Audio Engineering Society

Convention Paper

Presented at the 120th Convention
2006 May 20–23 Paris, France

This convention paper has been reproduced from the author's advance manuscript, without editing, corrections, or consideration by the Review Board. The AES takes no responsibility for the contents. Additional papers may be obtained by sending request and remittance to Audio Engineering Society, 60 East 42nd Street, New York, New York 10165-2520, USA; also see www.aes.org. All rights reserved. Reproduction of this paper, or any portion thereof, is not permitted without direct permission from the Journal of the Audio Engineering Society.

Simple, Ultralow-Distortion Digital Pulse Width Modulator

Bruno Putzeys[1]

[1] Hypex Electronics B.V., Groningen, The Netherlands
bruno@hypex.nl

ABSTRACT

A core problem with digital Pulse Width Modulators is that effective sampling occurs at signal-dependent intervals, falsifying the z-transform on which the input signal and the noise shaping process are based. In a first step the noise shaper is reformulated to operate at the timer clock rate instead of the pulse repetition frequency. This solves the uniform/natural sampling problem, but gives rise to new non-linearities akin to ripple feedback in analogue modulators. By modifying the feedback signal such that it reflects only the modulated edge of the pulse train this effect is practically eliminated, yielding vastly reduced distortion without increasing complexity.

1. INTRODUCTION

Digital pulse width modulators are commonly used to drive H bridge output stages in power D/A converters (commonly referred to as "digital amplifiers"). Although the use of analogue controlled modulators with feedback is more warranted in this line of work, the present analysis highlights problems which occur equally in analogue and digital modulators. Digitally implemented, the level of complexity is not higher than that of the simplest prior art methods, while offering performance well beyond that available from the most complicated designs known to date. Academic interest aside, the new work could have applications in small-signal D/A conversion.

2. TERMINOLOGY AND DEFINITIONS

PWM	Pulse Width Modulation, Pulse Width Modulated, Pulse Width Modulator
NTF	Noise Transfer Function (noise shaper analogue of Error Transfer Function)
DC	(prefix) That which can be fully characterised using a time-invariant stimulus
f_s	Audio input sampling rate. f_s=48kHz unless stated otherwise.
f_r	The PWM pulse repetition frequency. $f_r = 8f_s$ unless stated otherwise
f_c	The PWM timing clock. Defines the timing resolution of the PWM signal.

3. HISTORY

In the eyes of the DSP engineer, the most obvious way of obtaining a high-quality analogue signal capable of driving a loudspeaker would be digitally generating a pulse-width modulated data stream based on PCM audio signals and "amplifying" this using a switching power stage (invariably assumed to be perfect). During the 90's this idea has led to numerous efforts to design the best possible PCM-to-PWM conversion algorithm.

3.1. f_r sampled method

The basic structure assumed by most modulators is that of an upsampling filter followed by a counter circuit.

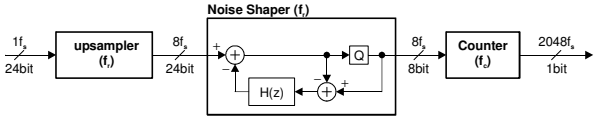

Figure 1: Basic f_r sampled noise shaped PWM generator

A noise shaper is inserted to ally realistic counter rates (f_c) with good in-band noise performance. Error-feedback noise shapers execute very efficiently in DSP, consisting of no more than an FIR filter, an addition and a truncation. The counter circuit can be made either to modulate one edge or two edges. Single-edged PWM signals offer a theoretical advantage over double-edged PWM in that any asymmetry between rise and fall times of the analogue waveform does not lead to errors. In reality, the variation of edge shape with output current is far greater than the symmetry error.

It was quickly realised that this algorithm is not ideal because the output is sampled at signal-dependent instants while the input is sampled at regular intervals. This effectively constitutes a variable delay or phase modulation. Phase-modulating a signal with itself generates distortion that becomes worse at higher signal frequencies. The problem is not limited to the signal. The shaped HF noise is subjected to the same nonlinearity and part of it demodulates into the signal band. A result of this is that higher orders of noise shaping actually worsen SNR performance.

