Center for Computer Research in Music and Acoustics July 1987

Department of Music
Report No. STAN-M-41

HIGH LEVEL FACTORS AND THL MUSICAL SALIENCY OF AUDITORY PHENOMENA

by

John R. Pierce & Earl D. Schubert

CCRMA
DEPARTMENT OF MUSIC
Stanford University
Stanford, California 94305

Center for Computer Research in Music and Acoustics July 1987

Department of Music
Report No. STAN-M-41

HIGH LEVEL FACTORS AND THE MUSICAL SALIENCY OF AUDITORY PHENOMENA

by

John R. Pierce & Earl D. Schubert

This is the text of a proposal submitted to the National Science Foundation for research. For some auditory phenomena, including residue pitch, what is heard in a musical context when played over speakers in a reverberant room seems not to agree with the published outcome of experiments in which subjects wearing headphones made various matches or adjustments. Four initial experiments are proposed, 1) to investigate the conditions under which residue pitch is clearly musical pitch, 2) to investigate and endeavor to remedy the "buzziness" of computer-generated sounds, 3) to investigate musical perception of combination tones, and 4) to investigate the recognition of a constant spectrum when he fundamental is changed.

1 PROJECT SUMMARY

Center for Computer Research in Music and Acoustics
Department of Music

Stanford University
Stanford, California 94305

Dr. John R. Pierce, Dr. Earl D. Schubert

HIGH LEVEL FACTORS AND THE MUSICAL SALIENCY OF AUDITORY PHENOMENA

For some auditory phenomena, including residue pitch, what is heard in a musical context when played over speakers in a reverberant room seems not to agree with the published outcome of experiments in which subjects wearing headphones made various matches or adjustments. Further, some auditory stimuli are consistently judged differently by different listeners. Here we propose to investigate the musical saliency or applicability of various auditory phenomena by playing recordings of musically acceptable sounds in a musically acceptable context to musically trained people. The material we record and play will in general evoke clear and definite responses, though not always the same responses from different subjects. We propose four initial experiments, and others to follow. The first will investigate the conditions under which residue pitch is clearly musical pitch. The second will investigate and endeavor to remedy the "buzziness" of computer-generated sounds. The third will investigate musical perception of combination tones. The fourth will investigate recognition of a constant spectrum when the fundamental is changed. Among other topics we hope to investigate are salient effects of phase, various aspects of masking, and other phenomena which come to our attention during the course of the work.

2 PROJECT DESCRIPTION

HIGH LEVEL FACTORS AND THE MUSICAL SALIENCY OF AUDITORY PHENOMENA

This proposal focuses on the musical saliency of auditory phenomena. The approach is to play stimuli to subjects in a musical context and a reasonably musical environment. The results sought will be clear cut but we expect them to differ from subject to subject. Subsequent matching experiments may be carried out, but only after their pertinence has been clearly established in a musical context.

The initial experiments chosen are to find when the residue pitch is equal to the pitch heard by musicians, to investigate and endeavor to remedy the "buzziness" of computer-generated sounds, to investigate the musical saliency of combination tones, and to investigate the recognizability of a constant spectrum when the fundamental is changed.

2.1 General Background

A number of salient auditory phenomena evoke unequivocal responses from subjects with various degrees of musical training. These responses can be evoked in a musically natural setting, by playing a good recording in a moderately reverberant room. Yet, the response, though consistent for a particular subject, may differ among subjects, so that averaging results over subjects could hide the very information that is sought. It would conceivably be highly useful for the composer, especially of electronic music, to know whether certain phenomena evoke different responses from different listeners.

In order to illustrate the field in which we propose to work, we will note that among such salient auditory phenomena are judgments as to whether the pitch of the second of a pair of notes is higher or lower than the first [Deutsch, 1986], the relation between handedness and what is perceived in the octave illusion [Deutsch, 1983], and the preferred intonation of the tones of a triad [Roberts and Mathews, 1984]. In a presentation on residue pitch and difference tones [Pierce, 1986] there was general agreement on some residue pitches and differences on others, but responses differed markedly from the literature on residue pitch.

