Pallab Midya and Matt Miller
Motorola Labs
Schaumburg, IL, USA

Mark Sandler
King's College London
Strand, London, UK

Presented at the 109th Convention 2000 September 22-25 Los Angeles, California, USA

AES

AN AUDIO ENGINEERING SOCIETY PREPRINT

Integral Noise Shaping for Quantization of Pulse Width Modulation

Pallab Midya and Matt Miller
Motorola Labs, Schaumburg, Illinois

Mark Sandler
King's College, London

Abstract:

Integral Noise Shaping (INS) is introduced as a method for quantization of Pulse Width Modulation (PWM) that takes into account the fact that it is PWM that is being quantized. The quantization errors are noise shaped by integrating the ideal and quantized PWM waveforms analytically. The resulting quantizer can noise shape the rising and falling edges of two sided PWM through the same quantization process. This effectively doubles the oversampling ratio without any changes to the switching or clock frequency. This method results in a tremendous improvement in SNR. It is particularly useful for high end digital audio amplifiers.

1. Introduction:

Traditionally audio amplifiers are analog amplifiers usually operating in Class AB mode of operation. They tend to have low power conversion efficiency and consequently have large size and weight. With the advent of digital technology and in particular digital audio sources, it is highly desirable to perform the audio amplification function in the digital domain. Digital audio amplifiers using class D output stages tend to have a higher power conversion efficiency and a fixed switching frequency.

Digital Audio Power Amplification has become, for at least one if not two decades, the Holy Grail of Digital Audio research. Early publications on the subject were mostly by this author [1], and all relied on the use of PWM (Pulse Width Modulation) and Class D output stages. This work was firmly grounded in so-called Natural Sampling PWM, which is the normal form for analog signals. Key theoretical spec-

tral derivations for this and for Uniformly Sampled PWM appear in [2] while [3] presents a useful description of the various categories of Class D waveforms.
The early publications [1] presented conceptual systems which were clearly impractical, relying on modulator clock speeds of tens of GHz. The incorporation of Noise Shaping techniques as in [4] brought this down to several tens of MHz and allowed demonstration systems to be constructed, measured and compared against the simulation and theoretical predictions [5].

However, there were still distortions due to the underlying modulation process, and the next significant step was to incorporate non-linear processing to compensate for this. Several approaches were published in the early 1990's, including most notably, [6], [7], [8], [9]. All of these still relied on PWM as the core process for converting from standard digital representation (PCM – Pulse Code Modulation).

The next major development was in ways to use SDM (Sigma Delta Modulation) within a Digital Power Amplifier strategy. Two papers which describe different though related approaches to this are [10] which modifies the SDM so that fewer +1/-1 or –1/+1 transitions take place, and [11] which is a form of hybrid PWM/ SDM structure.

A recent research development is the proposal [12] to use Logan's Click Modulation [13]. This technique offers some advantages in that it is inherently linear in the signal band, but it still incurs a heavy DSP computation rate. The lower switching rates used in click modulation systems also imply the need for ultra-fast frequency roll-off in the passive L-C recovery filters which are necessary for all Class D output stages (see Fig. 1).

This paper deals with a new approach to noise-shaping the oversampled signal in a PWM-based Digital Amplifier. Unlike analog PWM, digital PWM is quantized and is generated by counting a high speed clock. Figure 1 shows a block diagram of a PWM based digital audio amplifier. The Integral Noise Shaping (INS) quantization algorithm described here dramatically improves the performance of the quantizer block.

Digital audio amplifiers based on quantized PWM typically handle one sample of the input data at a time [9]. These systems are either single sided PWM wherein one edge is fixed while the other is modulated, or symmetric PWM wherein both edges are modulated but the center of the pulse is fixed. Past methods of two-sided PWM

have noise shaped the rising edges and falling edges separately [5]. The effective oversampling ratio can be doubled if a stable, two-sided PWM specific noise shaping algorithm could be developed. INS is such an algorithm.

The paper starts with a brief overview of the theory behind PWM noise shaping loops, a description of the problems associated with applying traditional noise-shaping to PWM, and a description of the INS algorithm and its advantages. The following section illustrates the derivation of the INS algorithm, and presents a small signal analysis of the INS loop. The next section presents some ideas concerning stability of the loop, and also illustrates some design examples. The following section describes simulation results and compares performance of the INS algorithm to traditional PWM noise shaping loops. Finally, some concluding observations are presented.

2. PWM Noise Shaping:

PWM noise shaping loops are used to improve the in band signal to noise ratio of quantized PWM signals. The input to the loop is a high resolution digital signal which represents a series of duty ratios. The output of the loop is a lower resolution digital signal which also represents a similar series of duty ratios. The noise shaping loop forces the noise produced from quantizing the duty ratio signal to fall mostly out of band. PWM noise shaping loops have traditionally treated the duty ratio input signal just as if it were PCM - that is, it is simply a uniformly sampled data sequence. In fact, the PWM data sequence represents the time position of the PWM signal transitions. In this paper, integral noise shaping refers to the fact that the pulse width duty ratios, the input to the noise shaping loop, are treated as quantities that represent a signal that is continuous in time. This is a marked departure from the prior art and results in vast improvement in the performance of the system. The improvement comes at the price of more complex arithmetic that must be performed within the loop and significant changes in the way the loop is analyzed and evaluated for stability. The reason behind the performance improvement is quite simple. Noise shaping works by providing a data directed adjustment to the quantized output sample based on previous accumulated quantization error. The basic problem with PWM is that the position or timing of the correction is non-uniform. Treating the PWM duty ratio signal as if it were a uniformly sampled PCM sequence results in phase distortion of the correction. The Integral Noise Shaping algorithm takes

into account this difference between PCM and PWM, and slight modifications to the noise shaping loop eliminate the phase distortion problem.

Noise shaping can be implemented in part by computing an integral of all past error between the ideal signal and the quantized signal. It is possible to integrate PWM waveforms analytically to any order. The previous quantization errors between the ideal and quantized PWM waveforms can be noise shaped taking into account that we are dealing with PWM pulses.

Figure 2 shows a conceptual block diagram of the INS quantizer. Even though the actual algorithm is performed in discrete time in the digital domain, this figure shows the analog continuous time representation. The difference between the ideal and the quantized PWM is integrated and the result is used to correct for the error during the next sampling period.

In figure 3 the concept is taken to the digital domain such that the computation is done in discrete time at the switching frequency rate. Figure 3 shows an example schematic of an INS PWM quantizer for a CD data system. The CD data is over-sampled by 16X to 705.6kHz and then natural sampling of the interpolated signal is performed. Natural Sampling is a technique by which all in band harmonics due to the conversion from PCM to PWM are eliminated [9, 14]. This is similar to analog PWM which does not have in band harmonics.

The PWM duty ratio with high resolution, e.g. 24 bits, is noise shaped by adding a correction to the low resolution quantized pulse width based on integrals of the difference between the high and low resolution PWM waveforms. These integrals are calculated analytically. The noise shaped duty ratio is quantized to 7 bits for each half of the switching cycle to give a left half duty ratio and a right half duty ratio.

The integrals are calculated at the beginning of each half of the PWM cycle. The correction term is added directly to the high resolution duty ratio and the result is quantized to determine the low resolution duty ratio output. The duty ratio output controls a counter that provides the PWM output signal.

At this point, it is appropriate to begin analyzing the system in a more quantitative fashion. We will start with a derivation of the algorithm from a continuous time standpoint, proceed to a discrete time model, and then follow with a small signal analysis of the loop.

2.1 INS Derivation:

The error between the input and the output is integrated multiple times. Here a fourth order system is shown, but other orders are possible and appropriate based on the system requirements.

$$I_1(t) = \int_{-\infty}^{t} (X(\tau) - Y(\tau))d\tau \quad [1]$$

$$I_2(t) = \int_{-\infty}^{t} (I_1(\tau))d\tau \quad [2]$$

$$I_3(t) = \int_{-\infty}^{t} (I_2(\tau))d\tau \quad [3]$$

$$I_4(t) = \int_{-\infty}^{t} (I_3(\tau))d\tau \quad [4]$$

The variables corresponding to two sided PWM are illustrated in figure 4. The integrals are defined at discrete time intervals and computed analytically. This is possible because X and Y are pulse signals having values equal to zero and unity. Note that uppercase is used to represent continuous time variables and lower case is used for discrete time variables. All the computations are done in real time with the discrete time variables only. The discrete time variables are defined for the left and right side integrals at the beginning of the respective cycle. The integrals can be expressed in closed form. The values for the left half cycle are as follows.

$$il_1(n) = I_1(nT_s) = ir_1(n-1) + ((1 - yr(n-1)) - (1 - xr(n-1))) \quad [5]$$

$$il_2(n) = I_2(nT_s) = ir_2(n-1) + ir_1(n-1) + \frac{((1 - yr(n-1))^2 - (1 - xr(n-1))^2)}{2} \quad [6]$$

$$il_3(n) = I_3(nT_s) = ir_3(n-1) + ir_2(n-1) + \frac{ir_1(n-1)}{2} + \frac{((1 - yr(n-1))^3 - (1 - xr(n-1))^3)}{6} \quad [7]$$

$$il_4(n) = I_4(nT_s) = ir_4(n-1) + ir_3(n-1) + \frac{ir_2(n-1)}{2} + \frac{ir_1(n-1)}{6} + \frac{((1 - yr(n-1))^4 - (1 - xr(n-1))^4)}{24} \quad [8]$$

