

Uniform Quantization of Signals with Clipped and Truncated Laplace Distributions

Jerald L. Bauck*

July 4, 2020

Abstract

Many natural and man-made signals including much of speech and music are well-modeled by the Laplace distribution. Methods of synthesizing such signals are available and sometimes preferred over traditional test signals. However, sometimes those Laplace-like signals are clipped or do not exhibit infinitely-long tails. These situations are analyzed to determine their variances with an application of estimating signal-to-noise ratio as they are quantized by an analog-to-digital converter.

1 Background

MANY natural and man-made signals including much of speech and music are well-modeled in amplitude by the Laplace distribution [1]. Such signals, before processing, are frequently converted to digital form by a uniform quantizer—an analog-to-digital converter (ADC). Laplace signal models can also be used in testing and simulations thus requiring the synthesis of suitable signals [2]; in some usage cases the Laplace signal can be a substitute for or used in conjunction with the more common sine and uniformly distributed noise signals, providing for example more realistic real-world estimates of signal-to-noise ratio (SNR) of various equipment including ADCs and communication channels than provided by the more traditional signals [2].

Laplace signal models are characterized by a high sharp peak near the origin and long tails that are fatter than gaussian tails. However, those tails are infinite which is not always a good modeling feature and which do not fit into the finite range of a normal ADC. In addition, actual signals do not have infinitely long tails and in that respect the Laplace model fails, if modestly. Actual signals can have a Laplace-like distribution except for truncated tails, and synthesized Laplace-like test signals can also have finite tails by design or default, thus making them perhaps a better model than the unmodified Laplace shape. With these considerations there are two practical signal distributions that are derived from the Laplace distribution: a long-tailed distribution that has been clipped by the ADC or associated electronics and a truncated, finite-tailed distribution that has not

been clipped.

A naive use of the Laplace variance for the signal power in an estimate of the signal-to-noise ratio (signal-to-distortion ratio) of the quantizer can yield a number that is too high in cases where a clipped or truncated model is more appropriate. The error in SNR estimate depends upon the range and resolution of the quantizer and the width of the Laplace distribution that is assumed as a signal model. As it happens, the naive application of the usual Laplace variance might be suitable in some situations but the informed engineer knows when that approximation is appropriate.

The general equation for SNR in quantizer applications is

$$\text{SNR} = \frac{\sigma_X^2}{\sigma_E^2} \quad (1)$$

where variance σ_X^2 is the power of the signal and σ_E^2 is the power of the error signal due to quantization. The ratio is commonly expressed in decibels, $\text{SNR}_{\text{dB}} = 10 \log_{10}(\text{SNR})$. X shall be replaced herein by L for Laplace, C for clipped, and T for truncated, both for variances and probability density functions (PDF).

To be specific, when considering the tails of Laplace-like signals relative to the quantizer range there are several possibilities. One possibility is that the tails are small enough that ignoring the problem and using σ_L^2 causes little error in the resulting SNR estimate. Another possibility is that the signal is clipped either before or by the quantizer so that the signal power in the tails is captured in two particular quantizer regions, usually the lowest and highest such regions but sometimes other regions depending on how the signal was clipped. A third possibility is that the quantizer is designed in such a way that its lowest and highest regions are open, capturing at least in effect any signal values that extend respectively to $-\infty$ and $+\infty$; as

*Originally published on July 2, 2020. Available at engrxiv.org. This paper has not been peer reviewed. Jerry Bauck welcomes comments and can be reached by e-mail at this address after left-shifting each alpha character by one position, e.g., f becomes e: kfssz@cbvdl.ofu. © 2020 Jerald L. Bauck




Figure 1: Laplace probability density function for $\beta = 0.1$ with hypothetical lower and upper clipping or truncation levels indicated by x_m and x_M .

considered here, this is effectively the same as clipping the signal before encountering a quantizer with closed outer regions. Another possibility is that the signal does not have long tails but instead has distinctly finite tails that are well-modeled by the truncated distribution. And finally, there is the possibility that the signal has finite tails and is nevertheless clipped; this situation can be analyzed by the tools provided but the details are omitted.

2 Clipped Laplace Distribution

The Laplace probability density function (PDF) is

$$f_L(x) = \frac{1}{2\beta} e^{-|x-\mu|/\beta} \quad (2)$$

where $\beta > 0$ and for present purposes the mean $\mu = 0$ which is usually not restrictive for real-world signals. The variance, discussed below, is $\sigma_L^2 = 2\beta^2$. A plot for the central part of this PDF with $\beta = 0.1$ is shown in Figure 1.