Music is commonly listened to in a moderately reverberant room. Recorded music, including recorded computer or electronic music, is commonly listened to as played over speakers in a moderately reverberant room. Many well-known auditory phenomena have been investigated chiefly under laboratory conditions in which the subject listens over headphones in a very quiet environment and, under instruction, makes adjustments for matches or gives

responses for signals that are very different from musical sounds in a musical environment.

Our proposal is for an extended study of the actual relevance of various auditory phenomena and data to the concerns of the composer or the instrument designer. Phenomena must be explored and data gathered for sounds from speakers in typically reverberant rooms as well as for sounds heard through earphones. Sounds must be heard by musically trained subjects in a musical context or a context not alien to music. Both the subject and the experimenter must be aware of what they hear, and not merely of knob settings. Sound examples must be recorded digitally, so that there is a permanent record of stimuli, and in case of dispute results can be demonstrated.

In Helmholtz's great book, psychophysical and psychological experiments are motivated by and intertwined with musical knowledge and observations; indeed, the book endeavors to explain musical phenomena through experiment.

Pierce's book, The Science of Musical Sound [1983], presented new experimental material, in recorded and textual form, on various musically salient phenomena, including stretched octaves with stretched partials, but it also drew on much published auditory work without establishing the salience of the work described in the elucidation of musical phenomena. Much of the data cited was drawn from experiments far removed from the experience of music. There has been commendable experimental work that is closely related to the musical experience [Deutsch, 1975, 1982 ; Bregman 1981, 1984].

The admirable set of sound examples produced by Green and others [Green, 1978] enable a listener to evaluate a number of striking auditory phenomena by listening, including critical bandwidth, echoes, masking, temporal integration, categorical speech perception, lateralization, loudness, and combination tones. It is interesting to note, however, that only one of these examples makes use of a melodic sequence of tones.

Nonetheless, a very extensive literature gives inadequate insight and guidance in the generation and admixture of musically useful sounds. Indeed, experiments, even when they can be replicated, appear in some cases to be misleading concerning the perception of sounds in a musical context.

In part this is because a composer experiences a sound field in a reverberant space. The field may be produced by individual instruments of great complexity, with complicated radiation patterns and different and somewhat varying locations. The simulation of such complicated sound fields has been attempted through the use of two or more speakers and artificial reverberation. This is an important and complex subject which will be secondary in the following discussion. Nonetheless, we note that most music is heard and judged, not by listening through headphones, but by listening in a reverberant space.

With the advent of the electrical generation of waveforms and their presentation through headphones, psychoacoustic research, aside from that which dealt with architectural acoustics, came to deal largely with isolated stimuli presented to the subject over headphones, and often in a random order. This occurred because the primary aim of most auditory research was to comprehend and model the basic operation of the auditory system. In many cases, the subject was not asked what he heard; rather, he was asked to adjust for some sort of match or perception by turning a knob or some other control. This is far from listening in a room for features of a musical sound in a musical or near-musical context.

With the advent of computer-generated sounds around 1960 [Mathews, 1969], and the subsequent developments in digital synthesis, composers have at their disposal means, literally, of generating any possible sound wave.

Here we raise a question, and propose to seek an answer. How much of the immense literature on auditory phenomena is salient in a musical contest? How much is meaningfully applicable to the task of the composer or instrument maker? Are parts of the literature misleading?

Pierce has made unsettling observations concerning the perception of residue frequencies in a musical context [Pierce, 1986].

One well known experiment is to match the pitches of tones with two sinusoidal components with a constant and equal separation of the frequency components. Ordinarily, the pitch of the lowest tone is varied and the residue pitch is determined by matching to an adjustable sinusoidal signal. Commonly, a residue pitch match is found at or near the difference frequency. Usually, there are multiple matches at various harmonics of the difference frequency [Smoorenburg, 1970; Schouten, 1962].

