The Vorbis Encoder's Scalar Quantizer and the Energy Preservation Problem

Sebastian Gesemann

March 2, 2004

Abstract

This paper describes scalar quantization with a focus on the implementation of the current Vorbis encoder (version 1.0.1) and explains the reason of the well-known "high frequency boost" problem. Further, a solution is provided which is based on a smart selection of quantization thresholds for specific sample distributions. Finally the solution's impact on the signal-to-noise ratio is discussed.

1 Introduction

Many of the Vorbis fans (including me) complained about the fact that the current Vorbis encoder seems to amplify higher frequencies. This is especially a problem in double blind tests because it enables users to distinguish between the original and the encoded sound file.

Due to the long time existence of this problem my motivation was high enough to investigate on my own. So, here I am – trying to express my thoughts on this issue.

Currently the Vorbis encoder always quantizes the frequency samples of the residue vectors linearly by rounding to the nearest integers. The psychoacoustic model determines a signal-to-noise ratio for each of the signal's frequency bands and controls the floor curve which is used to create the residue vectors (a frequency adaptivly scaled version of the spectral vectors). Usually the model tries to save bits in the upper parts of the spectrum and uses low SNRs for those regions.

I prepared a 120 second monophonic white noise audio file and encoded it using the current Vorbis reference encoder (quality level 4). The analysis (figure 1) shows the mean energy (Y-axis) versus frequency (X-axis) of 3 signals (original, encoded and error signal). We can clearly see that the lower the SNR gets the higher the increase of energy will be. In the following section details of the quantization process are presented.

2 Scalar Quantization

In scalar quantization we map signal samples (later on referred to as x_t) to a smaller finite set of indexed scalars and code the index only. Let C be the set of indexed scalars and Q the quantizer mapping:

$$C = \{c_i \in \mathbb{R} \mid i \in \mathbb{N}, 0 \leq i < n\}$$

$$Q : \mathbb{R} \mapsto C$$

In general a quantizer tries to find a c_i for each x_t such that

$$\sum_t |x_t - Q(x_t)|^2$$

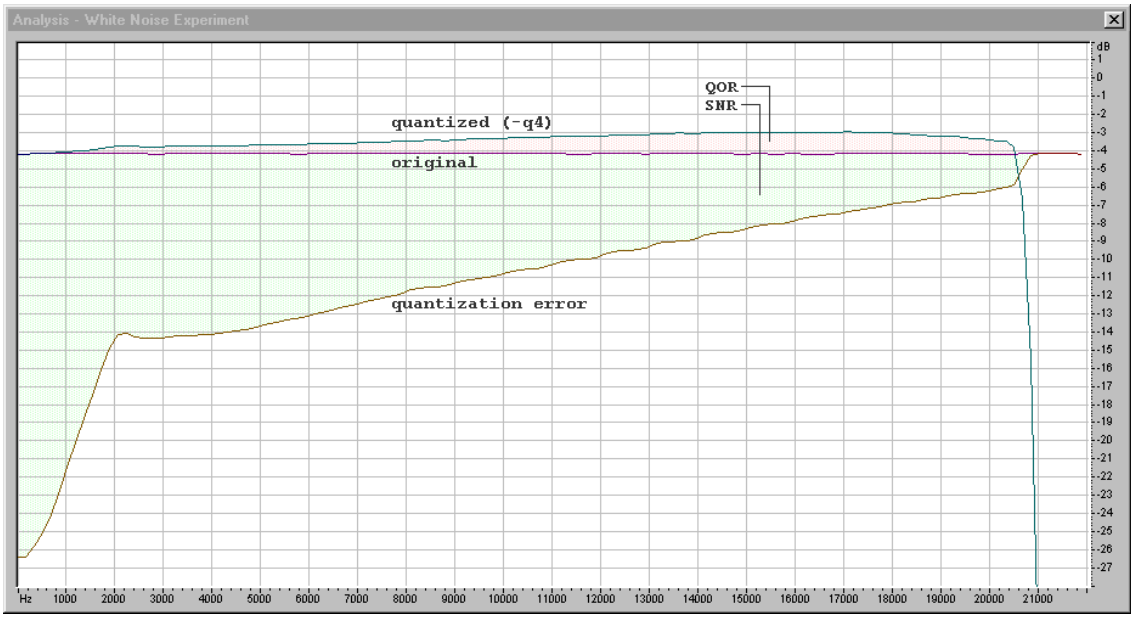

Figure 1: The white noise experiment

is minimized (quantization error energy). The ratio

$$\frac{\sum_t x_t^2}{\sum_t |x_t - Q(x_t)|^2}$$

is usually called signal-to-noise ratio (SNR). It is easy to show that if we define Q as