Modulating both edges greatly reduces distortion by having two sampling instants moving in opposite directions, but the distortion is still significant, at least in "digital" terms. In fact, a double-sided "naïve" modulator already has sufficient performance for the system THD+N to be dominated entirely by any practical power stage – a fact that had obviously escaped the attention of the digital PWM vanguard.

Plot 1 (in annex) shows the spectrum out of a simple 4th order noise shaper. Plot 2 shows the spectrum produced by a single edged modulator processing the signal of plot 1 while plot 3 shows the double edged variant. The double edged variant has significantly better distortion and noise performance than the single edged variant. Either way, the low-frequency spectrum is still largely plastered over by demodulated HF noise.

3.2. Improved f_r sampled method

In a first improvement [1], the low-frequency error is corrected by predicting what the continuous-time equivalent of the input signal would be at the resulting PWM sampling instant. In effect, the 8fs PCM signal is pre-distorted in such a way that the modulation distortion is cancelled.

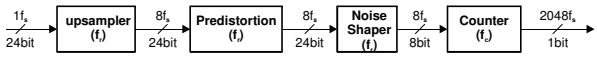

Figure 2: noise shaped PWM generator with predistortion.

Distortion products can be rendered arbitrarily small. The demodulation of shaped noise introduced by the PWM modulation is not addressed though. Not correcting for the demodulation of shaped noise turns the design of the noise transfer function into a delicate balancing act, as detailed in the same paper.

3.3. f_r sampled method with error model in the noise shaper

Both errors can be addressed simultaneously if the error introduced by the PWM process is placed inside the noise shaper loop.

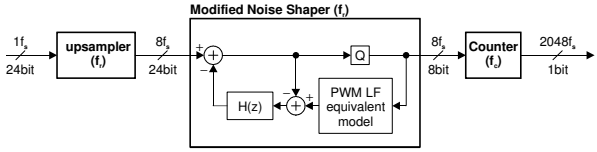

Figure 3: Low-frequency error model included into the noise shaper loop

In this way, the noise shaper "knows" what the spectral content produced by the PWM stage will be and thus apply its NTF to the correct variable. Doing this is less trivial than it seems. Depending on the signal, the PWM

process can sample anywhere from immediately to nearly a full sample later. The PWM model must include the effects of this variable time shift. Likewise, the noise shaper loop should be made not to overreact to late feedback, thus have some predictive ability [2].

This complication arises artificially, because the noise shaper uses a uniform sampling domain as its frame of reference and tries to correct for what is happening on another frame of reference that wobbles about relative to it. The complication disappears when the noise shaper is rewritten either to use the modulated edges as its frame of reference or yet another frame of reference that is uniform and common to both the input sample rate and the modulated signal. The latter proves quite a bit simpler than the former, because a common sampling rate is already present in the process. It is the high frequency timing clock f_c.

4. F_C SAMPLED METHOD

When the noise shaper is modified to operate at f_c, its loop gain function should be kept fixed relative to f_r. The error-feedback noise shaper explicitly controls only the poles in the loop function. The zeros follow automatically. Frequency-scaling requires control over poles and zeros separately. One implementation that offers all the required degrees of freedom is the common delta sigma coder. It also offers the advantage of easy hardware implementation, quite important given the much higher sampling rate. Its poles and zeros are set such that the loop gain is roughly equal to that produced by the original f_r sampled structure. The modulation stage becomes a digital implementation of the schoolbook analogue modulator.

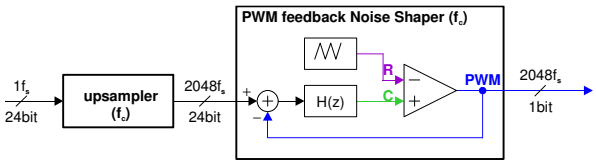

Figure 4: "delta-sigma-style" noise shaper with PWM inside the loop

The output from the control function is compared to a sawtooth or triangle reference. A triangle is found to be most workable. Single-edged modulation is also possible but requires a set/reset circuit after the comparator.

The complexity of a gate-level implementation is similar to that of a delta-sigma coder, around 3000 logic gates.