Pierce found that in playing an ascending octave scale with the lowest component of tones made up of two or three sinusoids with equal differences in frequency, what is heard is an ascending scale, rather than a sequence of residue pitches, all of the same or nearly the same frequency which is at or near to a least common divisor of the frequency components, as traditionally observed. Further, when immediately following the playing of the ascending scale, the first note of the scale was played followed by the last note, a rise of pitch was heard, despite the fact that the frequency ratios of the first and last notes had the same common divisor and presumably should give rise to the same residue pitch. Here we have a competition between a weak residue phenomenon reported from laboratory experimentation and a much stronger musical percept.

Residue pitch is not in all cases a weak phenomenon. It plays an important part in musical perception. The lowest note on the piano keyboard has a fundamental frequency of 27.5 Hz, and the perception of its pitch must depend more on the combination of

the harmonics than on the fundamental. However, with a high first partial and succeeding partials separated by at least harmonic intervals, musical pitch will be judged as that of the partial of lowest frequency, as in the case of the glockenspiel. Further, in experimenting with trains of all-positive pulses, which have both odd and even harmonics, and sequences of alternating positive and negative pulses, which have only odd harmonics, at pulse rates around 60 pulses per second a pitch is matched on pulse rate rather than fundamental frequency. This phenomenon was noted [Flanagan and Guttman, 1960] and interpreted as a match on pulse rates, but it is also a match on the spacing of harmonics. We have reason to believe that the match holds when the peak factor is reduced by well-known means [Schroeder, 1970], so there are no "pulses", and that it also holds when only six partials are present, in any relative phase. What can we say about residue pitch that will be true and useful to musicians?

Helmholtz attributed roughness or dissonance to the beating of partials. Plomp [Plomp and Levelt, 1965; Plomp, 1976] carried this further, and gave quantitative curves for roughness, or <u>tonal dissonance</u>, as a function of the frequency ratio of two sinusoidal components. This was proposed as a way of explaining the relative consonance or dissonance of chords.

Two complex tones may have different onset times, as in suspension or anticipation, and this may conceal or suppress the listener's sense of roughness [Wright and Bregman, in press]. In a single complex tone, partials rise simultaneously or nearly simultaneously, and interaction among partials lying within a critical bandwidth can make the tone sound very rough, or buzzy, or "electronic". Through additive synthesis, it can be shown that excluding any partials nearer to one another than a quarter of an octave avoids such buzziness. We propose to carry out further investigation into this very salient phenomenon and its musical uses.

Pierce [Pierce, 1986] also played two sinusoids of equal intensity together, with frequencies around 2,000 Hz and with difference frequencies of a few hundred Hz. When the difference frequency was used to "play a tune", at low levels no tune was heard; all that was heard was a roughness of timbre. This despite the fact that the literature on residue pitch indicates the perception of a residue pitch [Smoorenburg, 1970]. However, when the level was raised, a tune was heard, inside the ear rather than coming from the speakers. This was clearly an f2-f1 combination tone. The composer Johannes Goebel has asserted [private communication] that he uses f2-f1 combination tones deliberately in his compositions. These are produced by special metalophones.

Difference tones, or Tartini tones, have long been known and heard in music. Their musical salience has been questioned [Plomp's chapter, in Deutsch, 1960]. In the literature on auditory phenomena, the combination tone 2f1-f2 is heard at much

lower levels of primaries than the f2-f1 tone, yet it may be less important in music through being closer in frequency to the component sinusoids and thus less salient. The role of combination tones in music merits the further investigation that we propose.

The material above covers the background of the experiments that we propose to do initially. We wish, however to call to attention other areas in which musical saliency is both important and incompletely investigated.

In listening with headphones, experimenters have found a considerable sensitivity to the phases of the partials of a tone. Thus, time-reversed click doublets have been distinguished when the separation of the clicks is less than a millisecond [Resnick and Feth, 1975; Schroeder and Mehrgardt, 1982]. Reverberation tends to make unnoticeable the effects of constant differences in the phase between harmonics.

