ON THE ROLE OF WIENER FILTERING IN QUANTIZATION AND DPCM

Geir E. Øien and Tor A. Ramstad

Norwegian University of Science and Technology
Department of Telecommunications
N-7491 Trondheim, NORWAY
{oien,tor}@tele.ntnu.no

ABSTRACT

The additive noise model of a scalar quantizer facilitates the use of a Wiener-type receiver/reconstruction filter performing signal estimation and quantization noise attenuation. We contrast the two important cases of fixed-rate pdf-optimized quantization, and uniform quantization followed by entropy coding. Assuming a stationary input signal with known second-order statistics, it is shown that the Wiener filter response is unity in the pdf-optimized case, but is a nontrivial filter in the uniform quantization case. The theory is then extended to the DPCM case, where we consider a closed-loop structure with both types of quantizers. We derive Wiener filter responses and theoretical coding gain expressions covering both cases. Simulation results confirm that a coding gain is obtained for the uniform quantization case by introducing a Wiener synthesis filter instead of the ordinary inverse filter used in DPCM, most significantly at low rates and for highly correlated input signals.

1. INTRODUCTION

A scalar quantizer (SQ) is a mapping from the real line to a finite set of representation values according to a nearest-neighbor rule [1]. Let us denote the input signal to an arbitrary SQ at time k by $x(k)$, and its nearest neighbor in the set of representation values by $\widehat{x}(k)$. We shall assume that $\{x(k)\}$ is a discrete-time, wide-sense stationary (WSS) stochastic process [1] with zero expected value and power spectral density $S_{xx}(\omega)$.

The *quantization noise* is by definition $q(k) = \widehat{x}(k) - x(k)$. It may be viewed as the output of an additive noise source whose properties depend on the quantizer characteristic. Whatever the type of quantization used, the noise may generally be assumed to be a WSS process. We denote its power spectrum $S_{qq}(\omega)$ and its variance σ_q^2. For most practical cases the noise can be assumed to be white, i.e. $S_{qq}(\omega) = \sigma_q^2$.

In the case of *pdf-optimized* SQ [1] the additive noise is uncorrelated with the *output* signal, but correlated with the input. Assuming real-valued signals, this crosscorrelation is given by

$$\begin{aligned} R_{xq}(l) &= E[x(k)q(k-l)] \\ &= E[(\widehat{x}(k) - q(k))q(k-l)] = -R_{qq}(l) \end{aligned} \quad (1)$$

where $R_{qq}(l)$ is the quantization noise autocorrelation, equal to $\sigma_q^2\delta(l)$ if the quantization noise is white.

For *uniform* SQs with *midpoint representation values* on the other hand, a very good assumption, at least at high rates (i.e. 2 or more bits per sample), is that the noise is uncorrelated with the *input* signal. This is easily verified by experiments. The uniform case is important since it is known that uniform SQ followed by ideal entropy coding outperforms pdf-optimized SQ in the rate-distortion sense (also if entropy coding is applied) [1]. In this case, $R_{xq}(l) = 0$ for all lags l.

2. WIENER FILTERS AND QUANTIZATION

By definition, the *Wiener filter* is the least-mean-squared-error (or LMSE) optimal linear operator for estimating a desired signal, given an observed signal which is correlated with the desired signal in some known way. Assuming the desired signal $u(k)$ and the observed signal $v(k)$ to be WSS processes, the Wiener filter is ideally (with no causality or order constraints) given by [2]

$$W(\omega) = \frac{S_{uv}(\omega)}{S_{vv}(\omega)} \quad (2)$$

where $S_{uv}(\omega)$ is the cross spectral density associated with $u(k)$ and $v(k)$. This filter minimizes the variance of the reconstruction error signal $r(k) = u(k) - (w * v)(k)$, where the $*$ operator denotes discrete-time convolution.

In the case of SQ, $u(k) = x(k)$ and $v(k) = \widehat{x}(k)$. Exploiting the known auto- and crosscorrelation properties of the signals involved, we may derive the Wiener filters corresponding to the two types of quantization considered.

Starting out with the crosscorrelation $R_{x\widehat{x}}(l)$, we get

$$R_{x\widehat{x}}(l) = R_{xx}(l) + R_{xq}(l), \quad (3)$$

or, in the frequency domain,

$$S_{x\widehat{x}}(\omega) = S_{xx}(\omega) + S_{xq}(\omega). \quad (4)$$

A similar simple derivation yields

$$S_{\widehat{x}\widehat{x}}(\omega) = S_{xx}(\omega) + 2S_{xq}(\omega) + S_{qq}(\omega). \quad (5)$$

In general the Wiener filter minimizing the reconstruction error variance based on the observed quantized signal is thus

$$W(\omega) = \frac{S_{xx}(\omega) + S_{xq}(\omega)}{S_{xx}(\omega) + 2S_{xq}(\omega) + S_{qq}(\omega)}. \quad (6)$$

