaps mask reverberation, arpeggios and jumps upward of more than a third tend to unmask it. For example if a sequence of C, E, G is played at a tempo of 5 notes per second, reverberation from the C is clearly audible during the G, although the reverberation from the E is masked. In this example the E serves to mark time. By the time it is finished the melody has jumped to a different critical band, and enough time has elapsed that the brain can identify the reverberation from the C as RR.



We tested a tape of various pieces played on a synthesizer with no gaps between the notes, and found that at direct to reverberant ratios typical of musician self support the reverberation was nearly inaudible. The lack of spaces between notes raises the threshold for detection, as does making the reverberation come from only from the medial direction. This directional dependence is not significant in practice, since by 180ms or so after the direct sound nearly all acoustic spaces are directionally diffuse.

As reported in (1) naive subjects can be asked to match reverberance by varying the direct to reverberant ratio between two different sources and the same reverberant profile. The perceived loudness depends on the type of source. These results suggest that masking is critical to our perception of reverberance. Speech masks reverberance very little, and direct to reverberant ratios of 20dB sound identically reverberant to solo music with a ratio of 10dB. Orchestral and string quartet needed an additional 4 or 5 dB of reverberant level to be equally reverberant. Even with orchestral music it is the author's experience that quite different direct to reverberant ratios are desirable, depending on the composer and the instrumentation. We have spent a great deal of time working with musicians at the Tsai Center at Boston University to determine the ideal reverberant level. Solo piano requires about 3dB less level at an RT of 1.5 seconds, than a string quartet at an RT of 1.7 seconds. A romantic symphony requires another 3dB, and an RT of 2.0. There is an obvious tradeoff between the reverberant level and the predelay, as adjusted by the "spread" control on the electronic system. Once the right level has been found there is almost always unanimous agreement that the degree of reverberance is correct for the piece. At the chosen level the reverberation is partially masked by the music - neither inaudible nor audible all the time.

With partial masking the range of adjustment for optimum reverberant level is small. Slightly too little and the reverberation is seldom audible - slightly too much and the reverberation will start to mask inner voices in the music. When level is adjusted for optimum for highly masking material (Bruckner) a slight increase in level makes the RR seem very loud. Although for untrained listeners the loudness of RR appears to depend on masking it may be possible to separate the absolute loudness of RR from its audibility. One need only learn to wait for gaps in the music to judge absolute loudness.

Recent results

We have made a computer model of human hearing for reverberation. The model consists of a series of 1/3 octave filters followed by level detectors which can find the start and ending points of musical notes. We can play anechoic recorded music into the computer and ask how often is reverberation from the various notes in the music audible, and how often is it masked. The model is mostly written in MATLAB, and is available to anyone

interested. (Inquire by email.)




Figure 5: Fraction of time reverberation is audible for a section of full orchestra, as a function of direct to reverberant ratio (in dB), and third octave



band. For example, here is a graph of the fraction of the time reverb is audible as a function of frequency and reverberant level for a short segment of orchestral music. Notice that when the direct to reverberant ratio is about one the reverberation is highly audible, but as the reverberant level decreases the masking increases rapidly. In fact, a one decibel decrease in reverberant level causes a nearly four decibel decrease in reverb audibility.

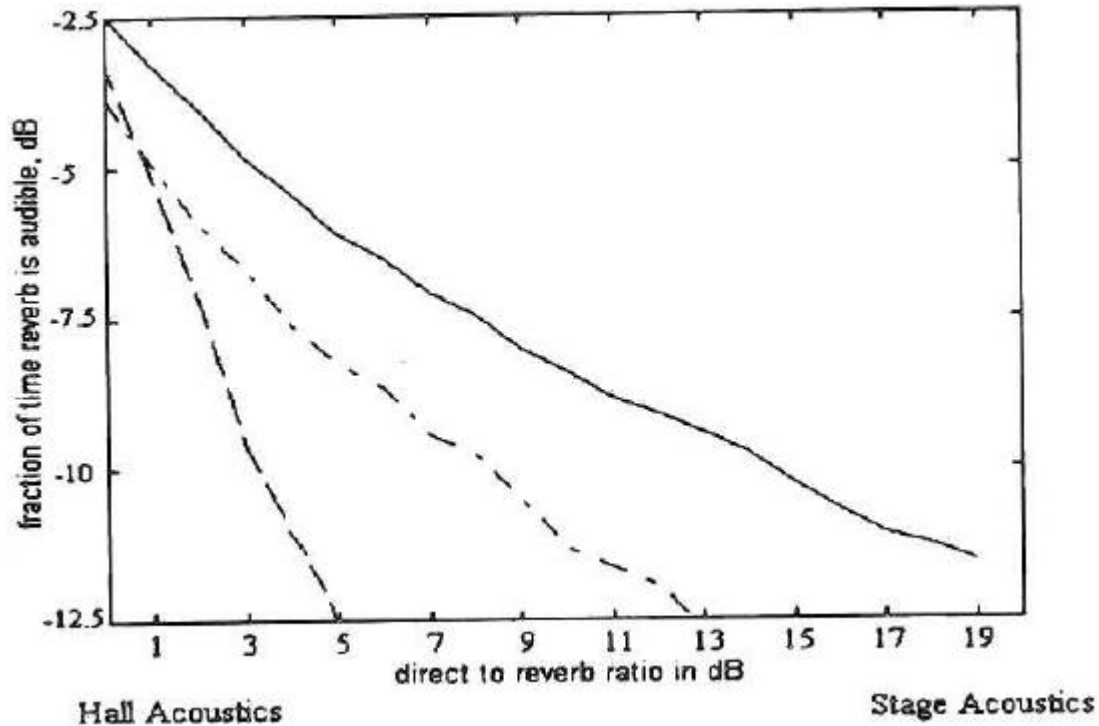


Figure 6: fraction of time reverb is audible in dB ($10 \cdot \log_{10}$), for three different types of music. Average of bands from 500Hz to 2kHz.