The feedback signal is no longer an approximation of the PWM signal. It is the PWM signal itself. At first sight the result should now be distortion-free and have no demodulated noise in the base band. Plot 4 shows that the latter appears true, but the former much less so: the distortion performance is far from ideal. Note that this modulator structure can be viewed as a discrete-time model of a continuous-time PWM modulator with feedback around it [3]. Distortion mechanisms arising in a continuous-time modulator have been previously identified [4] and arise in this algorithm as well.

4.1. DC linearity

A look at the signal at the comparator inputs shows what happens.

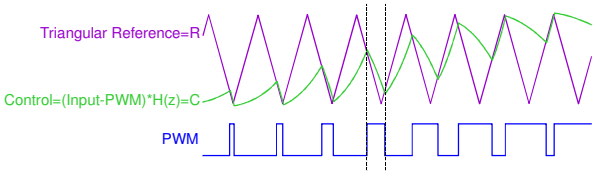

Figure 5: Waveforms in figure 4

The loop function provides only limited attenuation of the carrier component of the PWM signal. The residual is curvy and, more importantly, signal-dependent.

For a modulator to be linear DC-wise, the slope of the difference signal around the crossings has to remain constant over modulation index. In the absence of feedback this condition is automatically satisfied by the triangle wave. Note that the comparator is uninterested in what the signal does outside the crossings. It is not a necessary condition for the slope of the difference signal to be linear, although it would certainly be sufficient.

4.2. Variable bandwidth in double-edged PWM

That linearity leaves much to be desired can also be seen when comparing the outband noise at large (plot 4) and small (plot 5) modulation indexes. This effect is partly but not entirely explained by the DC non-linearity of the modulator. At low modulation indexes, the edges are approximately equidistant and effective sampling rate is twice f_s. At high modulation indexes, the edges move closer pair-wise, increasing redundancy. At 100%

modulation, the effective sampling rate has reduced to once f_s. The loop is effectively trying to control a system with a signal-dependent bandwidth. This too introduces distortion that is lower than that produced by the DC non-linearity but frequency dependent in nature. It is sure to show up if only the DC linearity problem is addressed.

Keeping the effective modulator bandwidth constant requires that we let not the rising edge know what the falling edge is doing [5]. By all means the simplest way to get there is not to modulate one edge at all so we can safely ignore it [6].

5. MODIFYING THE FEEDBACK SIGNAL

Being able to ignore the unmodulated edge in a single-sided modulator is at once the key to keeping the ripple more constant in amplitude and in phase compared to the modulated edge.

A single-edged PWM signal can be split up into two staircases.

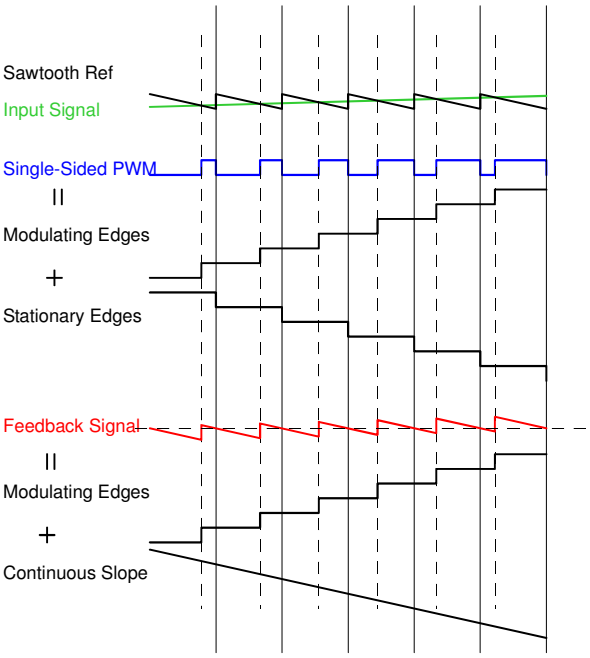

Figure 6: Breakdown of a single-edged PWM signal into two staircases

One staircase consists of only the rising (modulated) edges, the other the falling (unmodulated) edges. The unmodulated staircase contains no information and may just as well be replaced by a continuous linear slope. The signal is then put together again. Apart from f_r and its harmonics, the signal is perfectly identical to the original PWM and can be used as feedback signal in its stead.