The values for the right half cycle are as follows.

$$ir_1(n) = I_1\left(\left(n+\frac{1}{2}\right)T_s\right) = il_1(n) + ((xl(n)) - (yl(n))) \quad [9]$$

$$ir_2(n) = I_2\left(\left(n+\frac{1}{2}\right)T_s\right) = il_2(n) + il_1(n) + \frac{((xl(n))^2 - (yl(n))^2)}{2} \quad [10]$$

$$ir_3(n) = I_3\left(\left(n+\frac{1}{2}\right)T_s\right) = il_3(n) + il_2(n) + \frac{il_1(n)}{2} + \frac{((xl(n))^3 - (yl(n))^3)}{6} \quad [11]$$

$$ir_4(n) = I_4(nT_s) = il_4(n) + il_3(n) + \frac{il_2(n)}{2} + \frac{il_1(n)}{6} + \frac{((xl(n))^4 - (yl(n))^4)}{24} \quad [12]$$

The corrected or predistorted duty ratio is defined for the left and right half cycles as follows.

$$zl(n) = xl(n) + k_1 il_1(n) + k_2 il_2(n) + k_3 il_3(n) + k_4 il_4(n) \quad [13]$$

$$zr(n) = xr(n) + k_1 ir_1(n) + k_2 ir_2(n) + k_3 ir_3(n) + k_4 ir_4(n) \quad [14]$$

Note the correction is computed using the values of the integrals at the beginning of the half cycle. This is an approximation since, in reality, the integral of the error term continues to accumulate up until the correction to the quantized PWM signal is actually made. Since the correction is a small, discrete quantity that changes slowly from sample to sample, the nonlinearity caused by uniform sampling is minor and does not result in significant distortion or additional noise.

The corrected duty ratio is quantized to fit an integral number of clock cycles. The result, which consists of a desired signal component and a quantization error component, is fed back to the input summing junction to be used for computing future corrections.

$$yl(n) = Quantize(zl(n)) = zl(n) + el(n) \quad [15]$$

$$yr(n) = Quantize(zr(n)) = zr(n) + er(n) \quad [16]$$

As described so far, the INS loop is obviously highly non-linear. This is not surprising considering the nature of the PWM signal and the fact that the corrections that are imposed upon the quantized PWM are time adjustments rather than amplitude adjustments. However, the INS loop has its roots in traditional noise shaping techniques, and it would seem that the loop might be analyzed using linearized models in a fashion similar to that used to analyze normal PCM noise shaping loops. Such an analysis is described in the next section.

2.2 INS Small Signal Loop Analysis:

We start with the familiar error feedback noise shaping loop shown in figure 5. The loop consists of an input and feedback summer, a quantizer, Q, and a finite impulse response filter, H(z). The transfer function of this loop in the z domain is easily found to be:

$$V(z) = U(z) + (1-H(z))E(z) \qquad [17]$$

Where V(z), U(z), and E(z) are the z transforms of the output signal, the input signal, and the quantization error respectively. In the first order low pass case, $H(z)=z^{-1}$ and the quantization error is simply high pass filtered with a zero at DC. Therefore, the output of the loop v(n) is simply the input to the loop plus an error term which is the high pass filtered quantization noise. Normally, the input to a loop of this sort is a sequence of numbers corresponding to the pulse code modulation (PCM) representation of an analog signal. However, this type of loop is also useful if u(n) is the natural sampled pulse width modulation (PWM) representation of an analog signal. In this case, u(n) is simply a sequence of numbers with each number representing a duty ratio. Unfortunately, because of the non-linear nature of PWM, the resulting output, v(n) -- a sequence of quantized duty ratios -- produces distortion terms and a higher noise floor when converted back to an analog signal than in the case when u(n) and v(n) are in PCM format.

Part of the problem stems from the loss of phase information in the loop processing. In an effort to fix this, we go to a modified version of a sigma-delta loop shown in figure 6. In this case, the quantizer is assumed to be a block which simply adds quantization noise to its input. This loop is similar to a typical sigma-delta loop but has an additional summing node where u(n) is added to the output of the noise-

shaping filter. For the first order case, $H(z)=z^{-1}/(1-z^{-1})$, a simple digital integrator, and the transfer function and performance are the same as the previous loop.

In order to improve the performance of higher order noise shaping loops when processing PWM signals, additional feedback terms are added as input to the noise shaping filter. These additional terms contain phase information that is normally lost in the error feedback PWM noise-shaping loop. The forms of these additional terms are derived from the higher order continuous time integrations of the difference between u(t) and v(t). The addition of these terms results in a higher degree of difficulty in analyzing the stability of the loop and also add a degree of computation complexity in implementation, but the benefits are significant.

We will first consider the second order case. Higher order cases are similar, but the algebra becomes much more tedious. The loop we will analyze is illustrated in figure 7. In figure 7, the noise-shaping loop is shown in an exploded view. Normally the filter would consist only of two cascaded integrators with their outputs summed with gains of $k_1=2$ and $k_2=1$. However, we have modified the filter with the addition of a non-linear term $[u^2(n)-v^2(n)]/2$. We will analyze this loop in the small signal sense - that is, we assume that v(n) is never very different from u(n) and the u(n) is near the middle of its range between 0 and 1. This means that we will approximate the $[u^2(n)-v^2(n)]/2$ term with a linear term. While we're at it, we'll write the approximations for higher order terms as well since they will come in handy when analyzing the higher order loops.

$$u^2-v^2 = u - v \quad [18]$$

$$u^3-v^3 = 3(u - v)/4 \quad [19]$$

$$u^4-v^4 = (u - v)/2 \quad [20]$$

$$u^5-v^5 = (u - v)/4 \quad [21]$$

Now getting back to our second order system, the expressions pertaining to $I_1(n)$ and $I_2(n)$ are the following.

$$I_1(n) = I_1(n-1) + u(n-1) + v(n-1) \quad [22]$$

$$I_1(z) = \frac{z^{-1}}{1-z^{-1}}[U(z) - V(z)] \quad [23]$$

$$I_2(n) = I_2(n-1) + I_1(n-1) + \frac{1}{2}u(n-1)u(n-1) - \frac{1}{2}v(n-1)v(n-1) \quad [24]$$

$$u(n-1)u(n-1) - v(n-1)v(n-1) = (u(n-1) + v(n-1))(u(n-1) - v(n-1)) \quad [25]$$

$$(u(n-1) + v(n-1)) \cdot (u(n-1) - v(n-1)) \approx 2u(n-1)(u(n-1) - v(n-1)) \quad [26]$$

$$2u(n-1) \approx 1 \quad [27]$$

$$I_2(z) = \frac{z^{-2}}{(1-z^{-1})^2}[U(z) - V(z)] + \frac{1}{2}\frac{z^{-1}}{(1-z^{-1})}[U(z) - V(z)] \quad [28]$$

Now we can write the transfer function for the entire loop.

$$V(z) = U(z) + \frac{E(z)(1-z^{-1})^2}{1 + \left(k_1 + \frac{1}{2}k_2 - 2\right)z^{-1} + \left(\frac{1}{2}k_2 - k_1 + 1\right)z^{-2}} = U(z) + NTF(z)E(z) \quad [29]$$

Here we see that the output V(z) is simply U(z) plus an error term. The change from a classic noise shaping loop is that the denominator adds poles to the high pass noise transfer function. The poles in the noise transfer function have the effect of reducing the maximum gain of the noise transfer function. As it turns out, this is beneficial. One could choose values for k_1 and k_2 which result in a noise transfer function that has no poles, in which case, the resulting noise transfer function would be equivalent to the error feedback loop noise transfer function. A similar approach of pole elimination could be used for higher order systems as well. Unfortunately, for most higher order PWM systems, the resulting loop is unstable. This brings us to the topic of noise transfer function design.

3.1 Noise Transfer Function Design:

In designing the noise transfer function, two main criteria must be met. The loop must be stable, and the noise transfer function should be well behaved in terms of its overall shape. In order to ensure stability of the noise shaping loop, the poles of the noise transfer function must be within the unit circle of the complex plane, and the maximum gain of the transfer function must be limited relative to the resolution of

the quantizer. Making sure the poles are inside the unit circle is straightforward. The second requirement for stability is similar to the situation for high order single-bit sigma-delta modulators. For single-bit modulators, it has been found that the maximum noise transfer function gain should be set to about 1.5 regardless of the order of the loop [15]. Since the noise transfer function is nothing more than a high pass filter, this constraint on the gain generally sets the cutoff frequency for the high pass filter and therefore determines the maximum attainable signal to noise ratio for the noise shaping loop. For multi-bit quantizers, the gain can be allowed to increase. If the quantizer resolution is high enough, then the maximum noise transfer function gain can be allowed to increase to the point where the poles in the noise transfer function are no longer necessary. However, this is generally not the case with our PWM noise shaping loop. For example, with a fourth order, two-sided INS system having a 7 bit quantizer for each half of the switching cycle, the maximum noise transfer function gain needs to be less than 5 to ensure stability. This result was derived empirically. For a similar fourth order PCM system, one would expect to be able to allow the maximum noise transfer function gain to be at least 16. The reason for the difference has to do with the fact that the INS system contains nonlinearities which are not accounted for in the small signal analysis. Remember that we arrived at this transfer function by assuming linear approximations for the nonlinear terms in the loop. The difference between the approximations and the actual nonlinear terms can be quite significant in the large signal sense.

