Non-Subtractive Dither

Stanley P. Lipshitz, Robert A. Wannamaker, and John Vanderkooy
Audio Research Group, University of Waterloo, Waterloo, Ontario N2L 3G1, CANADA, Tel. (519)885-1211
and
J. Nelson Wright
Acuson, 1220 Charleston Rd., Mountain View, CA 94039-7393, USA, Tel. (415)969-9112

Abstract

A mathematical investigation of quantizing systems using non-subtractive dither is presented. It is shown that with a suitably-chosen dither probability density function (pdf), certain moments of the total error can be made signal-independent and the error signal rendered white, but that statistical independence of the error and the input signal is not achievable. Some of these results are known but appear to be unpublished. The earliest references to many of these results are contained in manuscripts by one of the authors [JNW[1]] but they were later discovered independently by Stockham and Brinton[2,3], Lipshitz and Vanderkooy[4], and Gray[5]. In view of many widespread misunderstandings regarding non-subtractive dither, it seems that formal presentation of these results is long overdue.

1 Introduction

Consider a typical mid-tread quantizer with a transfer characteristic of the form

$$f(x) = \Delta \left\lfloor \frac{x}{\Delta} + \frac{1}{2} \right\rfloor , \qquad (1)$$

where the floor operator, $\lfloor\ \rfloor$, returns the greatest integer less than or equal to its argument, and Δ represents one least significant bit (LSB) of the quantized output. Such an inherently non-linear operation always introduces a strongly input-dependent error into the signal unless preventive measures are taken. This error can represent significant signal modification for low-level or simple (e.g., sinusoidal) input signals, and will manifest itself as signal-dependent harmonic and intermodulation distortion in the output power spectrum.

One way to control the characteristics of the error is through the use of dither, a statistically independent random signal added to the input prior to quantization. A subtractive dithering scheme[6] exists which allows the error to be rendered statistically independent of the input signal. Unfortunately, such a scheme is impractical in many applications, since the dither signal must be available at both ends of the channel, requiring either its transmission or the existence of synchronized dither generators. A more frequently used scheme employs a non-subtractive dither signal, as shown in Fig. 1. We proceed to analyze in detail the properties of this architecture.

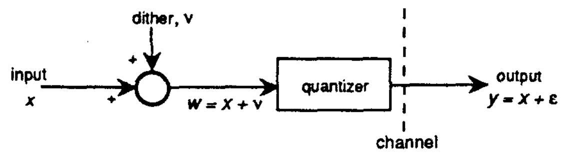

Figure 1: Schematic of a non-subtractively dithered quantizer.

2 First-Order Statistics

The dependence of the total error, ε, on the input, x, is described by the conditional pdf, $p_{\varepsilon|x}(\varepsilon|x)$, which represents the probability of a particular error, ε, given a specified input value, x. In order to derive an expression for this pdf, we note that the total input to the quantizer, $w = x + \nu$, has a conditional pdf

$$p_{w|x}(w|x) = p_\nu(w - x) , \qquad (2)$$

where $p_\nu(\nu)$ is the pdf of the dither signal. Furthermore, if w lies between $-\Delta/2 + k\Delta$ and $\Delta/2 + k\Delta$, then the total error is $\varepsilon = -x + k\Delta$, so that

$$p_{\varepsilon|x}(\varepsilon|x) = \sum_{k=-\infty}^{\infty} \delta(\varepsilon + x - k\Delta) \int_{-\frac{\Delta}{2}+k\Delta}^{\frac{\Delta}{2}+k\Delta} p_\nu(w - x)\, dw . \qquad (3)$$

This reduces to

$$p_{\varepsilon|x}(\varepsilon|x) = [\Delta \Pi_\Delta(\varepsilon) \star p_\nu(\varepsilon)] \sum_{k=-\infty}^{\infty} \delta(\varepsilon + x - k\Delta) , \qquad (4)$$

where the "$\star$" denotes convolution and where Π_Γ, the rectangular window function of width Γ, is defined by

$$\Pi_\Gamma(x) = \begin{cases} \frac{1}{\Gamma} & , -\frac{\Gamma}{2} < x \le \frac{\Gamma}{2} , \\ 0 & , \text{otherwise} . \end{cases} \qquad (5)$$