The modified signal is a pseudo-sawtooth wave with the steps coinciding with the modulated edges of the original PWM. The zero-crossings of the signal coincide with the stationary edges. This initially counter-intuitive feature greatly simplifies implementation.

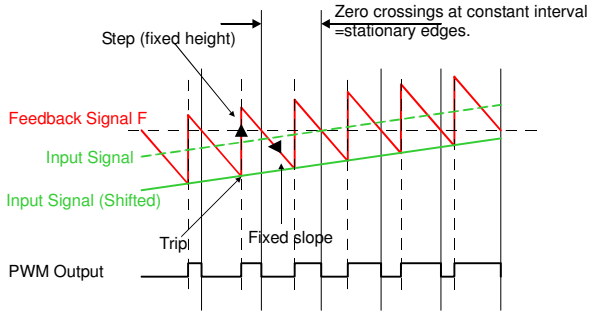

Figure 7: Direct generation of pseudo-sawtooth feedback signal. Subsequent derivative of PWM

Instead of going through the process of comparing the control signal C to a sawtooth and then stripping the stationary edges from the resulting PWM, the pseudo-sawtooth signal is generated directly from C and the PWM signal derived from there.

The input signal is offset by an amount equal to half the step size. The ratio of slope to step size sets f_r. The pseudo-sawtooth signal F is generated using a counter which is decremented by one every clock cycle and incremented by the step size when it crosses the input signal. The PWM signal is derived by slicing F.

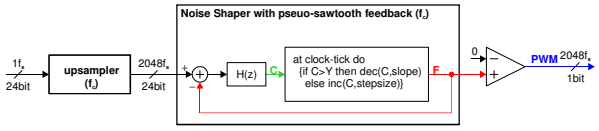

Figure 8: Direct generation of "pseudo-sawtooth" feedback signal inside a delta-sigma loop.

A plot from the closed loop in operation shows that the residual is now very nearly constant in shape and amplitude, and its constant phase relationship with the sampling instant is obvious.

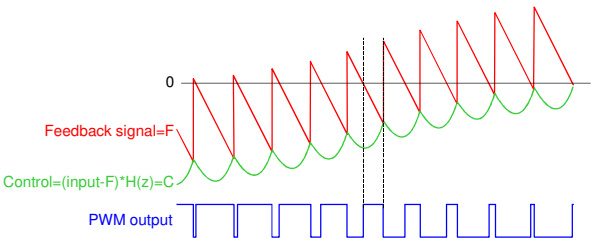

Figure 9: Control signal, pseudo-sawtooth and PWM with the loop closed.

The two conditions for linearity are now fulfilled: the spacing of the sampling instants is highly constant and the slope of the comparator input is the same at every crossing.

As a bonus, the pseudo-sawtooth is not amplitude limited. The PWM signal clips when the pseudo-sawtooth no longer has zero crossings, but this happens outside the loop. This modulator circuit can never overload.

Plot 6 shows the output spectrum. The loop function is identical to that used in figure 4. All distortion products have disappeared from view. Comparing the noise spectrum with that obtained with a small input signal (plot 7) verifies that loop gain and bandwidth are indeed completely constant.

It is noted that the noise shaper effectively refers only to the real (signal-dependent) sampling instants and no reference is made to the uniformly spaced f_r sampling domain. Ideally this would require the upsampling filter to produce samples at f_c. In practice, performance does not suffer if linear interpolation is used from f_r upward.

The complexity of the current modulator is identical to the original PWM feedback circuit, i.e. about 3k logic gates.

6. A PROTOTYPE DESIGN FOR SMALL-SIGNAL D/A CONVERSION

As said before, this effort is rather pointless if the planned use is to drive a power stage. The single-edged modulation and the very low distortion can be taken advantage of in continuous-time small signal D/A conversion stages. An implementation geared towards this will have a higher f_r and a lower f_c. The sample design operates at f_c=49.152MHz and f_r=2.8224MHz. A 7th order loop is employed in order to clear out the bottom 80kHz, providing full 192kHz compatibility. The design is a worst-case implementation with a pipeline delay at every integrator and single-bit coefficients.