Now for a few words about what is meant by a “well behaved” noise transfer function. It has been found in actual implementations of INS PWM DACs that a noise transfer function with large resonant peaks behave poorly in terms of signal to noise ratio and audio quality. This is because high frequency noise components tend to mix nonlinearly with phase noise components of the sampling clock or other clocks in the system and alias back into the signal band. It has been found that noise transfer functions with flatter passband characteristics yield better performance. This tends to steer us towards butterworth or low ripple chebyshev filter types for the noise transfer function.

At this point, we can begin illustrating some of the concepts governing the design of an INS PWM loop by going through a few examples.

3.2 Design Examples:

We will start with noise transfer functions which place all of the high pass filter zeroes at DC. A fourth order example will be shown. Then, we will modify the loop so that noise transfer function zeroes can be placed at arbitrary frequencies and, again, a fourth order case will be shown.

We follow a procedure similar to that outlined in [15].

1. Choose the loop topology.
2. Calculate the noise transfer function.
3. Pick a filter type (butterworth, chebyshev, type II chebyshev, etc.) and cutoff frequency.
4. Compute loop coefficients by matching to the filter function.
5. Test for stability and SNR.
6. If stability test fails, reduce noise transfer function gain by reducing filter cutoff frequency.
7. Go back to step 4.

First we will consider a fourth order system with all noise transfer function zeroes at DC.

Analysis of higher order loops is done in a manner similar to the second order case, but the algebra becomes quite a bit more bulky, so a symbolic math tool such as Mathematica is used to do the algebra. The general structure of the loop is the same as illustrated above with H(z) implemented as a weighted sum of a cascade of integrators I_j with each performing the following function.

$$I_j(n) = I_j(n-1) + I_{j-1}(n-1) + \frac{1}{2}I_{j-2}(n-1) + \ldots + \frac{1}{(j-1)!}I_1(n-1) + \frac{1}{j!}(u^j(n-1) - v^j(n-1)) \qquad [30]$$

Using the same assumptions for a fourth order loop that we used in the second order case, it can be shown with a little help from Mathematica that the noise transfer function is the following.

$$\frac{48(z-1)^4}{A_0 + A_1 z + A_2 z^2 + A_3 z^3 + 48 z^4} \qquad [31]$$

Where the form of each of the A_n is a linear combination of the four integrator weighting coefficients k_n. For this particular case, the equations are as follows.

$$A_0 = 48 + 48k_1 + 24k_2 - 6k_3 + k_4 \quad [32]$$

$$A_1 = -192 + 144k_1 - 24k_2 - 30k_3 + 23k_4 \quad [33]$$

$$A_2 = 288 - 144k_1 - 24k_2 + 30k_3 + 23k_4 \quad [34]$$

$$A_3 = -192 + 48k_1 + 24k_2 + 6k_3 + k_4 \quad [35]$$

Now, we move on to pick a type of transfer function filter. As we said earlier, a butterworth type filter is a reasonable choice because of its monotonic characteristics. We choose a cutoff frequency of fs/8 and find the denominator z coefficients using any number of filter design tools such as Matlab, and then after appropriate scaling, solve for the k values by matching coefficients. The resulting k_n values are:

$$k_1 = 1.3810,\ k_2 = 1.1452,\ k_3 = 0.5969,\ k_4 = 0.1634 \quad [36]$$

Using these values, we get a peak noise transfer function gain of 2.88. The noise shaping system using two sided PWM, a switching frequency of 352.8kHz, and a 7 bit quantizer for each side of the PWM pulse, yields an SNR of 105.4dB over a frequency range of 20kHz.

The above system places all of the noise transfer function zeroes at DC. It is well known that improvements in the performance of noise shaping loops can be had if the NTF zeroes are distributed across the bandwidth of interest. This can be done with a simple modification to the loop shown in Figure 7. We simply add a negative feedback path across each pair of integrators which creates a resonant pole in the feedforward transfer function, and therefore a corresponding zero appears in the noise transfer function. The algebraic relations describing this new modified loop are a little more complicated. The noise transfer function becomes.

$$\frac{(-48)(z^2 - 2z + G_1 z + 1)(z^2 - 2z + G_2 z + 1)}{A_0 + A_1 z + A_2 z^2 + A_3 z^3 - 48z^4} \quad [37]$$

Now the form of each of the A_n is again a linear combination of the four integrator weighting coefficients k_n. For this particular case, the equations are as follows.

$$A_0 = -48 + 48k_1 - 24k_2 - 18k_3 + 23k_4 \quad [38]$$

$$A_1 = 192 - 48G_1 - 48G_2 - 144k_1 + 24G_1k_1 + 48G_2k_2 + 24k_2 - 24G_2k_2 + 54k_3 - 6G_1k_3 + 5G_2k_3 - 23k_4 + 3G_1k_4 \quad [39]$$

$$A_2 = -288 + 96G_1 + 96G_2 - 48G_1G_2 + 144k_1 - 48G_1k_1 - 48G_2k_1 + 24G_1G_2k_1 + 24k_2 - 24G_2k_2 - 6k_3 + 6G_1k_3 + 18G_2k_3 - 3G_1G_2k_3 - 47k_4 + 3G_1k_4 \quad [40]$$

$$A_3 = 192 - 48G_1 - 48G_2 - 48k_1 + 24G_1k_1 - 24k_2 - 30k_3 + G_2k_3 - k_4 \quad [41]$$

Where the new terms G_n are the feedback gains from the output of the second integrator stage to the input of the first, and the output of the fourth stage back to the input of the third.

Now, we move on to pick a type of transfer function filter. In this case, we want a high pass filter with stopband attenuation ripple. A type II Chebyshev or elliptic filter is suitable. In this case, we choose a stopband attenuation and bandwidth, and the peak noise transfer function gain that results should remain below 5. For a 20kHz bandwidth, this results in k_n values of

$$k_1 = 1.5492,\ k_2 = 1.0651,\ k_3 = 0.6343,\ k_4 = 0.3141 \quad [42]$$

and G_n values of

$$G_1 = 0.0047,\ G_2 = 0.027 \quad [43]$$

Using these values, we get a peak noise transfer function gain of 3.78. The noise shaping system using two sided PWM, a switching frequency of 352.8kHz, and a 7 bit quantizer for each side of the PWM pulse, yields an SNR of 118.8dB over a frequency range of 20kHz.

4. Simulation Results:

At this point, we present the results of simulations of some INS PWM loops and compare the results to similar simulations of traditional noise shaping loops. Figures 8 through 13 contain the results of simulations of noise shaped PWM using conventional noise shaping and the Integral Noise Shaping algorithm. Figure 14 contains the result of a noise shaped PCM simulation. All of the simulations are of a 960Hz sinusoidal signal with a resolution limited only by the simulation software - a corresponding dynamic range greater than 300dB. For the PWM simulations, natural sampling is performed iteratively before the PWM is quantized - again, the starting dynamic range is greater than 300dB. The simulations were done this way to enable a fair comparison of the two PWM noise shaping algorithms.

Figure 8 shows a 1,693,440 point DFT of a quantized PWM signal which is generated using conventional noise-shaping. Each sample point of the PWM signal is either positive or negative one. The 1,693,440 point sequence represents 5880 PWM switch cycles with the quantized PWM signal being single sided and quantized to 8 bits. Each quantized PWM cycle consists of the 256 points corresponding to 8 bit quantization plus 16 points at the beginning and end of each cycle to guard against intersymbol interference in an actual switching amplifier application. Conventional third order noise shaping is performed on the PWM signal and the result is a signal to noise ratio of 95.5dB over a 20kHz bandwidth with flat weighting.

Figure 9 shows the result for a PWM signal quantized to 8 bits but this time using the INS algorithm to perform fourth order noise shaping on a two sided PWM signal. Unlike conventional noise shaping, the INS algorithm continues to produce SNR improvements for fourth order and higher order loops. The INS algorithm is also easily adapted for two-sided PWM, and most of the improvement in performance stems from this. The noise transfer function in this simulation approximates a fourth order Butterworth filter with a cutoff frequency of 44,100Hz. The resulting SNR is 105.4dB - nearly 10dB better than conventional noise shaping.

