Proceedings of the Institute of Acoustics

MASKING OF CODING/DECODING ERRORS IN AUDIO DATA COMPRESSION SYSTEMS

Michael A. Gerzon

Technical Consultant, 52 Walton Crescent, Oxford OX1 2JQ

1. INTRODUCTION

Recently, a variety of different systems have been proposed for reducing the data rate of high quality digital audio to between two and four bits per sample [1]. All such systems introduce encoding/decoding errors, and rely on these errors being masked by the wanted signal in order to achieve high quality results. Despite the proposed systems substantially meeting the conventional requirements for good masking of errors, many experienced audio professionals have expressed reservations about such systems. The aim of this paper is to present evidence that conventional masking criteria are indeed inadequate, and to note that effective masking thresholds can be reduced by as much as 30 dB in some circumstances.

Specifically, we note that well-known psychoacoustic phenomena suggest that conventional models of spectral masking break down when there is a substantial degree of cross-correlation between the error-signal and the wanted signal. We specifically examine the effect of correlated errors that cause gain modulation of the wanted signal. Such errors are audible even at a level of 36 dB below the signal, although their audible effect is not that of a spurious noise, but is that of a change of the character of the wanted signal. We then go on to demonstrate that all Shannon-efficient coding systems using Max quantisers (i.e. quantisers with minimum error energy) cause significant gain modulation effects that, in currently proposed systems, are well above even a conservative audibility threshold derived from the literature.

These results are of some importance, since they suggest that existing approaches to low-bit-rate audio data compression, based on using spectral masking of errors with a high coding efficiency, are going to lead to audible alterations in the character of signals akin to those encountered with mis-tracking analog noise reduction systems. Possible means of getting round this problem, at the expense of a small loss of theoretical coding efficiency, are discussed.

2. SPECTRAL MASKING

Providing an error is substantially coincident in time with the wanted signal, the general theory of spectral masking asserts that one can predict whether or not the error is masked by the wanted signal solely from a knowledge of the power spectrum of the wanted signal and that of the error-signal. Given the power spectrum of the wanted signal, one can compute from it a <u>threshold spectrum</u> such that if the error power spectrum lies below this threshold spectrum at all frequencies, then the error will be inaudible. Although the precise procedure for computing this threshold spectrum from the wanted-signal power spectrum is somewhat uncertain, it is generally agreed that the following is a good first approximation.

One can determine experimentally the spectral masking threshold for narrow bands of noise for a masking sine wave signal at each frequency and level. It is found that if the noise band is close to the sine wave frequency (say within the critical bandwidth - around 0.2 octaves at middle frequencies), then the noise is masked by the signal if it is between 4dB and 7dB lower in level, and the masked level of noise falls away rapidly as the difference of the frequencies increased. This masking curve resembles the energy spectral response of a moderately high Q tuned filter. To determine the masking threshold spectrum for a more complex signal than a sinewave, one normally convolutes the wanted-signal's power spectrum with a convolution kernel, varying with frequency, respresenting the masking threshold for the energy within a critical bandwidth of each frequency. If any error signal is below this computed threshold, or below the absolute threshold of hearing at any frequency, it is presumed to be inaudible.

It is not our intention here to give any detailed model along the above lines for spectral masking, only to indicate the general type of model that is widely used. Our reason for doing this is to note that such a model is clearly grossly incorrect in some situations, and that the spectral masking model described in this section cannot be relied on, whatever the subtle detailed modifications may be made in the way that the threshold spectrum is computed. We claim that the spectral masking model is conceptually flawed at a fundamental level.

3. CORRELATED ERRORS

We now demonstrate the fundamental flaw in spectral masking models: namely that, for a given masking signal, two error signals having the <u>same</u> power spectrum may be masked to very different degrees. In other words, the spectrum of the error is not enough, on its own, to predict whether or not it will be masked.