It is clear from Eq. 4 that the error cannot be made statistically independent of the input, since the functional dependence of the conditional pdf upon x cannot

[1] Wright, J.N., unpublished manuscripts, 1979 June-Aug.
[2] Stockham, T.G., private communication, 1988.
[3] Brinton, L.K., "Non-Subtractive Dither," *M.Sc. Thesis*, Dept. of Elec. Eng., Univ. of Utah, 1984 Aug.
[4] Lipshitz, S.P., and J. Vanderkooy, "Digital Dither," presented at the 81st Convention of the Audio Engineering Society, *J. Audio Eng. Soc. (Abstracts)*, vol. 34, p. 1030, 1986 Dec., preprint 2412.
[5] Gray, R.M., private communication, 1991.
[6] Schuchman, L., "Dither Signals and Their Effect on Quantization Noise," *IEEE Trans. Commun. Tech.*, vol. COM-12, pp. 162–165, 1964 Dec.

be eliminated by any choice of dither pdf. We will now show, however, that certain statistical moments of the error *can* be rendered independent of the input.

The m-th conditional moment of the error signal given x is defined as the expectation value of ε^m for a specified x:

$$E[\varepsilon^m|x] = \int_{-\infty}^{\infty} \varepsilon^m p_{\varepsilon|x}(\varepsilon|x)\, d\varepsilon \quad (6)$$
$$= \left(\frac{j}{2\pi}\right)^m \frac{\partial^m P_{\varepsilon|x}}{\partial u_\varepsilon^m}(u_\varepsilon, x)\Big|_{u_\varepsilon=0} \quad (7)$$

where $P_{\varepsilon|x}(u_\varepsilon, x)$ is the characteristic function of $\varepsilon|x$, the Fourier transform of its pdf. In the ensuing discussion, we adopt the following definition of the Fourier transform operator, $\mathcal{F}[\,]$:

$$\mathcal{F}[f](u) = \int_{-\infty}^{\infty} f(x)\, e^{-j2\pi ux}\, dx\,. \quad (8)$$

Taking the two-dimensional Fourier transform (with respect to both ε and x) of Eq. 4 and substituting into Eq. 7 yields the Fourier transform of $E[\varepsilon^m|x]$:

$$\mathcal{F}[E[\varepsilon^m|x]](u_x) = \left(\frac{j}{2\pi}\right)^m \frac{d^m G_\nu}{du_x^m}(-u_x) \times \sum_{k=-\infty}^{\infty} \delta\!\left(u_x - \frac{k}{\Delta}\right) \quad (9)$$

where

$$G_\nu(u_x) = \frac{\sin(\pi\Delta u_x)}{\pi\Delta u_x} \cdot P_\nu(u_x). \quad (10)$$

in which $P_\nu(u_x)$ is the characteristic function of the dither. If $E[\varepsilon^m|x]$ is to be independent of x we require that Eq. 9 reduce to a constant times a single delta function at the origin. This demands that

$$\frac{d^m G_\nu}{du_x^m}(u_x)\Big|_{u_x=k/\Delta} = 0 \quad ,\text{for } k = \pm1, \pm2, \ldots \quad (11)$$

in which case

$$E[\varepsilon^m|x] = E[\varepsilon^m] = \left(\frac{j}{2\pi}\right)^m \frac{d^m G_\nu}{du_x^m}(0) \quad (12)$$

independent of x. A dither pdf satisfying this condition for $m = 1$, so that the mean error is input-independent (and zero), is a rectangular pdf of the form

$$p_\nu(\nu) = \Pi_\Delta(\nu)\,. \quad (13)$$

A triangular-pdf dither resulting from the summation of two independent rectangular-pdf processes has the form

$$p_\nu(\nu) = \Pi_\Delta(\nu) \star \Pi_\Delta(\nu) \quad (14)$$

and satisfies the conditions for $m = 1$ and $m = 2$, so that both the mean and variance of the error are input-independent. In general, a dither generated from the summation of N independent rectangular-pdf processes satisfies the condition of Eq. 11 for $m \le N$.