Figure 10 shows the result for a PWM signal quantized to approximately 6.5 bits using conventional noise shaping. The motivation for doing this would be to reduce the frequency of the PWM counter master clock. As expected, the reduced number of quantization levels reduces the SNR. One would normally expect to lose around 9dB of SNR, but in the case of conventional noise-shaping, the SNR is reduced 13dB to 82.5dB.

Figure 11 shows a similar result for the INS algorithm. In this case, the resulting SNR is 96.2dB. This is considerably better than the result for conventional noise shaping, and, comparing to figure 9, the INS algorithm produces no more loss than one would expect from reducing the number of quantization levels.

Figure 12 shows the result for the 8 bit quantized INS algorithm with noise transfer function zeroes distributed across the signal band rather than simply at DC. In this case, the noise transfer function approximates a type II Chebyshev high pass filter with stopband transmission 80dB down from the peak over a bandwidth of 20kHz. As expected, the signal to noise ratio improves with respect to the situation of Figure 9 by more than 13dB to an SNR of 118.8dB.

Figure 13 shows the result for the 6.5 bit quantized INS algorithm with the same noise transfer function as in Figure 12. Again, comparing to Figure 11, the SNR improves by nearly 13dB to a level of 108.6dB.

Figure 14 shows the result of PCM quantized to 8 bits using a conventional fourth order noise shaping loop. The sampling rate is 352.8kHz, and the resulting SNR is 102.2dB which is actually less than the 105.4dB SNR of the INS PWM signal illustrated in Figure 9. Given the history of PWM, this seems surprising. However, consider the situation from an information content point of view. Each PCM sample can only take on one of 256 values. The two sided PWM signal is split into two halves with each having 128 possible values, so each PWM cycle can take on one of 16384 values. It would seem that the two sided PWM has the potential for much higher levels of SNR compared to a similarly oversampled PCM signal.

Finally, figure 15 shows the output spectrum of a PWM digital to analog converter using INS implemented on a 24bit DSP. The switching frequency is 352.8kHz and the clock frequency 101.6MHz. A third harmonic is visible which is the result of the limitation of the natural sampling block of figure 1.

5. Conclusions:

INS is introduced here as a technique to improve the SNR of PWM digital audio amplifiers. The technique is valid for other digital power amplifier applications outside the audio realm. Performing PWM specific noise shaping makes it possible to quantize both the rising and falling edge of the PWM waveform. Utilizing two sided PWM effectively doubles the oversampling ratio for a given PWM counter clock frequency, which alone produces a very significant improvement in the SNR. In addition, because the INS algorithm preserves the phase information contained in the PWM signal and uses this information to compute the corrections to the quantized PWM signal, significant improvements in the signal to noise ratio are achieved. These benefits are gained at the cost of a small increase in system complexity and computation overhead.

6. Patent Protection

The INS algorithm and methods for implementing it that are described in this paper are protected by a Motorola, Inc. patent application.

7. References:

[1] M.B.Sandler, "Towards a Digital Power Amplifier", 76th AES Convention, 1984 October 8-11, New York.

[2] H.S. Black: ‘Modulation Theory’, 1953, D. Van Nostrand Co, Inc., Princeton, N.J.

[3] J.D.Martin, “Theoretical Efficiencies of class D power amplifiers”, Proc IEE, Vol.117, No.6, June 1970, pp 1089-1090

[4] M.Gerzon & P.Craven: “Optimal Noise Shaping and Dither of Digital Signals”, 87th AES Convention, New York, 1989, preprint 2822

[5] R.Hiorns & MB.Sandler: “Power Digital to Analogue Conversion using Pulse Width Modulation and Digital Signal Processing”, IEE Proc Pt G, Vol.140, No.5, Oct.1993,pp329-338.

[6] P.H.Mellor, S.P.Leigh & B.M.G.Cheetham, “Reduction in Spectral Distortion in class D amplifiers by an enhanced sampling process”, Proc IEE Part G, Vol. 138, No. 4, Aug. 1991.

[7] M.O.J.Hawksford, “Dynamic model-based linearization of quantized PWM for applications in Digital to Analogue Conversion and Digital Power Amplifier Systems”, J. Audio Eng. Soc., Vol. 40, No. 4, April 1992, pp. 235-252.

[8] P.Craven, “Towards the 24 bit DAC: novel noise shaping topologies incorporating correction for the non-linearity in a PWM output stage”, J. Audio Eng. Soc., Vol. 51, No.5, May 1993.

[9] J.M.Goldberg & M.B.Sandler, “Pseudo-Natural PWM for High Accuracy Digital to Analogue Conversion”, Electronics Letters, 1 August 1991, Vol. 27, No 16, pp. 1491-1492

[10] J. Magrath, M. B. Sandler: Digital power amplification using sigma-delta modulation and bit flipping, J. Audio Eng. Soc., 45, (6), pp. 476–487 (1997).

[11] A.Magrath & M.Sandler, "Hybrid PWM/SDM Power Digital to Analogue Conversion", IEE Proc on Circuits, Devices and Systems, Vol 143, No.3, June 1996, pp. 149-156

[12] M. Streitenberger, H. Bresch, W. Mathis: A new concept or high-performance class-D audio amplification, 106th AES Conv., 1999 May 8–11, Munich, Germany, Preprint 4941.

[13] F. Logan, Jr.: Click Modulation, AT&T Bell Lab. Tech. J., 63, (3), pp. 401–423 (1984).

[14] J. M.Goldberg & M. B. Sandler, “A New High Accuracy, Pulse Width Modulation based Digital-to Analogue Converter/Power Amplifier,” IEE Proc. Circuits, Devices Syst., pp. 314-324, August 1994

[15] S. R. Norsworthy, R. Schreier, G. C. Temes, et. al., “Delta-Sigma Data Converters -- Theory, Design, and Simulation,” IEEE Press, Piscataway, NJ, 1997.