$$Q(x) = c_i \quad | \quad |x - c_i| \leq |x - c_j| \ \forall j$$

the quantization error will be minimized (and the SNR will be maximized). This is actually how the Vorbis encoder quantizes the samples, too. But how does this "HF boost" issue fit in here? Let's analyze the following ratio which I refer to as QOR (quantized-to-original ratio):

$$\frac{\sum_t Q(x_t)^2}{\sum_t x_t^2}$$

This is the ratio of the energy of the quantized signal to the original signal's energy. We can easily simulate the quantizer by replacing the samples x_t with a PDF function representing a specific distribution:

$$p : \mathbb{R} \mapsto [0, \infty)$$

$$\int_{-\infty}^{\infty} p(x)\ dx = 1$$

$$Prob(a \leq x_t \leq b) = \int_a^b p(x)\ dx$$

Let's also call

$$R : C \mapsto \mathcal{P}(\mathbb{R})$$

$$R(c) = \{x \in \mathbb{R} \mid Q(x) = c\}$$

the quantization regions. Now, we can compute the SNR and QOR using these formulas:

$$SNR = \frac{\int_{-\infty}^{\infty} p(x) x^2\ dx}{\sum_i \int_{R(c_i)} p(x) |x - c_i|^2\ dx}$$

$$QOR = \frac{\sum_i \int_{R(c_i)} p(x) c_i^2 \; dx}{\int_{-\infty}^{\infty} p(x) x^2 \; dx}$$

I've done a simulation of the Vorbis encoder's scalar quantizer for two different sample distributions (gaussian and triangular). To obtain some SNR/QOR pairs for each distribution I scaled the PDF function at several levels (see figure 2) (The gaussian distribution is a quite realistic approximation for the usually noisy upper spectrum part).

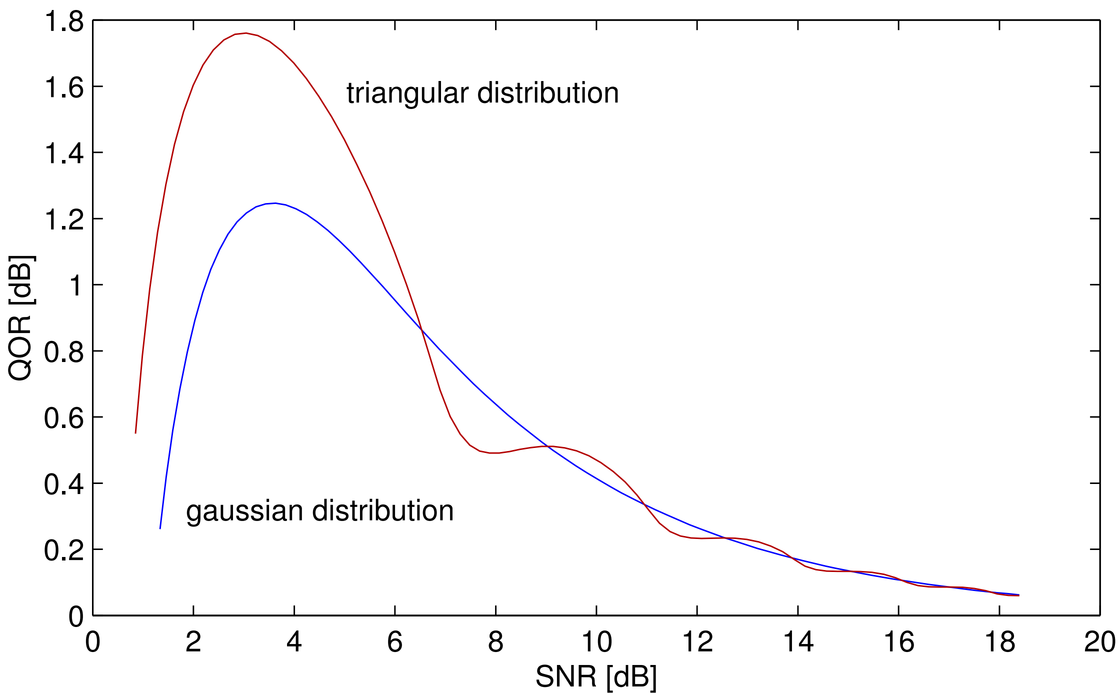

Figure 2: SNR/QOR relationship for a linear scalar quantizer

Please compare the theoretically derived figure 2 to the real-life analysis in figure 1. They really seem to match. For example: At 12000 Hz the Vorbis encoder quantized the samples with a SNR of roughly 5.5 dB and got a QOR of 1 dB – just like it is the case for a gaussian distribution (figure 2, blue line).