Exploiting the crosscorrelation expressions given in Subsection 1, this expression reduces to $W_{\text{pdf-opt}}(\omega) = 1$ for the pdf-optimized case, and

$$W_{\text{uniform}}(\omega) = \frac{S_{xx}(\omega)}{S_{xx}(\omega) + S_{qq}(\omega)} \quad (7)$$

for the uniform, midpoint-representation-value quantization case. Thus there is nothing to gain by postfiltering a pdf-optimally quantized signal, but there *is* something to gain by postfiltering a uniformly quantized signal. The reconstruction error variance after Wiener filtering of a uniformly quantized signal is given by [2]

$$\begin{aligned} \sigma_r^2 &= E[(x(k) - (\widehat{x} * w_{\text{uniform}})(k))^2] \\ &= \sigma_q^2 - \frac{1}{2\pi} \int_{-\pi}^{\pi} \frac{S_{qq}(\omega)}{1 + \frac{S_{xx}(\omega)}{S_{qq}(\omega)}} \, d\omega \\ &< \sigma_q^2 \text{ for all } S_{qq}(\omega) \neq 0. \end{aligned} \tag{8}$$

Without Wiener filtering, the reconstruction error variance is, of course, simply σ_q^2.

3. WIENER FILTERING AND CLOSED-LOOP DPCM

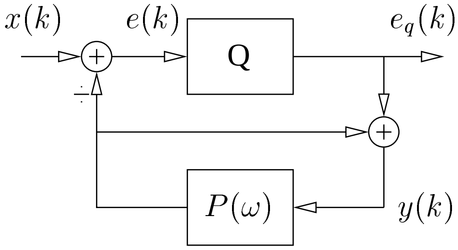

Fig. 1. DPCM encoder structure.

The encoder in a closed-loop Differential Pulse Code Modulation (DPCM) coder [1] is depicted in Figure 1. $P(\omega)$ is the *predictor* filter, while the filter $H(\omega) = 1 - P(\omega)$ is called the *prediction error* filter. The prediction error $e(k)$ is quantized and the quantized prediction error $e_q(k)$ is both output to subsequent storage or transmission, and input to the predictor filter. The predictor filter is typically optimized to provide minimum prediction error variance σ_e^2. If $P(\omega)$ is properly fitted to the input signal statistics, the closed-loop DPCM structure thus removes correlation and decreases signal variance, making it possible to quantize the prediction error efficiently by means of simple scalar quantization with few bits.

The corresponding *decoder* in a traditional DPCM codec filters the quantized prediction error by the *inverse* prediction error filter,

$$G(\omega) = \frac{1}{1 - P(\omega)} = \frac{1}{H(\omega)}, \tag{9}$$

to reconstruct a signal $y(k)$. Closed-loop DPCM, as shown here, has the advantage over *open-loop* DPCM [1] that the end-to-end reconstruction error, $r(k) = x(k) - y(k)$, is directly equal to the quantization noise $q(k) = e(k) - e_q(k)$ introduced in the prediction error signal. Thus we have introduced a *coding gain* relative to using direct quantization of $x(k)$ (PCM). However, as shall be seen, this is *not* the same as saying that the closed-loop DPCM structure as described above is optimal with regard to minimal end-to-end reconstruction error variance. There is something to gain by examining the components in the structure with a critical eye.

The closed-loop DPCM encoder structure in Figure 1 is easily shown to be equivalent to the system in Figure 2. Here $q(k)$ is the quantization noise added to the prediction error signal $e(k)$ inside the closed DPCM loop.

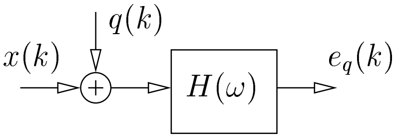

Fig. 2. DPCM encoder model with additive quantization noise.

We thus see that a closed-loop DPCM system is equivalent to a system where the input signal is contaminated by an additive noise source $q(k)$ (which is typically white), and where both signal and noise are subsequently shaped by the prediction error filter $H(\omega)$. However, this structure is again equivalent to the system in Figure 3.

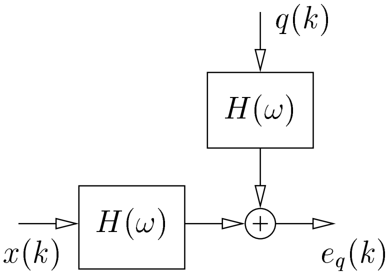

Fig. 3. Equivalent DPCM encoder model.

The equivalent additive noise source in Figure 3 has the effective power spectral density $S_{qq}(\omega)|H(\omega)|^2$.