Figure 1. Digital Audio Amplifier Schematic

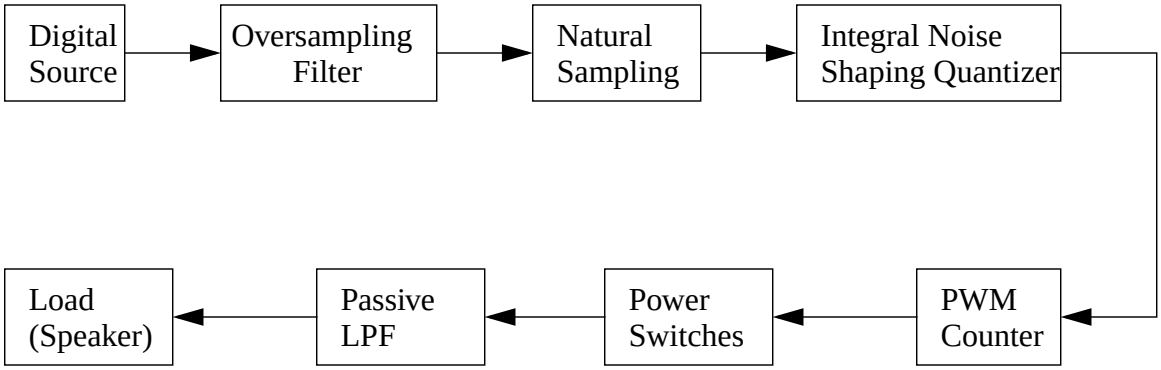

Figure 2. Integral Noise Shaping digital PWM Quantizer

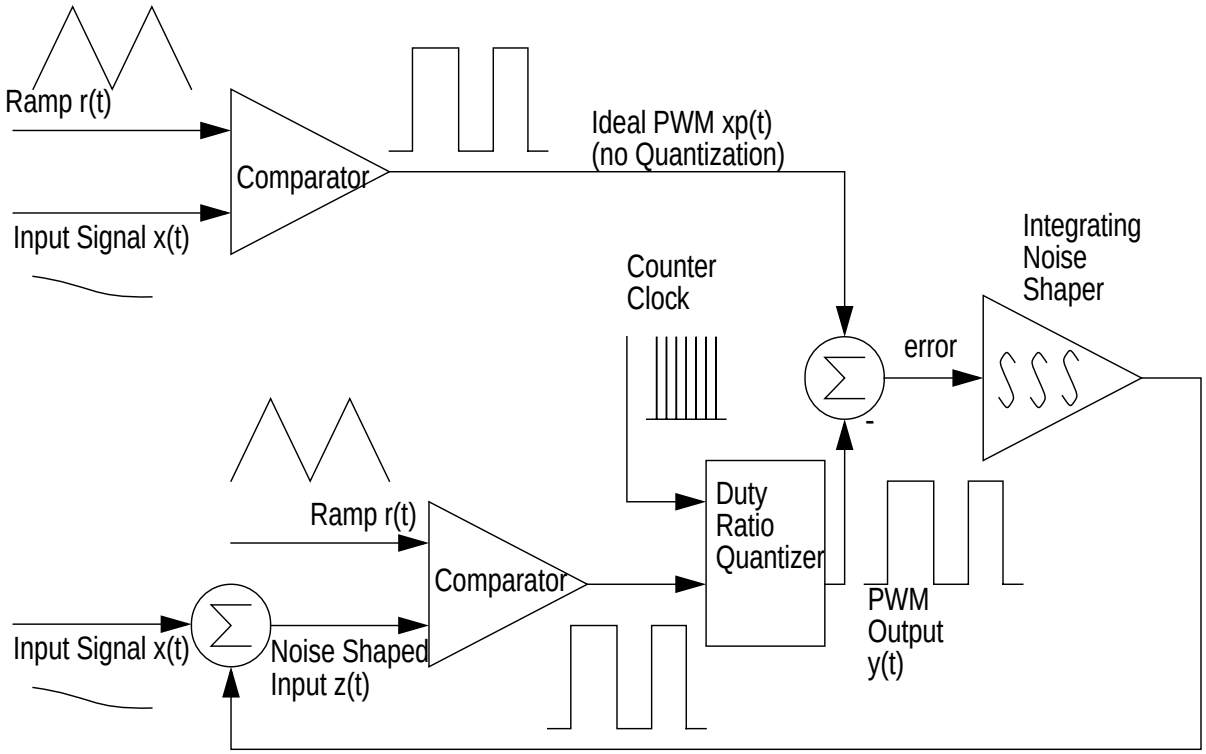

Figure 3: Implementation of INS Quantizer

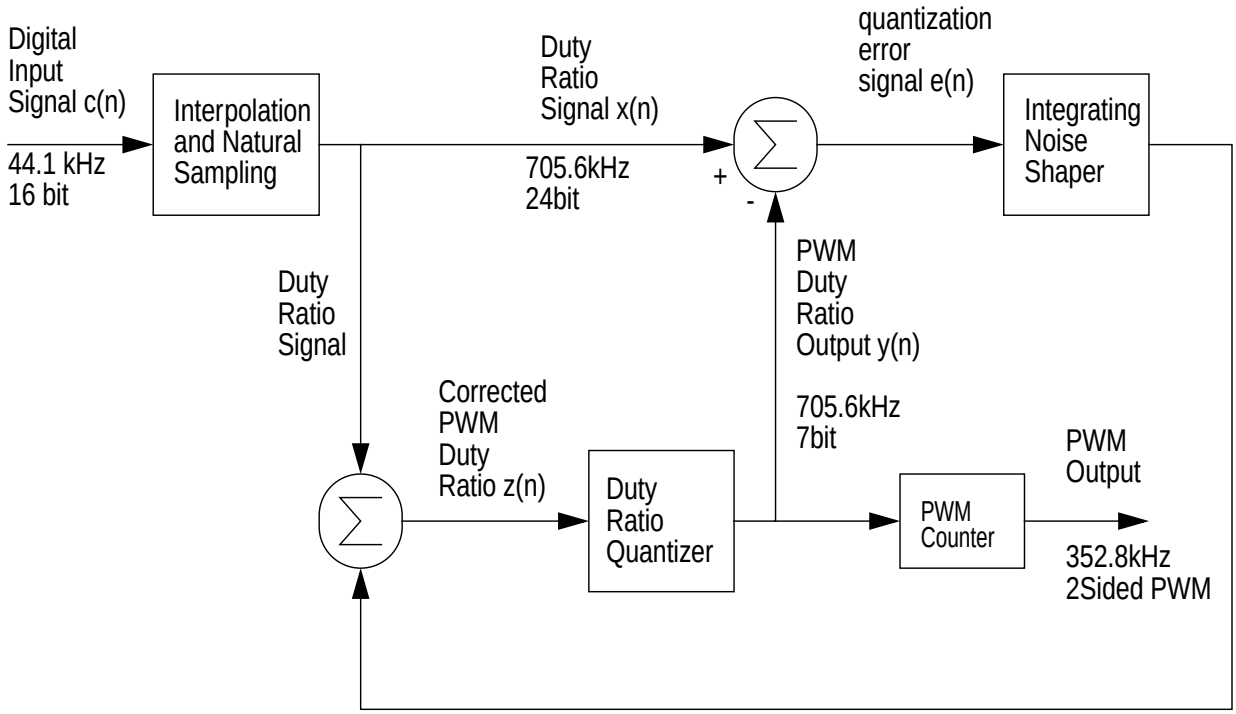

Figure 4. Timing Diagram for 2 Sided Integral Noise Shaping Quantizer

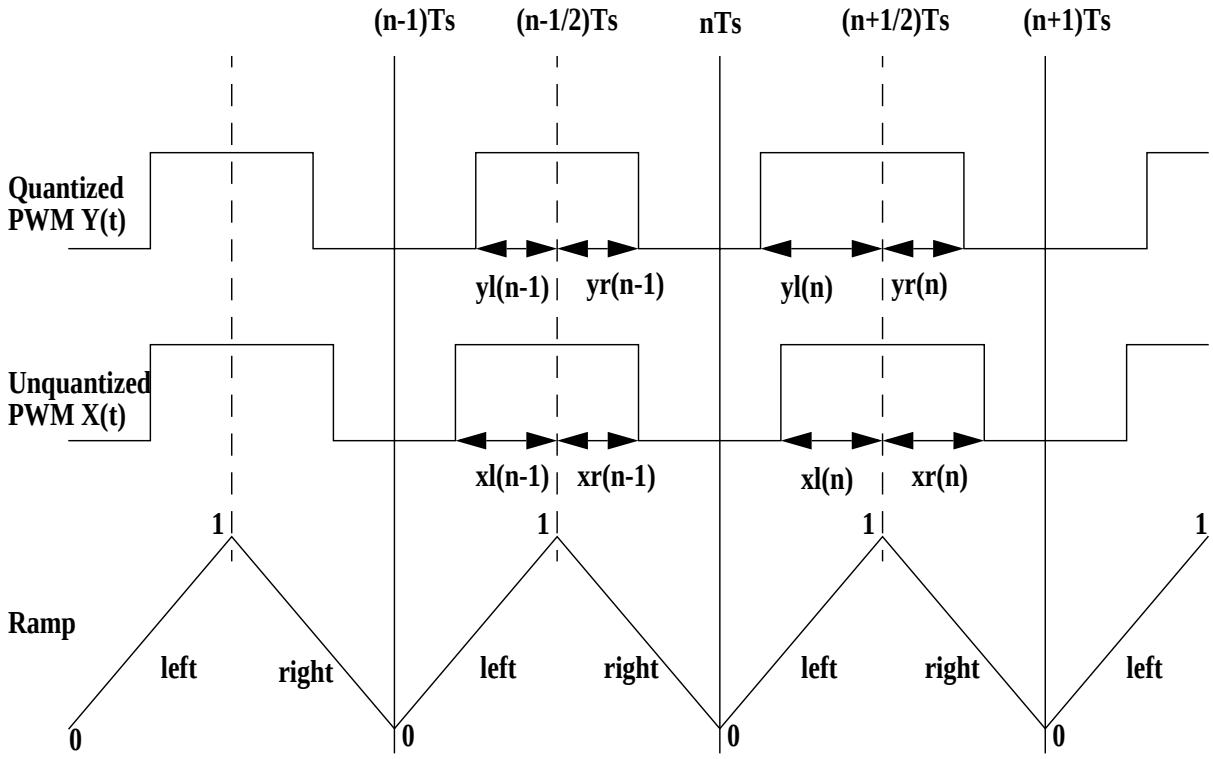