The demonstration of this is remarkably simple. It is known that gain changes in a sine wave of the order of 1 dB are audible, and that at mid frequencies and levels, gain changes of as little as 0.3 dB can be heard. Consider a wanted signal that is a sinewave of frequency F :

$$\sin(2\pi Ft)$$

and an error signal of the form

$$a(t)\sin(2\pi Ft)$$

where $a(t)$ is a low-frequency waveform, having no frequencies above around 20Hz or so. Then the effect of the error signal is to modulate the gain of the sinewave by $1 + a(t)$. If the amplitude of $a(t)$ varies between ±0.016, then there is a gain modulation of about ±0.15 dB, giving a gain variation of 0.3 dB. For a sufficiently long duration of these two extreme gains, the resulting 0.3 dB gain change will be audible. However, the energy of the actual error signal here is 36 dB below the wanted signal. Thus we have shown that, in some circumstances, an error signal within the critical bandwidth of the wanted signal can be heard at a level 36 dB below the wanted signal. This is a level about 30 dB below the conventional masking threshold.

There is nothing very controversial about this observation, since one's common sense would prevent one using spectral masking theory in such a situation. However, this example shows that one has to be very careful in using spectral masking theory to ensure that one is only applying it to errors of a suitable kind - and it is evident from the above that errors due to amplitude modulation of the wanted signal are not "of a suitable kind". We shall show in the next section that, unfortunately, conventional Shannon-efficient audio data compression systems do produce substantial amplitude modulation effects.

What is it about amplitude modulation errors that render them less liable to masking than noise-like errors of a similar frequency? Although we cannot produce entirely definitive answers, it seems likely that a large part of the answer lies in the degree of cross-correlation between the error and the wanted signal. The short-term cross-correlation between two signals f(t) and g(t) at time t may be defined as the integral

$$\int w(t-t')f(t'-\tfrac{1}{2}\tau)g(t'+\tfrac{1}{2}\tau)\,dt' \;=\; C_w(\tau,t)$$

where the weighting function w is a positive function of total integral 1 which is typically nonzero for around 50 ms. As the duration of the weighting function gets longer, this approaches the ordinary cross-correlation. The Fourier transform of the short-term cross-correlation with respect to the τ variable is termed the (short-term) cross-spectrum of f(t) and g(t) at time t. The short term cross-spectrum of f(t) with itself is, of course, the conventional power spectrum of f(t) windowed by the weighting function w(t).

The real and imaginary parts of the cross-spectrum of the wanted signal with the error signal provides the information missing from a knowledge only of the spectra of the wanted and error signals. Until demonstrated otherwise, we can continue to assume that, when the cross-spectrum of signal with error is zero, then spectral masking criteria can be used safely. However, amplitude modulation errors cause the real part of the cross-spectrum to become non-zero, and phase-modulation errors of the wanted signal cause the imaginary part of the cross-spectrum to become non-zero. It is necessary to determine by experiment separate empirical masking curves for the real and imaginary parts of the cross-spectrum. From the arguments given earlier, we know that, depending on frequency, the masking threshold for the real part of the cross-spectrum can be as much as 36 dB below the level of the wanted-signals power spectrum at the same frequency.

The masking threshold for the imaginary part of the cross-spectrum is less certain. Although random phase modulation is known to be significantly less audible than the same sideband error energy in the form of random amplitude modulation , we cannot safely extrapolate this result to the case where amplitude or phase modulation is strongly signal-dependent, as empirical experience in the design of dynamic filters demonstrates. In these devices, phase-modulation effects have been found to be at least as disturbing as amplitude modulation effects of similar magnitude.

Until more experimental data is available on the audibility of cross-spectral error components is available, we cannot give precise figures about cross-spectral masking thresholds, but we can say as a matter of general experience that these thresholds are considerably lower than for uncorrelated errors - perhaps 20 dB or 30 dB lower.