Consider asymmetric clipping of signal $x(t)$ with the lower clipping level at $-x_m$ and the upper clipping level at x_M with $x_m > 0$ and $x_M > 0$. The clipped signal is defined as

$$x_c(t) = \begin{cases} x_M & x(t) \geq x_M \\ x(t) & -x_m < x(t) < x_M \\ x_m & x(t) \leq -x_m \end{cases}$$

The probability of overload, that is, clipping, on the negative side is $P_m = e^{-x_m/\beta}/2$ found by integrating (2) from $-\infty$ to $-x_m$ and similarly the probability of overload on the high side is $P_M = e^{-x_M/\beta}/2$ with a total overload probability $P_t = P_m + P_M$. With $x_m = 1$ and $x_M = 1$, $P_t = e^{-1/\beta}$ which is plotted as a function of β in Figure 2. With a selector func-




Figure 2: Probability of Laplace overload (clipping) as a function of β when $x_m = -1$ and $x_M = 1$.

tion defined as

$$c(x) = \begin{cases} 1 & -x_m < x < x_M \\ 0 & \text{otherwise} \end{cases}$$

the clipped PDF is

$$f_C(x) = P_m \delta(x + x_m) + c(x) \frac{1}{2\beta} e^{-|x|/\beta} + P_M \delta(x - x_M)$$

where the weighted Dirac impulses $\delta(\cdot)$ represent the probability masses due to clipping.

Consider the integral

$$I(x) = \frac{1}{2\beta} \int x^2 e^{-|x|/\beta} dx \quad (3)$$

the integrand of which is plotted in Figure 3 for $\beta = 0.1$. The indefinite integral is

$$I(x) = \begin{cases} e^{x/\beta} \left(\beta^2 - \beta x + \frac{x^2}{2} \right) & x < 0 \\ 2\beta^2 - e^{-x/\beta} \left(\beta^2 + \beta x + \frac{x^2}{2} \right) & x \geq 0 \end{cases} \quad (4)$$

which is plotted in Figure 4 again for $\beta = 0.1$. Definite integrals will be represented as $I(x) \big|_a^b$.

The variance of the clipped PDF is

$$\begin{aligned} \sigma_C^2 &= \int_{-\infty}^{\infty} x^2 f_C(x) dx \\ &= P_m x_m^2 + I(x) \big|_{-x_m}^{x_M} + P_M x_M^2. \end{aligned}$$

The integral is

$$\begin{aligned} I(x) \big|_{-x_m}^{x_M} &= 2\beta^2 - e^{-\frac{x_m}{\beta}} \left(\beta^2 + \beta x_m + \frac{x_m^2}{2} \right) \\ &\quad - e^{-\frac{x_M}{\beta}} \left(\beta^2 + \beta x_M + \frac{x_M^2}{2} \right). \end{aligned}$$




Figure 3: Integrand for computing Laplace variance σ_L^2 with $\beta = 0.1$.




Figure 4: Indefinite integral $I(x)$ of (3) related to the Laplace variance σ_L^2 with $\beta = 0.1$. The asymptotic value to the right is $2\beta^2$.

Identifying P_m and P_M within the expression, the clipped variance becomes

$$\sigma_C^2 = 2\beta^2 - e^{-x_m/\beta}(\beta^2 + \beta x_m) - e^{-x_M/\beta}(\beta^2 + \beta x_M). \quad (5)$$

The variance of the clipped signal is reduced from that of an unclipped signal with the same β because the impulsive masses in $f_C(x)$ are concentrated nearer to the origin than the extended regions of the tails from which they derive. As $x_m/\beta \rightarrow -\infty$ and $x_M/\beta \rightarrow \infty$, $\sigma_C^2 \rightarrow 2\beta^2 = \sigma_L^2$. The truncated variance is plotted against β in Figure 5 when $x_m = x_M = 1$ and against $x_M = x_m$ when $\beta = 0.1$ in Figure 6

3 Truncated Laplace Distribution

The truncated Laplace PDF is derived from the Laplace PDF generally by asymmetrically truncating the tails to the




Figure 5: Clipped Laplace variance σ_C^2 versus β with $x_m = x_M = 1$.