Figure 5. Error Feedback Structure of Conventional Noise Shaping Quantizer

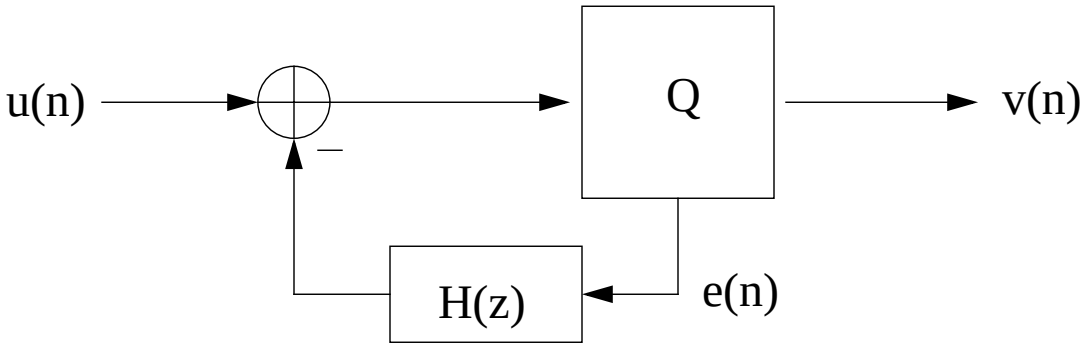

Figure 6. Feedback Structure of Integral Noise Shaping Quantizer

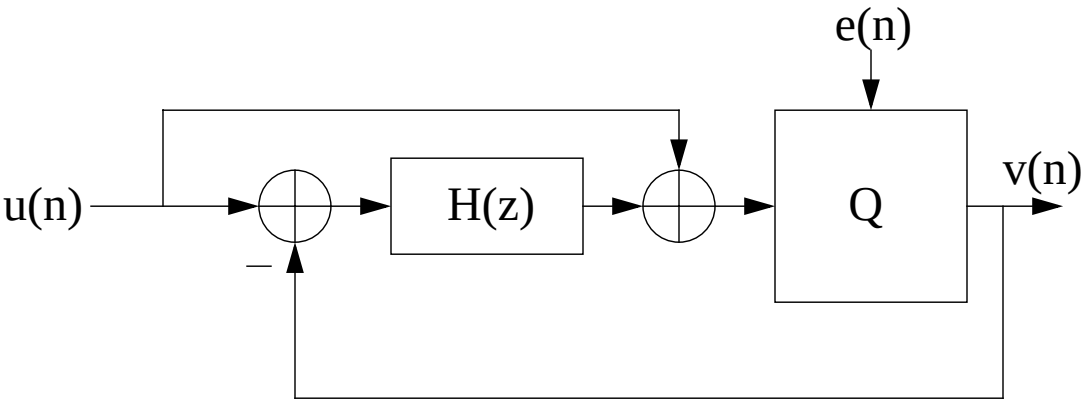

Figure 7. Structure of Integral Noise Shaping PWM Quantizer

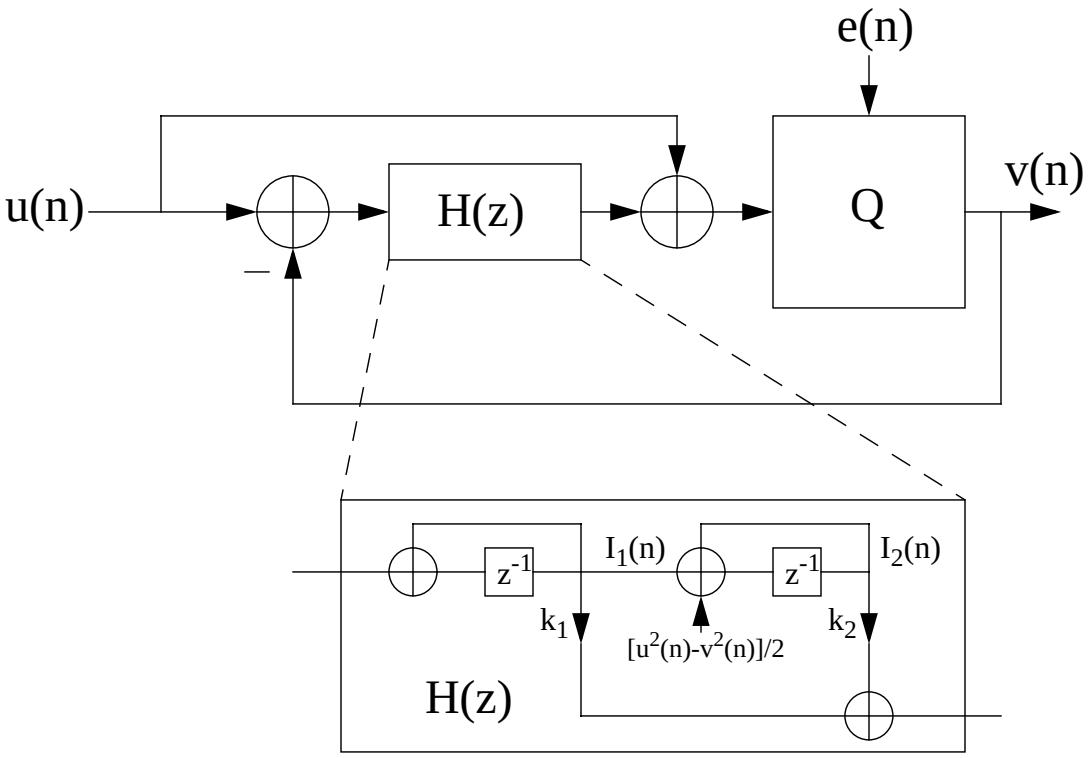