The audible effect of cross-spectral error components is generally quite different to that of uncorrelated errors. The latter have the quality of being an added unwanted sound to the wanted sound. Errors in the real part of the cross-spectrum cause an "unstable" quality, sometimes described as "pumping", which is familiar in mistracking analog noise reduction systems or dynamic processors. Errors having an imaginary cross-spectrum with the signal suffer from an effect known subjectively as "phasing", where audible alterations of pitch of some signal components are heard. In general, cross-spectral error components manifest themselves as an alteration in the quality of the wanted signal, and not as an added separate sound. It is possible that critical high-quality program material, such as that recorded with simple stereo microphone techniques in natural acoustics, conveying an accurate portrayal of ambient distance cues for sound sources, may be damaged by much smaller cross-spectral errors than those discussed above (gain errors well below 0.1 dB of gain modulation might be audible in such critical cases).

In the absense of detailed experimental information, it is advisable to keep gain modulation error well below 0.3 dB, and it is possible that it might be wise to prevent them altogether if at all possible.

4. GAIN MODULATION IN EFFICIENT CODING

There is a long-established theory, known as Shannon Rate-Distortion theory [2], that allows one to determine how low a bit rate a given signal, with a known power spectrum, can be coded into if one puts an upper bound on the spectrum of the error signal. Although this theory has some difficulties (it strictly applies only to Gaussian signals, and is more difficult in the non-Gaussian case), it is a good guide to how efficient a coding system can be made, and practical coding systems, such as those using principal-value transform coding [3] or adaptive differential pulse code modulation (ADPCM) [4] can come quite close to the Shannon bounds in performance.

However, the Shannon theory, which we shall not attempt to summarise here, is based on coding signals to achieve a given error-signal power spectrum at the lowest possible bit rate, without taking any account of whether or not there is a non-zero cross-spectrum between the wanted signal and the error signal.

There seems to have been much wishful thinking on the part of many workers in audio data compression systems in that they seem to have assumed, or hoped, that either cross-correlation of the error is unimportant, or else that the error in a Shannon-efficient coding system has the form of an amplitude-modulated (but not necessarily Gaussian) random noise signal. We now show that not only is the first assumption false (as shown in section 3), but that the second is untrue also. We are unaware of any result in the existing literature about whether or not the error signal in Shannon-efficient coding is cross-correlated with the wanted signal. If it is known, it is certainly not generally known to workers in the field.

We shall now outline a proof that the error signal in a Shannon-efficient coding system, and in systems practically approximating efficient coding, is cross-correlated with the wanted signal, and also that the magnitude of this cross-correlation is large enough to be of concern in proposed systems.

Our argument is a simple geometric one, which to most readers will be a handwaving, but hopefully plausible, argument. However, to those with a knowledge of Hilbert space techniques, we note that the argument is actually a mathematically rigorous one in Hilbert spaces of signals.

We can represent the wanted signal by a vector or arrow S whose length from the origin is equal to the square root of the signal energy. See figure 1. Denote the length (i.e. square-root of energy) of S by $\|S\|$, and similarly for other signals. An encoding/decoding or quantising system will produce a modified signal which we shall write as QS, and this can be represented as a second vector pointing in a different direction to S, since it is a different signal. The error signal

$$\varepsilon S = S - QS$$

can be represented as a third vector along the third side of the triangle formed by the vectors S and QS as in figure 1.