Further, this means this effect is not a bug in the Vorbis encoder it is rather a side effect of a linear scalar quantizer that minimizes the energy of the quantization error in low-SNR situations. So, the question is: How do we prevent this (natural) boost effect ? ...

3 Quantization Thresholds

If we equally shift all quantization thresholds by d away from zero we decrease the QOR while keeping the SNR as large as possible (see figure 3 for an example).

For each sample distribution (type & variance) a d can be determined such that the QOR equals 0 dB. In figure 4 the optimal shift d (Y-axis) is shown for the source's mean deviation sigma (X-axis) in case of a gaussian sample distribution.

Of course, the modification of quantization thresholds affects the signal-to-noise ratio. But my personal experiments have shown that those slight changes result in a rather negligible impact (see figure 5 for more details). In the worst case we'll loose 0.3 dB of the SNR which is tolerable.

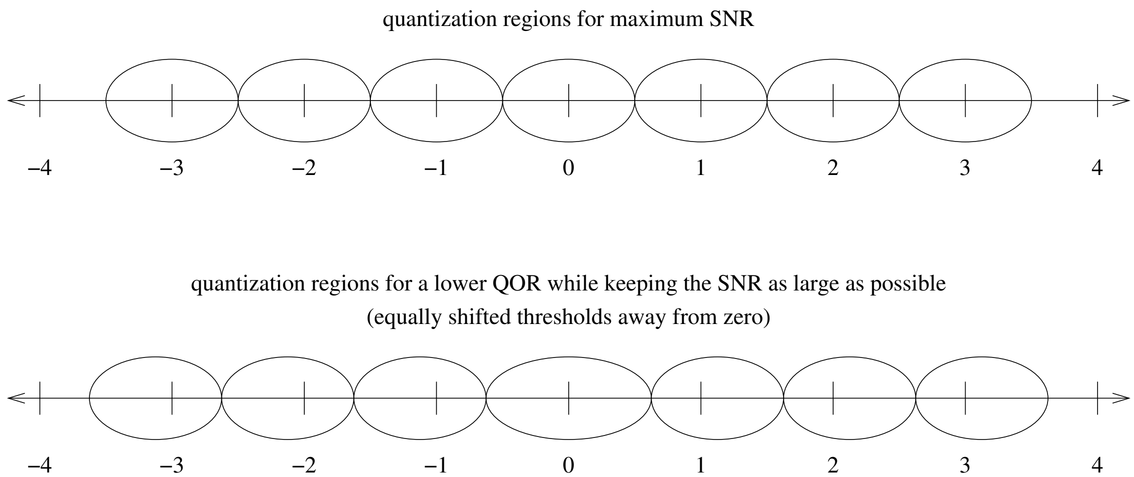

Figure 3: Shifting the quantization thresholds to compensate for a positive QOR

4 Updating the Vorbis encoder's implementation

The residue samples are coded in partitions with an appropriate set of code books that match the partitions' sample distribution. Therefore classification codes are stored in the bit stream. One could easily assign an appropriate threshold shift d to each of those classes and use it for quantization. This could be done offline before compile time.

Quantization with the modified thresholds can be reduced to ordinary quantization by adjusting the residue samples (simply shifting the values towards zero by d).

5 Conclusion

In this paper I've shown the "why" of the high-frequency-boost effect and how to get rid of it. The implementation should not be too complicated. For further details on this topic I can be contacted at <sgeseman *at* upb *dot* de>.