_____ = solo recorder

-. = string quartet

- - - - - = full orchestra

We can average the frequencies from 500Hz to 2kHz. Figure 6 shows the fraction of the time reverb is audible as a function of level for three different types of source material - solo recorder, string quartet, and orchestra. Notice the increase in masking for recorder is low, giving significant audibility to the performer while he is playing. Audibility decreases much more steeply for string quartet, which makes the selection of the ideal reverberant level much more critical. Orchestral music is the most critical. You must have the reverberant level just right. Unfortunately I have not yet tried these experiments on popular music.




Figure 7: Full orchestra masking for three different reverb profiles. Average of bands from 500Hz to 2000Hz.

The degree of masking depends on the decay profile. We can see this in figure 7, where masking is plotted against reverb level for three different profiles - 2.0s RT, 3.0s RT, and 3.0s RT plus 120ms pre-delay. Notice that as reverb time and pre-delay increase the reverberation becomes much more audible. For recorded music, using a larger value of "spread" increases the reverberation audibility, allowing a lower reverberant level to be used. In concert hall acoustics increasing the reverberation time through hard floors and seats, and adjusting the hall geometry to increase the late arriving energy will also increase the running reverberation.

Concert Halls

Reverberant level in modern recordings is adjusted by ear. This adjustment is critical, but is done by engineers with remarkable consistency. The reverberant level in most recordings is adequate to give a nearly continuous audibility to running reverberation, and yet is low enough that the music is only very rarely obscured. This is often not the case in concert halls. In some modern halls many seats have lower running reverberation than recordings. The hall is audible, but early energy predominates over late energy. The sound is loud, colored, and not reverberant except during pauses in the music.

Many acousticians design the hall to reflect the energy from the orchestra into the audience. These reflections increase the loudness of the direct sound, but the audience is absorptive. Once the first reflections hit the public the sound energy they contain is gone. There will be no energy to come back to the audience later as running reverberation.




Figure 8: Old fashioned hall with low balconies and coffered ceiling. Note sound survives the first few reflections, and will be heard later as running reverberation.

first bounce, and returns in the direction of the stage. From there it can go back to the audience, and since the time delay is now greater than 150ms, it will be heard as running reverberation.

For example, here is a simplified drawing of an old-fashioned concert hall, with gently sloping balconies and a coffered ceiling. Notice that a great deal of the sound from the performer survives the

This old fashioned hall can be "improved" by making the balconies steeper to improve sight lines, and by adding reflectors to the ceiling to increase speech intelligibility in the upper balcony. The new hall will have the same reverberation time as the old hall. However a much larger percentage of the sound will be absorbed in its first encounter with the audience, and the running reverberation may be inaudible.




Figure 9: "Improved" hall with better sight lines, higher speech intelligibility in the balcony, and much lower reverberant loudness. The reverberation time is the same as figure 8, but the success of the hall with orchestral music is questionable.

In sum, you can design a hall for high reverberant loudness, but only at the cost of a less forceful direct sound, poorer sightlines, and/or harder seats. In general a hall designed for orchestral music will be richer but softer than an equivalent general purpose hall. If the hall is large - more than 2500 seats - you usually cannot have both a loud sound and high running reverberation without electronic help.



Questions to be answered

In (4) the suggestion is made that for orchestral music it is often optimal to make running reverberance stronger at low frequencies. Recent work at Boston University indicates this is indeed useful. The experiments reported in this paper have been primarily concerned with solo music and have used frequency independent reverberation. We need to know much more about the frequency dependence of the perception of RR. The dependence on presentation level also needs investigation. This paper is primarily concerned with RR when the direct sound is within -2dB of the total reverberant energy or higher. In another paper (8) we investigate the case where the reverberation is stronger - a case which occurs often in large halls.

Conclusions

Running reverberation is a perception related to a tendency to sort sound events into a fluctuating foreground and a relatively continuous background. The sorting process depends on the direct to reverberant ratio, and the decay profile. Reverberant level at times greater than 150ms are particularly important. Two measures based on the impulse response are proposed, one for direct sound levels found on stages, and one for the lower levels typical of halls. To a naive listener the apparent loudness of RR is influenced by the degree to which RR is masked by the music itself, and the musically ideal RR level is one where the perception is partly masked. It is possible to use a simple hearing model to predict the degree to which anechoic recorded music will mask reverberation. Such analysis shows the optimum reverberant level to be especially critical for orchestral music.

References:

1. "Reverberation Level Matching Experiments," William Gardner, Proceedings of the Wallace Clement Sabine Centennial Symposium, Cambridge MA 5-7 June 1994. These proceedings are highly recommended. They are available from the Acoustical Society of America
2. "Subjective loudness of running reverberation in halls and stages" David Griesinger, *ibid.*
3. "Progress in Electronically Variable Acoustics," David Griesinger, *ibid.*
4. "Psychoacoustics," E. Zwicker and H. Fasfl, Springer-Verlag, 1990
5. "Quantifying Musical Acoustics through Audibility," David Griesinger, Knudsen Memorial Lecture, Denver ASA meeting, 1993. Copies available from the author
6. "Subjective preference in relation to objective parameters of music sound fields with a single echo" Y. Ando, JASA vol. 62 p1436-1441 (1977)
7. "Accents in equitone sequences" D-J Povel and H. Okkerman, Perception & Psychophysics 30 (6), p565-572 (1981)
8. "Further investigation into the loudness of running reverberation" David Griesinger, Proceedings of the Institute of Acoustics conference. Gatwick, UK Feb. 1995