Now a Shannon-efficient coding system, or the so-called "optimum" or "Max" quantisers [5] used to quantise the signal components in practical proposed coding systems, aim to ensure that the error-signal energy (suitably weighted) is minimised, i.e. that the length of the error vector εS is minimised. For a given direction of the quantised or coded signal vector QS, this means that the error vector εS is at right angles to the coded signal vector QS. (See figure 2). By Pythagoras' theorem, we have that

$$\|S\|^2 = \|QS\|^2 + \|\varepsilon S\|^2 ,$$

and also the cross-correlation of the error signal with the coded signal QS is zero, since the two vectors are at right angles to each other. However, since the error signal εS is <u>not</u> at right angles with the wanted signal S, there <u>is</u> a cross-correlation between the error and wanted signals. This can be seen from figure 2, where the component of the coded signal QS correlated with the wanted signal S is the orthogonal projection of QS onto S, which by simple geometry equals

$$(\|QS\|/\|S\|)^2 \; S ,$$

which has length $\|QS\|^2/\|S\|$. Thus the component of QS that is correlated with the wanted signal S has amplitude gain

$$(\|QS\|/\|S\|)^2 = 1 - \|\varepsilon S\|^2/\|S\|^2,$$

which varies according to the relative energy of the error signal. This proves that the correlated component of the coded/decoded signal is subject to amplitude modulation, and also that the component of the error signal εS correlated with the wanted signal has amplitude gain

$$\|\varepsilon S\|^2/\|S\|^2 .$$

Now Shannon Rate-Distortion theory [2] suggests that for a coding system using B bits per sample, then for efficient coding,

$$\|\varepsilon S\|^2/\|S\|^2 = 4^{-B} .$$

For practical coding systems that quantise signal components using Max quantisers (see [5]), the degree of amplitude modulation is somewhat greater. The data of J. Max [5] show that, for Gaussian signal statistics, a Max quantiser produces a reduction of signal gain for the correlated component of the coded signal of approximately 4^{-B+2} dB, i.e. about 4 dB gain reduction for a 1-bit Max quantiser, 1 dB gain reduction for a 2-bit Max quantiser, 0.25 dB gain reduction for a 3-bit Max quantiser, and about 0.06 dB gain reduction for a 4-bit Max quantiser.

Practical efficient coding systems quantise the different transform signal components using a variable number of bits (this is termed "dynamic bit allocation"), so that the amplitude gain of these components is not fixed, but varies dynamically in a signal-dependent way from moment to moment. Even systems using an average of 4 bits per sample will allocate as little as 1 or 2 bits to some audible signal components, so that we can conclude that currently used coding strategies produce a signal-dependent amplitude modulation of audible signal components of about 4 dB. This is over ten times the 0.3 dB known to be audible from standard results on audible gain changes.

5. REDUCING MODULATION EFFECTS

Given that we have a problem with amplitude modulation of the wanted signal, the question arises of what we can do about this effect. The first comment to make is that we need good psychoacoustic data on the audibility of signal--dependent gain modulation, not just on test tones, but also on high quality signals containing complex natural acoustic cues, since there is reason to believe that subtle variations in signal-envelope information may be important in the ears' interpretation of complex cues.

One can certainly say that reducing amplitude modulation effects means departing from very high Shannon coding efficiency - since we have shown that the amplitude modulation is a consequence of minimising error energy. Two main strategies exist for reducing amplitude modulation. The first strategy is simply to modify the reproduced gain of quantised signal components to avoid amplitude modulation. The second strategy is to use a dithered quantiser.

Modifying the gain of quantised signal components so that the correlated coded signal components avoid amplitude modulation means increasing the gain of the Max-quantised signal QS by a gain $\|S\|^2/\|QS\|^2$, which also has the effect of increasing the error energy by the same factor, and of increasing the noise-like component of the error energy uncorrelated with the wanted signal by the square of this amount. In practice, this means that an increased number of bits is needed to ensure effective masking of the low-bit transform components, requiring a slightly altered bit-allocation strategy.

However, there is a problem with this proposal - namely that it assumes that one actually knows what the optimal Max quantiser Q is. This depends on knowing reliably the instantaneous signal statistics. In practice, there is considerable uncertainty as the exact momentary signal statistics, and this uncertainty causes considerable unpredicatable amplitude modulation of the correlated component of the quantised signal. Although this requires further theoretical study, our provisional conclusion is that the practical amplitude modulation effects for components at low bit rates will remain serious, although a careful choice of quantiser characteristic can minimise the gain effect of quantiser/signal statistics mismatch - at the expense of a further loss of Shannon efficiency.