From Figure 3 it again becomes immediately clear that if the variance σ_r^2 of the end-to-end reconstruction noise $r(k)$ is to be truly minimized, the appropriate receiver filter is *not* the inverse prediction error filter as used in traditional DPCM. Rather, it is the *Wiener filter* corresponding to the signal degradations depicted in Figure 3. The Wiener filter is in this case given by [2]

$$W(\omega) = \frac{S_{xe_q}(\omega)}{S_{e_qe_q}(\omega)}. \tag{10}$$

For the case of *pdf-optimized quantization*, a straightforward derivation shows that this filter is in fact equal to the inverse prediction error filter (a direct analogy to the unity filter obtained for pdf-optimized quantization alone). Thus, the Wiener filter perspective coincides with classic DPCM theory in this case. However, for the case of *uniform quantization at high rates*, the quantization error may again be assumed to be uncorrelated with the input, and the Wiener filter reduces to

$$\begin{aligned} W(\omega) &= \frac{H^*(\omega) S_{xx}(\omega)}{|H(\omega)|^2 (S_{xx}(\omega) + S_{qq}(\omega))} \\ &= \frac{1}{H(\omega)} \frac{S_{xx}(\omega)}{S_{xx}(\omega) + S_{qq}(\omega)}, \end{aligned} \tag{11}$$

which is recognized to be the exact inverse filter cascaded with the "noise compensation" filter given by Equation (7), which is also used in the case of uniform quantization only. This is not surprising; the quantization error is exactly the end-to-end noise after inverse filtering in traditional closed-loop DPCM. The LMSE-optimal receiver filter sacrifices the perfect reconstruction property in order to attenuate the quantization noise. The end-to-end reconstruction error variance in the closed-loop DPCM case is given by Equation (8), just as in the case of pure uniform quantization. Note, however, that σ_q^2 in the closed-loop DPCM case is reduced by the prediction gain factor compared to the case of standalone SQ.

4. EXPERIMENTAL RESULTS

We shall illustrate the results when $x(n)$ is an *1st order autoregressive* (AR(1)) process with normalized autocorrelation coefficient ρ, which implies that [1]

$$S_{xx}(\omega) = \frac{\sigma_x^2(1-\rho^2)}{1-2\rho\cos\omega+\rho^2}. \tag{12}$$

Figure 4 depicts experimental and theoretical signal-to-noise ratio (SNR) improvement for such a process when using optimal closed-loop DPCM encoding with uniform quantization and a Wiener filter at the receiver. The improvement is plotted for $\rho = 0.9$ and $\rho = 0.95$, as a function of the SNR *without* noise compensation filtering. The SNR range for each curve corresponds to a prediction error entropy range from approximately 1.1 to 3.5 bits per sample.

The experiments were done in MATLAB and are based on gaussian signal sequences of length 50.000 samples. The Wiener filter responses were designed using the frequency sampling method. In practice, we observe a close correspondence between theory and simulations down to approximately 11 dB for $\rho = 0.9$, and 14 dB for $\rho = 0.95$. This corresponds to approximately 1.3 bits per sample in both cases. Below this point, the actual gain is somewhat less than what is predicted by theory, although the results continue to improve somewhat as the SNR is lowered further.

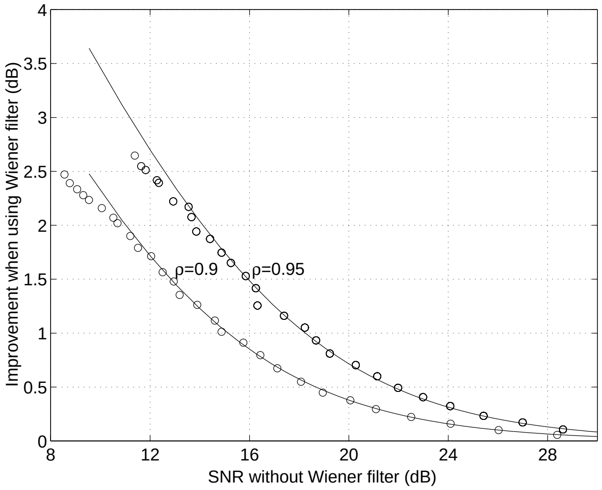

Fig. 4. Theoretical (solid lines) and experimental (circles) SNR improvement for AR(1) processes with autocorrelation coefficient $\rho = 0.9$ and $\rho = 0.95$.

5. CONCLUSIONS

Wiener filtering may significantly improve the fidelity of a quantized signal as long as the quantization noise is (close to) uncorrelated with the original signal. This model is valid for uniform scalar quantization with midpoint representation values at rates above approximately 2 bits per sample. A significant SNR improvement is experimentally observed also for rates down to 1.1 bit per sample, but below 1.3 bit per sample the improvement is not quite as large as that predicted by the model assuming no correlation between signal and noise.

Furthermore, postfiltering can never improve the coding result when pdf-optimized quantization is used, since the optimal filter response is always unity in this case.

The above conclusions hold both when the quantizer is standalone and when it is used as part of a closed-loop DPCM encoder. The optimal receiver filter in the latter case is simply the inverse prediction error filter in cascade with the same Wiener filter as derived in the case of pure quantization, and the additional coding gain due to Wiener filtering is the same as for quantization only.

Work is underway on extending these basic results to the case where a closed-loop DPCM signal is transmitted over an power-constrained, amplitude-continuous channel with additive channel noise. This is a good model for a certain promising joint-source-channel coding algorithm. In this case the transmit and receive filters may be jointly optimized to minimize overall reconstruction error. This model also includes open-loop DPCM on a transparent channel as a special case.

6. REFERENCES

[1] N. S. Jayant and P. Noll, *Digital Coding of Waveforms*. Prentice-Hall, 1984.

[2] A. K. Jain, *Fundamentals of Digital Image Processing*. Prentice-Hall, 1989.