Furthermore, it is not difficult to show that if $E[\varepsilon^m|x]$ is not a function of x, then

$$E[\varepsilon^m x^n] - E[\varepsilon^m]E[x^n] = 0 \quad (15)$$

for any choice of n. Hence a single rectangular-pdf dither is sufficient to render ε and x uncorrelated in the usual mathematical sense, while triangular-pdf dither also renders ε^2 uncorrelated with x.

3 Second-Order Statistics

We have explored the effects of non-subtractive dither upon the relationship between the error and the input signal, but the spectral character of the error has not yet been investigated. In order to do so we must consider the joint statistics of two error values, ε_1 and ε_2, separated by a time interval $t_1 - t_2 \neq 0$. Using a derivation analogous to the one used in the first-order case above, it can be shown that the joint moments of the two error values are given by

$$E[\varepsilon_1^m \varepsilon_2^n] = \left(\frac{j}{2\pi}\right)^{m+n} \sum_{k_1=-\infty}^{\infty} \sum_{k_2=-\infty}^{\infty} P_{x_1,x_2}\left(-\frac{k_1}{\Delta}, -\frac{k_2}{\Delta}\right) \times \frac{\partial^m \partial^n G_{\nu_1,\nu_2}}{\partial u_1^m \partial u_2^n}(u_1, u_2)\Big|_{u_1=-\frac{k_1}{\Delta}, u_2=-\frac{k_2}{\Delta}} \quad (16)$$

where

$$G_{\nu_1,\nu_2}(u_1, u_2) = \frac{sin(\pi\Delta u_1)}{\pi\Delta u_1}\frac{sin(\pi\Delta u_2)}{\pi\Delta u_2} P_{\nu_1,\nu_2}(u_1, u_2) \quad (17)$$

in which $P_{\nu_1,\nu_2}(u_1, u_2)$ is the joint characteristic function of the dither values at times t_1 and t_2. If ν_1 and ν_2 are statistically independent so that $G_{\nu_1,\nu_2}(\nu_1, \nu_2) = G_{\nu_1}(u_1)G_{\nu_2}(u_2)$, and if the dither statistics obey Eq. 11, then

$$\frac{\partial^m \partial^n G_{\nu_1,\nu_2}}{\partial u_1^m \partial u_2^n}(u_1, u_2)\Big|_{u_1=-\frac{k_1}{\Delta}, u_2=-\frac{k_2}{\Delta}} = 0 \quad \text{for } k_1, k_2 = \pm1, \pm2, \ldots \quad (18)$$

so that

$$E[\varepsilon_1^m \varepsilon_2^n] = E[\varepsilon_1^m]E[\varepsilon_2^n]\,. \quad (19)$$

In particular, simple rectangular-pdf dither with its samples statistically independent of one another ensures that

$$E[\varepsilon_1\varepsilon_2] - E[\varepsilon_1]E[\varepsilon_2] = 0\,, \quad (20)$$

so that the error samples are uncorrelated with one another and, hence, the power spectrum of the error is white.

We saw in Section 2 that triangular-pdf dither provides the further guarantee that this noise will have a constant variance (of $E[\varepsilon^2] = \Delta^2/4$, from Eq. 12), regardless of the input signal. In an audio signal such an error is perceived as a steady white noise with no audible relationship to the input[7], and it seems unnecessary to render any higher moments input-independent. For video applications, however, there is some evidence[3] that the third moment ($m = 3$) should also be made signal-independent.

[7] Vanderkooy, J. and S.P. Lipshitz, "Digital Dither: Signal Processing with Resolution Far Below the Least Significant Bit," *Proc. of the AES 7th International Conference: Audio in Digital Times*, Toronto, Canada, pp. 87–96, 1989 May.