Changing phase differences can be salient even in reverberant listening conditions. In his study of ensemble effects with violins, Dolson found that fluctuations in the phases of partials, fluctuations caused by beating, were a more powerful clue than the accompanying fluctuations in amplitude [Dolson, 1983].

In his FM synthesis of singing sounds, Chowning has shown [Chowning, 1980] that a common vibrato is necessary to fuse various complexes of partials together into a voicelike sound. McAdams has shown the power of a vibrato on individual partials in making them stand out separately [McAdams, 1984].

Phase in another sense, in the sense of relative time of onset, is a separate subject, and it turns out that time of onset is critically related to another important matter, the masking of high frequency musical sounds by sounds of lower frequency.

The phenomenon of the masking of high frequencies by low frequencies was first reported as a musical observation concerning orchestral performance [Mayer, 1876]. Subsequent work on masking has been concerned either with telephone transmission [Fletcher, 1953] or with various studies of the mechanisms of auditory perception. For sinusoids, tones of lower pitch can match tones of higher pitch, while tones of higher pitch cannot mask tones of lower pitch. For complex tones of equal intensity, tones of higher pitch may be much more salient (louder) than tones of lower pitch. How can the composer best be guided in taking masking into account?

Some interesting work has shown large effects in the relative time of onset of two nearly simultaneous notes [Rasch, 1978; Vos and Rasch, 1981; Gordon, 1984]. Is there a correspondingly important musical effect, and if there is a salient effect, does it lie in musical saliency, in the impression of simultaneity, or in expressiveness?

In all, in investigating the phenomena of hearing and the mechanisms of auditory perception, a huge amount of data on auditory phenomena has been gathered since the time of Helmholtz. Some of this has some relevance to the perception of music; some may have little or no relevance. It appears that in musical perception there is a competition between various aspects of auditory perception. It is important to know what phenomena "win out" under various circumstances.

As noted at the beginning of this proposal, the way to find out is to gather data involving the presentation to musically trained subjects of computer-generated sounds in a musical context, presented largely through speakers in reverberant rooms. Sounds must be heard in a musical context, or at least in a context not alien to music. Sound examples must be recorded digitally, so that in case of dispute results can be demonstrated.

CCRMA (the Center for Computer Research in Music and Acoustics) at Stanford is an ideal setting for such research. Excellent computer, sound processing and recording resources are available for generating and listening to musical sounds digitally, with up to four independent channels, and for making and playing digital recordings. There are several suitably reverberant rooms or studios in which sounds can be presented to subjects. The subjects can record their responses on prepared sheets.

We also plan to have at our disposal a personal-computer based "Acoustical Work Station" with A-to-D and D-to-A conversion and a hard disk on which a considerable body of test material can be stored and easily summoned up by the experimenter or the subject. In this case the computer can record the judgments of the subjects.

Subjects with a wide range of musical experience are available at CCRMA, including students, staff and visitors.

The first task to be undertaken will be the extension of work already undertaken [Pierce, 1986] on residue pitch and difference tones. This will involve a more detailed comparison between perceptions reported in the literature [DeBoer, 1976] and pitches perceived in musical contexts of tunes, scales, chords and impressions of musical pitch.

2.2 Proposed Experiments

It is proposed that all experiments be first explored in a preliminary way by listening to digitally generated sounds presented over one or more speakers in a moderately reverberant room. The project staff will produce sequences of sounds that are short enough not to be tiring or confusing but are long enough to give insight. They will alter and adjust such presentations until they appear to elicit a clear outcome.

The sounds will be in a musical context. This may involve presenting tones in such a way as to convey a melody which the subject is asked to recognize. The musical context may be a scale, ascending or descending, diatonic, chromatic or whole tone. In the case of a scale, the question may be, does the scale go up correctly in pitch. Or, it may be pairs of tones, with the subject asked which has the higher pitch.

Only after a salient effect has been investigated, will matching, such as pitch matching, be resorted to in order to investigate the exact parameters for a match.