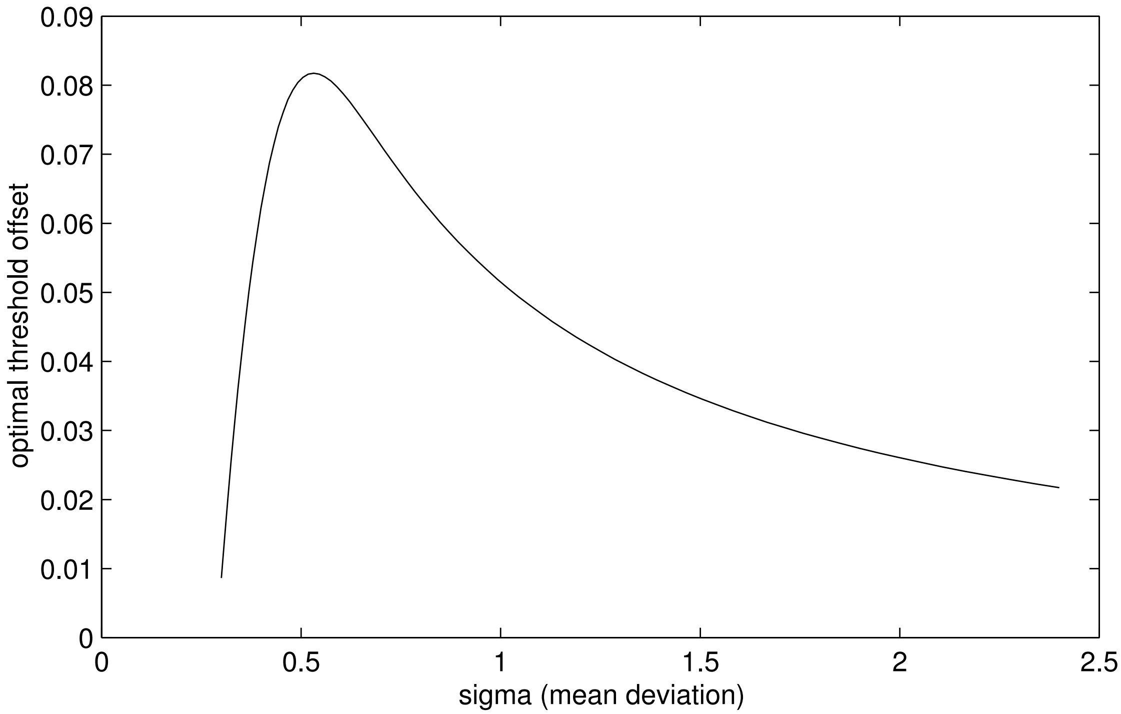

Figure 4: Optimal (QOR=0 dB) threshold shift d for a gaussian source

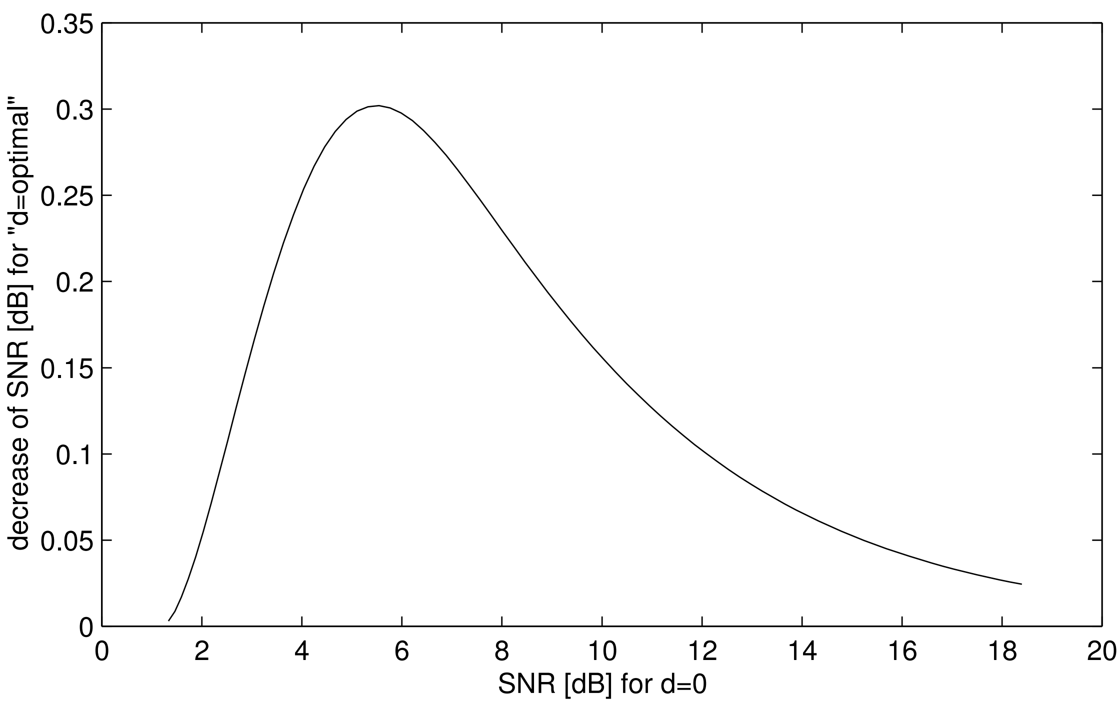

Figure 5: The solutions impact on the SNR