left and right, then adjusting the area to restore unity area which is required of a PDF. The area lost by truncating the tails is $P_m + P_M$ so the area of the truncated function is $1 - P_m - P_M$ implying that the truncated PDF is

$$f_T(x) = \begin{cases} 0 & x < -x_m \\ \frac{e^{-|x|/\beta}}{\beta(2 - e^{-x_m/\beta} - e^{-x_M/\beta})} & -x_m \leq x \leq x_M \\ 0 & x > x_M \end{cases}$$

The variance of the truncated PDF is

$$\begin{aligned} \sigma_T^2 &= \int_{-\infty}^{\infty} x^2 f_T(x) dx \\ &= \frac{1}{\beta(2 - e^{-x_m/\beta} - e^{-x_M/\beta})} \int_{-x_m}^{x_M} x^2 e^{-|x|/\beta} dx \\ &= \frac{1}{1 - P_m - P_M} I(x) \Big|_{-x_m}^{x_M}. \end{aligned}$$

Defining

$$\begin{aligned} B_m &= e^{-x_m/\beta} (2\beta^2 + 2\beta x_m + x_m^2) \\ B_M &= e^{-x_M/\beta} (2\beta^2 + 2\beta x_M + x_M^2) \end{aligned}$$

the variance becomes

$$\sigma_T^2 = \frac{4\beta^2 - B_m - B_M}{2 - e^{-x_m/\beta} - e^{-x_M/\beta}} \quad (6)$$

which is also reduced from σ_L^2 . As with the clipped case, as $x_m/\beta \rightarrow \infty$ and $x_M/\beta \rightarrow \infty$, $\sigma_T^2 \rightarrow \sigma_L^2$. The truncated variance is plotted against β in Figure 7 when $x_m = x_M = 1$ and against $x_M = x_m$ when $\beta = 0.1$ in Figure 8.




Figure 6: Clipped Laplace variance σ_C^2 versus $x_M = x_m$ with $\beta = 0.1$.




Figure 8: Truncated Laplace variance σ_T^2 versus $x_M = x_m$ with $\beta = 0.1$.

4 Signal-to-Noise Ratio

The SNR (1) requires the noise (distortion) power of the error of a uniform quantizer; the well-known result with identical lower and upper quantization ranges $-x_m = x_M = x_{\max}$ and quantized output represented by a w -bit word is

$$\sigma_E^2 = \frac{1}{3} \frac{x_{\max}^2}{2^{2w}}.$$




Figure 7: Truncated Laplace variance σ_T^2 versus β with $x_m = x_M = 1$.

If the entire Laplace signal range is assumed to fit in the ADC without significant error then

$$\text{SNR}_L = 2\beta^2 \frac{3 \cdot 2^{2w}}{x_{\max}^2}$$

and

$$\begin{aligned} \text{SNR}_{\text{LdB}} &= 10 \log_{10} 6 + 20 \log_{10} \beta + 20w \log_{10} 2 - 20 \log_{10} x_{\max} \\ &\approx 7.7815 + 20 \log_{10} \beta + 6.02w - 20 \log_{10} x_{\max}. \end{aligned}$$

Many times it is convenient to assume $x_{\max} = 1$ which drops the rightmost term. The pieces required to complete the SNR formulas for the situations studied are available above and will be indicated here for clarity. For the clipped Laplace case, (5) applies so that

$$\text{SNR}_C = \frac{\sigma_C^2}{\sigma_E^2} \quad (7)$$

and for the truncated Laplace case (6) applies and

$$\text{SNR}_T = \frac{\sigma_T^2}{\sigma_E^2}. \quad (8)$$

The forms of these variances do not lend themselves to convenient logarithmic expressions.

An application of (8) appears in [2] wherein methods of synthesizing classes of continuous deterministic bandlimited test signals with nearly Laplace distributions are described.

5 Version History

- July 2, 2020. First published.

- July 3, 2020. Corrected sign for x_m in expression for $x_c(t)$ and in text immediately following, and in expression for $f_T(x)$. Corrected signs on x_m and x_M in (5).
 - July 4, 2020. Re-rendered with properly sized Figure 5.
- [2] J. L. Bauck, “Generating continuous deterministic band-limited test signals with nearly laplace distribution,” *Submitted for publication, J. Audio Engineering Society*, 2020.

References

- [1] Free sound clips, sound bites, and sound effects. [Online]. Available: www.soundbible.com