2.21 Residue Pitch Experiments

As noted above, residue pitches reported in the literature are sometimes not heard as musical pitches.

The chief stimuli used in the residue pitch experiments will initially be complexes of six harmonics of equal amplitude. In the case of odd and even harmonics, these will be the first six harmonics, with frequencies f, 2f, 3f, 4f, 5f, 6f. In the case of tones with odd harmonics only, these will be f, 3f, 5f, 7f, 9f, 11f. In investigating residue pitches, in addition to tones with all 6 partials, tones will be used in which the lowest 1, 2, 3, or 4 partials are absent, so that if the frequency of the fundamental is heard it will be a residue pitch.

Even and odd harmonic tones with many more harmonics will be used, also, in generating sequences of pulses that are all positive, and in generating sequences of pulses that alternate in sign and have only odd harmonics. Also, signals with the same power spectrum as such pulses will be used, but with the peak factor drastically reduced by the method of Schroeder [Schroeder, 1970].

In these residue-pitch experiments, one experiment will be to present tones with missing lower harmonics in a diatonic ascending scale of the lowest harmonic present, to determine under what conditions a scale is heard, or whether, as indicated in some literature, a sequence of residue pitches near the pitch of the missing fundamental is heard. It may, for example, be a function of the number of partials in the complex or of their distance from the fundamental.

Another related experiment will be to present pairs of tones in different frequency ranges and with different numbers of partials as a pair of succeeding tones, and see if the pitch judgments follow the residue pitch as to high or low, or whether they more nearly follow the pitch of the lowest partial present.

In another residue pitch experiment we will use in a similar way tones of different numbers of harmonics, tones with even-and-odd harmonics vs tones with odd harmonics only, alternately in descending whole-tone scales. This will give the reaction of a

variety of subjects to whether or not two different tones which might be expected to have the same residue pitch do have musical pitches equal to their residue pitches.

Other experiments will involve presenting tones with a fundamental frequency (missing or present), but with even-and-odd partials vs odd partials alone in some of the sequences described above. For low fundamentals, we expect the pitches heard to correspond to the harmonic spacing rather than to the frequency of the fundamental, missing or present.

Only when we have found clear correspondenses of frequency will we resort to pitch matches in which the subject adjusts the pitch of one tone. Such matches will be done on the Acoustic Work Station.

2.22 Buzziness of Sounds

As we have noted above, in tones with partials higher than the 6th harmonic, higher partials may be so close together that the tones are self-dissonant, or buzzy (harsh or "electronic"). The effect is greatest for tones of low pitch (say, below A3), in which higher harmonics are salient. It is not easy to generate sounds with computers that are bright (have high partials) without being buzzy.

This buzziness can best be demonstrated by contrasting tones in which harmonics above the 6th have been deleted when they are less than a critical bandwidth (about a minor third) above the preceding partial, with tones which have pairs of succeeding partials above the 6th. Such comparisons can involve alternations of buzzy and non-buzzy tones or, a scale passage, sequential chords or arpeggios played alternatively with buzzy and non-buzzy tones.

It will be important to present contrasts of tones in different frequency ranges, randomly starting with the higher or the lower pitch, and determine whether a difference in harshness is indeed judged as greater for tones of low frequency.

It is also important to know how far down in level "missing" partials must be in order not to contribute annoying buzziness. To determine this, pairs of tones will be presented; both will be at the same pitch, but one will have the "missing" harmonics added at a lower level, ranging from -30 dB to higher levels, in steps that are found to be reasonable. In this experiment it is important that the buzziness of the tone with added "missing partials" always be compared with the non-buzzy tone. If we compare successive tones with higher and higher levels of missing partials, undetectable differences between successive tones can add up to a striking difference between the first and last tones when presented in direct succession.

In this experiment it is important to gather such comments as "annoying difference" from many subjects.