The second option is to use a subtractively dithered quantiser, as originally described by Roberts [6]. This involves adding a pseudo-random dither noise signal before the quantiser, and subtracting the same dither signal in the decoding process. Such subtractive dither ensures that the error signal is noiselike, and guarantees no amplitude modulation of the correlated component of the coded/decoded signal. There are two snags here. First, subtractive dither decreases the signal-to-noise ratio of an n-level quantiser to that of an (n-1)-level quantiser. Secondly, subtractive dither is only applicable to uniform or so-called linear quantisers with equal step sizes, which complicates the task of optimising the quantiser coding performance to match the signal statistics. In particular, it is important to limit the amplitude of signals to be coded so that they do not exceed the peak level of the quantiser - otherwise one gets clipping distortion (which also modifies amplitude gain).

We hope at a future time to publish detailed results on methods reducing amplitude modulation effects in low-bit-rate audio coding systems using masking, but space precludes dealing with the required theory here.

6. STEREO DIRECTIONAL MASKING

Another effect invalidating the use of conventional masking theory must be mentioned. For stereo signals, error signals lying in a percieved stereo direction different to that of the masking signal will be more audible than those in the same direction. To maximise stereo directional masking, coding systems should be designed to ensure that error signals have substantially the same stereo distribution as the masking wanted signal. This means either using a principal-component quantisation method for the full stereo signal, or a stereo matrix prediction filter. Existing systems simply quantise the stereo signal as separate mono channels, and so do not maximise directional masking.

7. CONCLUSIONS

We have shown that low-bit-rate audio coding systems produce amplitude modulation effects that are not masked according to established models for error masking, and also noted that stereo systems might also directionally unmask errors. These problems are likely to be serious for existing proposed systems, producing significant subjective "pumping" effects. Methods of reducing these problems, at the expense of a loss of Shannon-coding efficiency, have been discussed, but these require modified encoding/decoding strategies and algorithms.

8. REFERENCES

[1] M GERZON, 'The Gentle Art of Digital Squashing', Studio Sound 32 no. 5, pp.68-76 (1990 May)
[2] L D DAVISSON, 'Rate-Distortion Theory and Application', Proc. IEEE 60 pp.800-808 (1972 July)
[3] J J Y HUANG & P M SCHULTHEISS, 'Block Quantization of Correlated Gaussian Random Variables', IEEE Trans. Commun Syst. CS-11 pp.289-296 (1963 Sep)
[4] J MAKHOUL, 'Linear Prediction: A Tutorial Review', Proc. IEEE 63 pp.561-580, (1975 Apr.)
[5] J MAX, 'Quantizing for Minimum Distortion', IRE Trans. Inform. Theory, IT-6 pp.7-12 (1960 Mar.)
[6] L G ROBERTS, 'Picture Coding Using Pseudo-Random Noise', IRE Trans. Inform. Theory, IT-8 pp.145-154 (1962 Feb.)

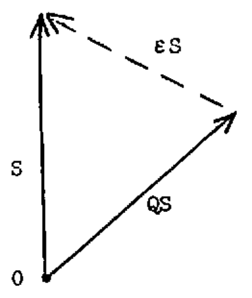

Figure 1. Vectors representing a signal S, a coded signal QS and the error signal εS.

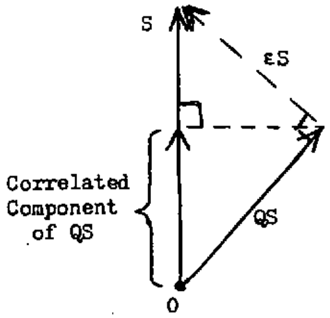

Figure 2. Vectors representing a signal S and coded signal QS for least rms error, showing projection of QS onto S.