Figure 8. 1693440 Point DFT of Conventional NS PWM Signal

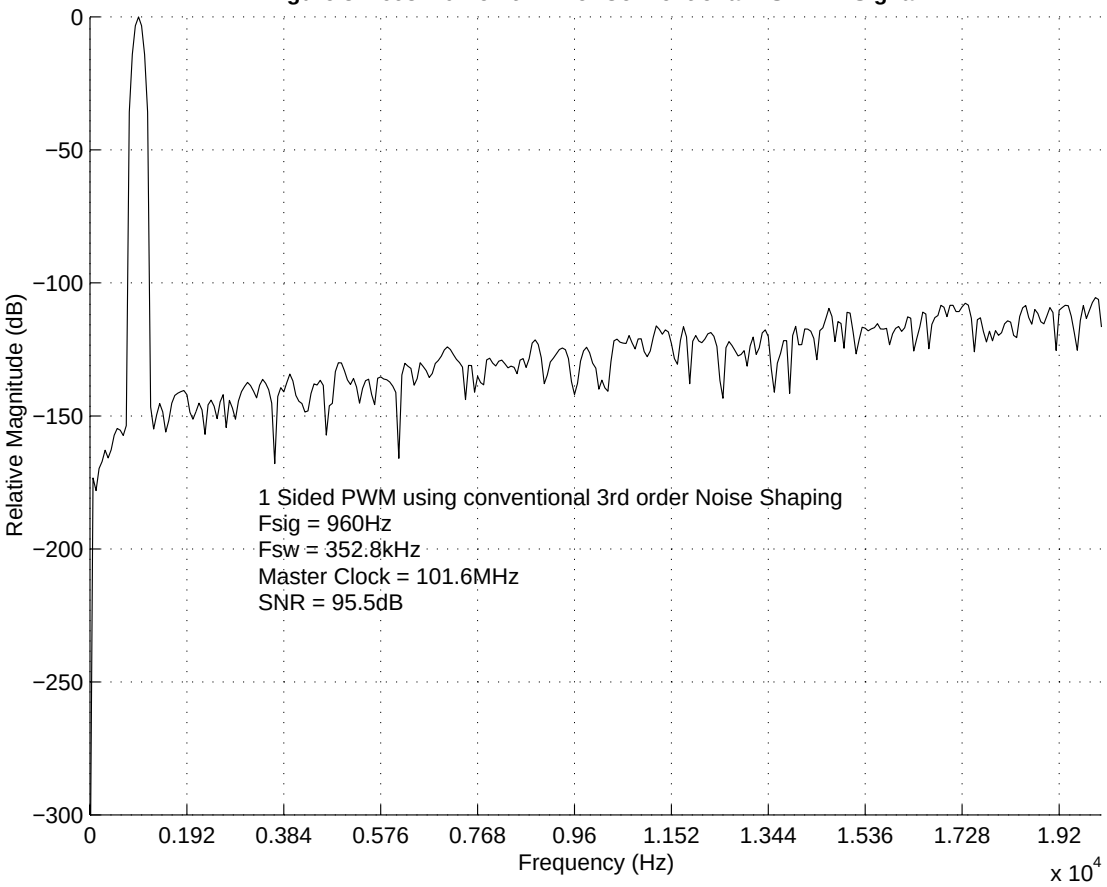

Figure 9. 1693440 Point DFT of INS PWM Signal

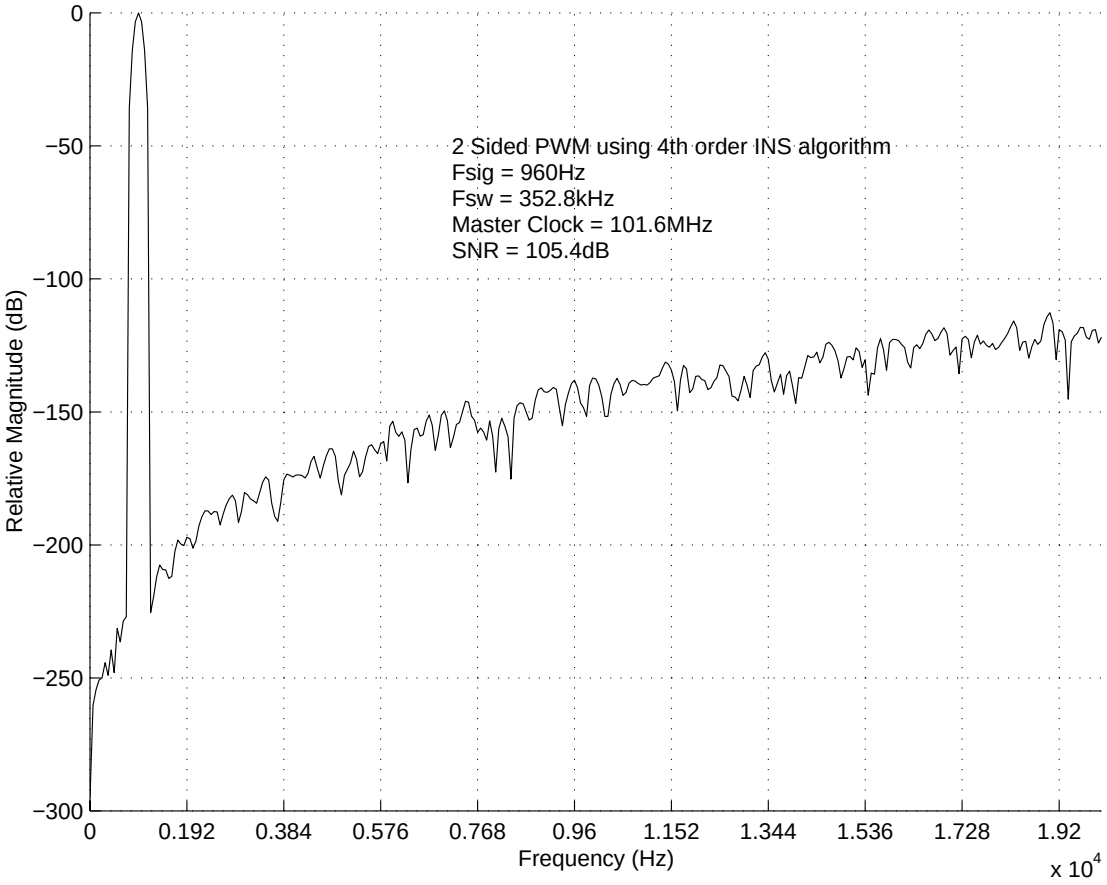

Figure 10. 564480 Point DFT of Conventional NS PWM Signal

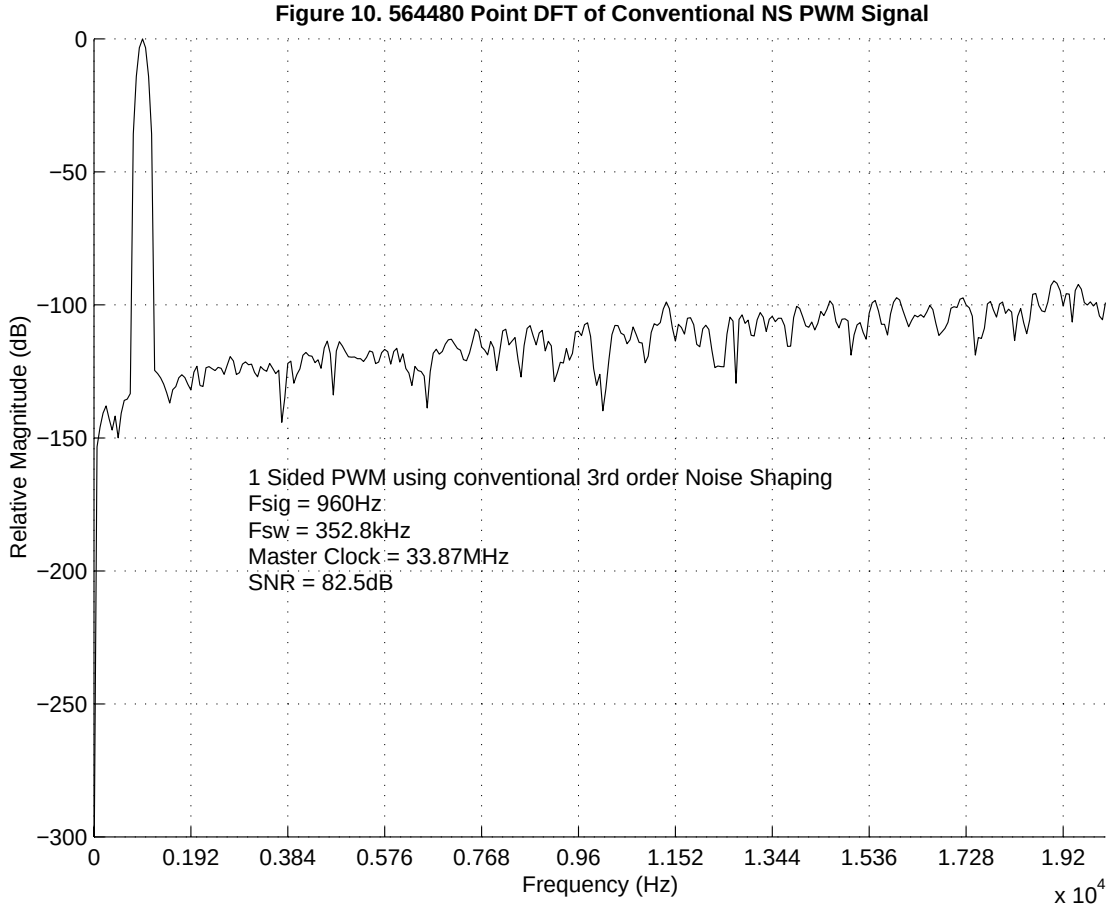

Figure 11. 564480 Point DFT of INS PWM Signal

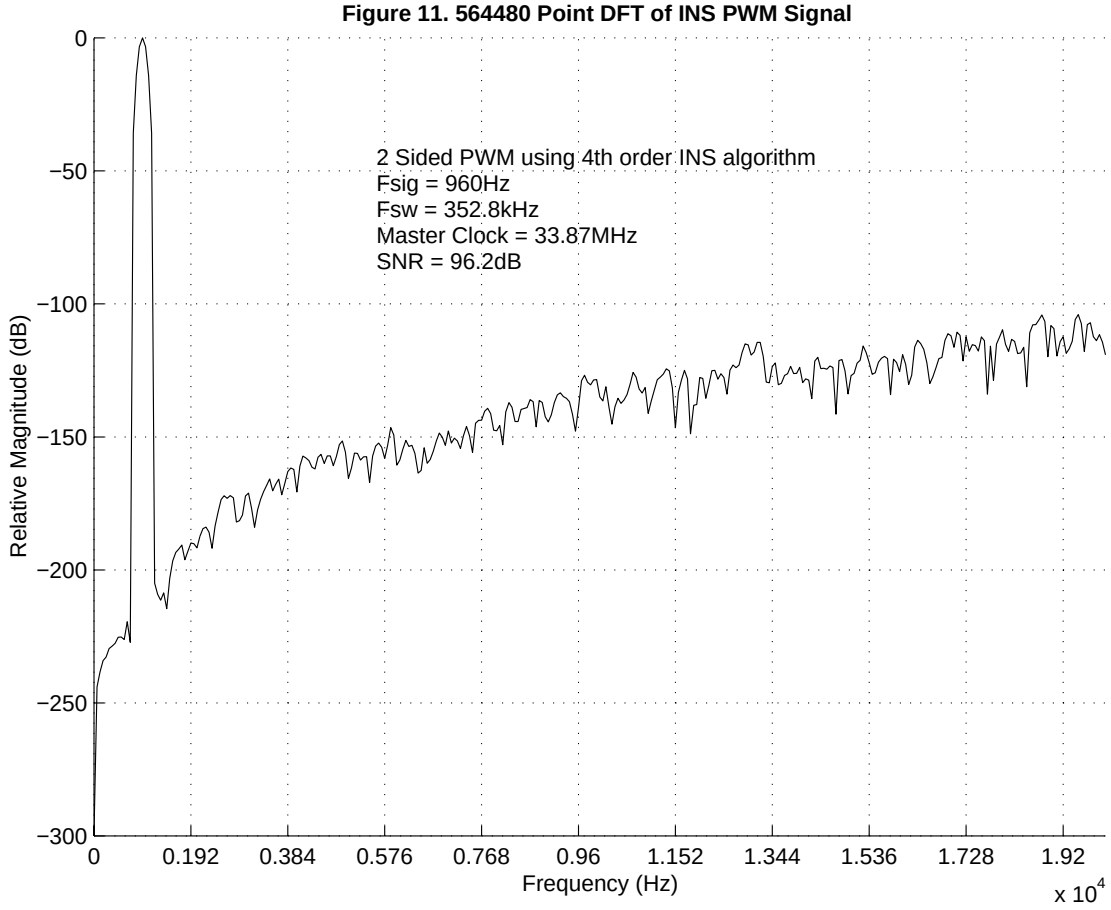