In the above description of experiments on buzziness, additive synthesis of steady sounds which are not musically useful will be used. It is important as experimentation is continued to seek more musically useful sounds which are bright without being buzzy.

2.23 Combination Tones

We have used pairs of sinusoids in the 2,000-Hz region to play a tune, using the difference frequency to produce the pitch of the tone as a difference-frequency pitch [Pierce, 1986]. The tune is not heard at low levels, and at higher levels it is heard, not as coming from the speakers, but in one ear or the other.

These preliminary experiments must be repeated and expanded. Particularly, under conditions in which the tone is heard, its musical pitch must be verified. It seem best to do this directly, by a musical subject adjusting the pitch of a musical tone (rather than a sine wave, whose pitch is less easy to judge and may depend on level).

The experiment must be carried out over various frequency ranges of the sinusoids, and with various difference frequencies, to see over what ranges the difference tone is heard. This must be approached by generating a large number of examples and playing them to a considerable number of musical subjects.

We have noted above that we believe the f2-f1 difference tone to be more musically salient than the 2f1-f2 combination tone, which is heard at lower levels. The trouble with the 2f1-f2 tone being heard is that in the conventional periodic complex tone it will be at the frequency of a partial that belongs in the harmonic series and therefore not generally expected to create a separate pitch. How much difference the presence of the nf1-(n-1)f2 tones make in timbre is a nearly moot point, since presumably they are routinely present in tones of moderate sensation level. However, since their frequencies coincide with other (frequently present) harmonic partials, they may contribute to timbre changes by their different vector interaction as different partials have their relative phase relations changed.

2.24 Recognition of Constant Spectrum

In addition to the three experiments outlined above, we propose at some point to carry out an experiment on recognition of constant spectrum.

Few natural sound sources maintain a completely constant relation among the partials over an appreciable (fundamental) frequency range. Such a possibility exists -- in fact, is quite convenient

-- for computer-generated sounds. Thus the question of the response to this type of spectral invariance becomes more interesting than for natural sources.

Schubert has shown (unpublished experiment) that for unfamiliar (synthesized) tones listeners cannot in general discern small changes in spectral envelope if the fundamental frequency between the two tones to be compared has changed by as much as a minor third. (How this relates to Green's comprehensive work on detection of small changes in spectral profile for constant frequency composition is also an interesting question.) We propose to investigate the recognition of this type of spectral constancy by changing the spectral envelope in controlled amounts in note sequences as contrasted with constant spectral composition for the same sequences. It is our feeling that the composer can profit from knowing how the listener responds to planned changes in spectrum and how these relate to the frequency changes characteristic of melodic patterns.

The basic experiment will be simple interval presentation with the only question being whether the timbre changed or remained constant for the two notes -- half the presentations involving no change of spectrum when the fundamental changes. Extensions of this will move to sequential chords or abbreviated scale passages, again ascertaining how great a change in spectral composition is required to detect the difference.

Representative timbres to be used is not an easy question to answer. Already envisioned are (1) a sloping spectrum with rate of slope being variable, (2) a spectrum with even partials weaker than odd and the odd-even balance gradually changed, (3) a formant-like spectral shape with the slopes of the formants the variable to be controlled, and (4) a dual-polarity pulse train with the spacing of the reversed pulse varied to modify the spectrum. Further controllable signals will be sought from the signal repertoires of resident composers.

2.25 Other Experiments

As the experimental work described above is brought to a satisfactory state, other phenomena will be investigated, including those mentioned above: the actual importance and utility of effects of phase, any salient effects of time of onset, especially in relation to masking, and other effects. It is our belief that new approaches and understanding, and perhaps new phenomena, will come to light in the course of the work proposed.

3 BIBLIOGRAPHY

Bregman, A. S., Asking the "what for" question in auditory perception. In M. Kubovy and J.R. Pomerantz, Perceptual Organization. Hillsdale, New Jersey: Lawrence Erlbaum, 1981.

Bregman, A. S., Auditory scene analysis. Proceedings of the Seventh International Conference on Pattern Recognition. Silver Spring, Md.: IEEE Computer Society Press. (Library of Congress No. 84-80909), 1984

Chowning, John M., Computer Synthesis of the Singing Voice, in Sound Generation in Winds, Strings and Computer. Johan Sundberg, ed., Royal Swedish Academy of Music, 1980.

DeBoer, E., The "Residue" and Auditory Pitch Perception, in Handbook of Sensory Physiology, Vol. V. Ed. by Wolf D. Keidel and William D. Neff, Springer Verlag, 1976.

Deutsch, D., Musical Illusions. Scientific American, 233, 92-104, 1982.

Deutsch, D. (Ed). The Psychology of Music. New York: Academic Press, 1982. See particularly chapters 4 and 9, which are by Deutsch herself.

Deutsch, D., The octave illusion in relation to handedness and familial handedness background, Neurophysologica, 21, 289-293, 1983.

Deutsch, D. A musical paradox. Music Perception, 3, 275-280, 1986.

Dolson, Mark. A Tracking Phase Vocoder and its Use in the Analysis of Ensemble Sound, Ph.D. Thesis, California Institute of Technology, 1983.

Flanagan, J. L, and Guttman, N. On the pitch of periodic pulses, J. Acoust. Soc. Am., 32, 1319-1328, 1960.

Flanagan, J. L. and Guttman, N. Pitch of periodic pulses without fundamental component. J. Acoust Soc. Am., 32, 1319-1328, 1960.

Gordon, J. W. Perception of Attack Transients in Musical Tones. Ph.D. Thesis, Stanford University, 1984.

Green, D. M., et. al. Auditory Demonstration Tapes. Prepared Dec. 15, 1978 at the Laboratory of Psychophysics, Harvard University, under National Science Founddation Grant SED76-24215.

Helmoltz, Hermann. The Science of Musical Sound as the Basis for a Physiological Explanation of Music. Dover Edition, 1954.

Mathews, M. V. The Technology of Computer Music. MIT Press, 1969.

Mayer, A. M. Researches in Acoustics. Philos. Mag., 2, 500-507, 1876.

McAdams, Stephen. Spectral Fusion, Spectral Parsing and the Formation of Auditory Images. Ph. D. Thesis, Stanford University, Stanford, California, May, 1984.

Pierce, John R. The Science of Musical Sound, Scientific American Library, 1983.

Pierce, J. R. Residue Pitch and Beats as Heard in a Musical Context, abstract of invited paper. Jour. Acoust. Soc. Am. Supplement 1, 80, S86, 1986. This work will be published in 1988, together with sound examples on compact disk, in an MIT Press book.

Plomp, R. Aspects of Tone Sensation. London: Academic Press, 1976.

Plomp,R. & Levelt, W. J. M. Tonal consonance and critical bandwidth. Jour. Acoust. Soc. Am., 38, 548-560, 1965.

Rasch, R.A. The perception of simultaneous notes such as in polyphonic music. Acustica, 40, 21-33, 1978.

Resnick, S. B. and Feth, L. L. Discriminability of time-reversed click pairs: intensity effects. Jour. Acoust. Soc. Am., 57, 1493-1499, 1975.

Roberts, L. A. and Mathews, M. V., Intonational sensitivity for traditional and nontraditional chords. J. Acoust. Soc. Amer., 75, 952-959, 1984.

Schroeder, Manfred, Synthesis of low-peak-factor signals and binary sequences with low autocorrelation, IEEE Trans. on Information Theory, IT-16, 1970

Schroeder, M. R., and Mehrgardt, S. Auditory masking phenomena in speech. In R. Carlson and B. Grandstrom (Eds.) The Representation of Speech in the Auditory System. Elsevier, 1982.

Vos, J. and Rasch, R., The perceptual onset of musical tones. Perception and Psychophysics, 29 (4), 1981, 323-335.

Wright, J.K. and Bregman, A.S. Auditory stream segregation and the control of dissonance in polyphonic music. Contemporary Music Review, in press.