Figure 12. 1693440 Point DFT of INS PWM Signal w/ optimized zeroes

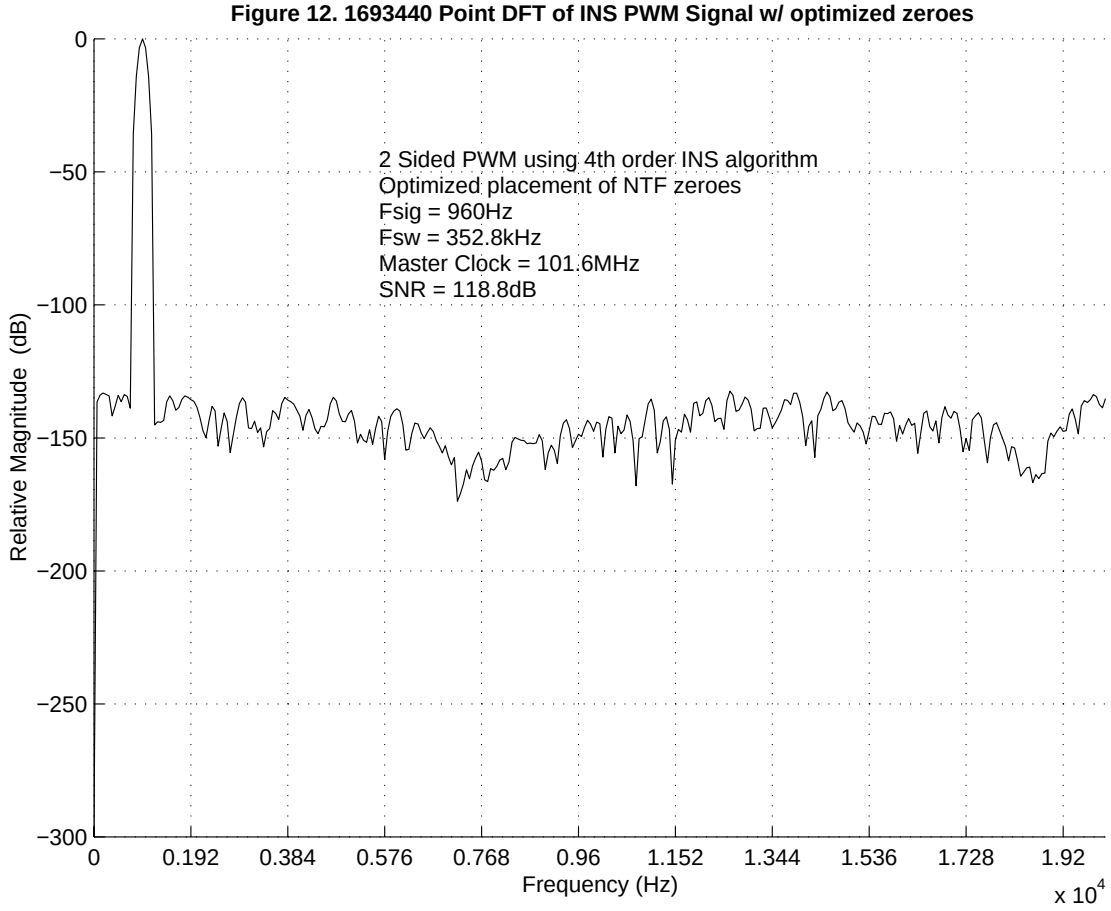

Figure 13. 1693440 Point DFT of INS PWM Signal w/ optimized zeroes

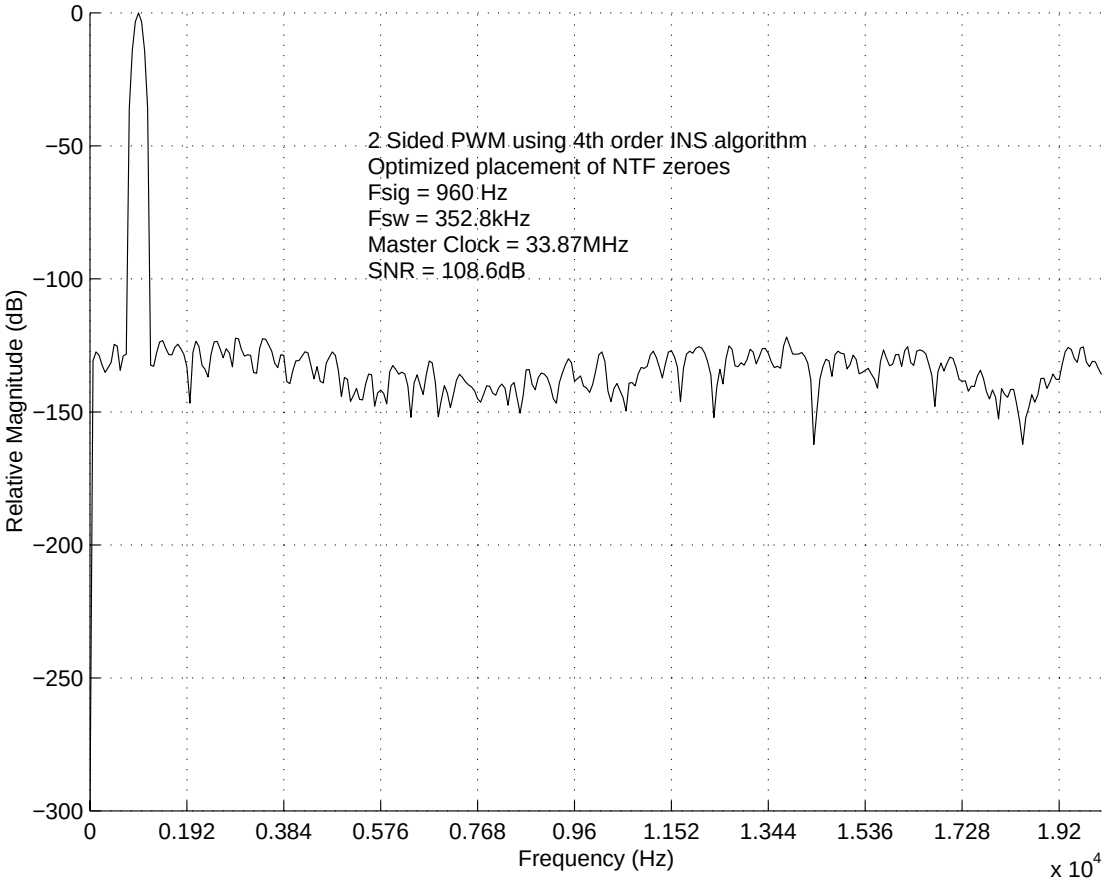

Figure 14. 47040 Point DFT of noise shaped PCM signal

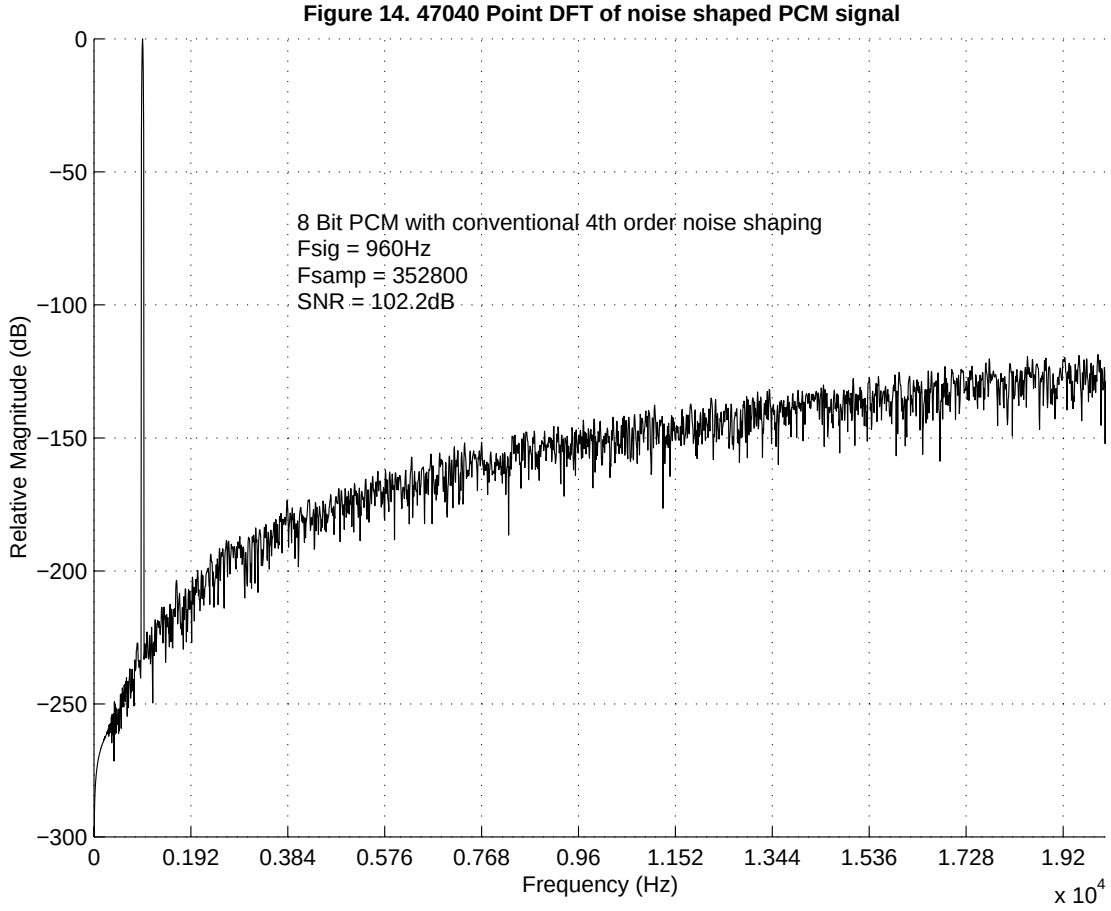

Figure 15. Spectrum of DSP PWM output signal