Plots 8 to 11 show the output spectra of this design under several signal conditions. The noise floor in the plot 8 is dominated by windowing artefacts. The performance obtained can be summed up as:

- THD+N (20kHz bandwidth) < -135dB
- THD+N (80kHz bandwidth) < -129dB

Obtaining matching performance from the analogue circuitry is considered very difficult but not impossible.

7. CONCLUSION

A method for digital PWM generation is presented that is as simple as a common delta-sigma coder while offering audio performance well beyond that of much more complicated designs.

(Note: Patent has been applied for)

8. REFERENCES

[1] Sandler et al., “Ultra-Low Distortion Digital Power Amplification”, AES Preprint 3115, Presented at the 91st AES Convention

[2] Peter Craven, “Toward the 24-bit DAC: Novel Noise-Shaping Topologies Incorporating Correction for the Nonlinearity in a PWM Output Stage”, J.AES Vol 41/5 pp. 291-313.

[3] Lars Risbo, “Discrete-time Modeling of Continuous-time Pulse Width Modulator Loops”, Paper 3-5, Presented at the 27th AES Conference.

[4] Bruce Candy, “Improved Analogue Class-D Amplifier with Carrier Symmetry Modulation”, AES Preprint 6260, Presented at the 117th AES Convention.

[5] Matthew 6:3

[6] Matthew 5:30

9. ANNEX: SPECTRAL PLOTS

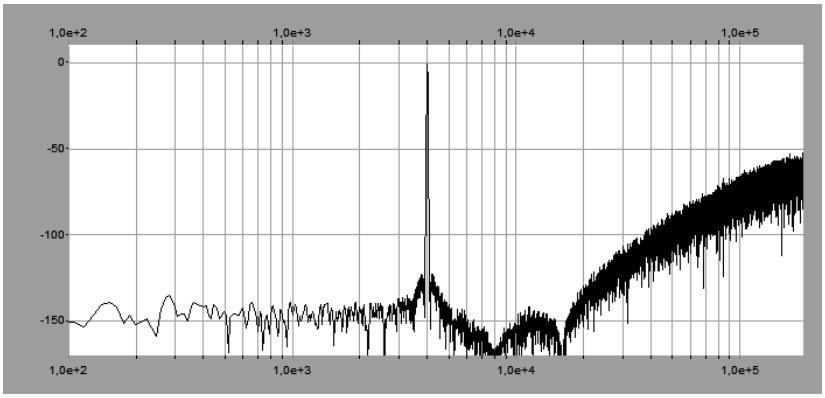

Plot 1: Output spectrum after the noise shaper of figure 1 with -1dB stimulus, before pulse-width modulation.

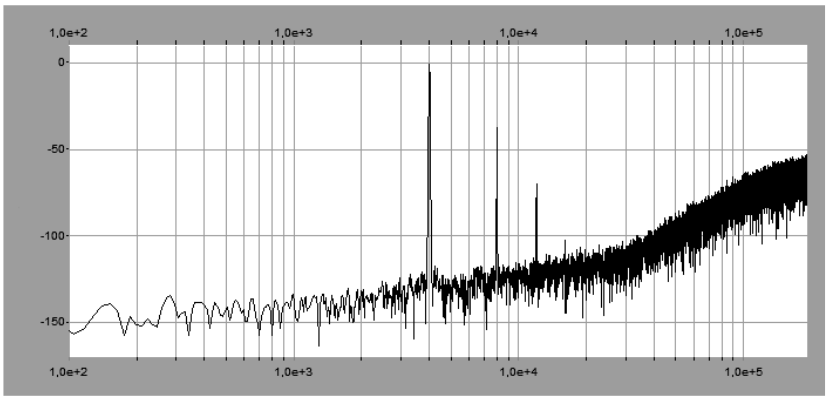

Plot 2: Output spectrum of circuit from figure 1 using a single-sided modulator

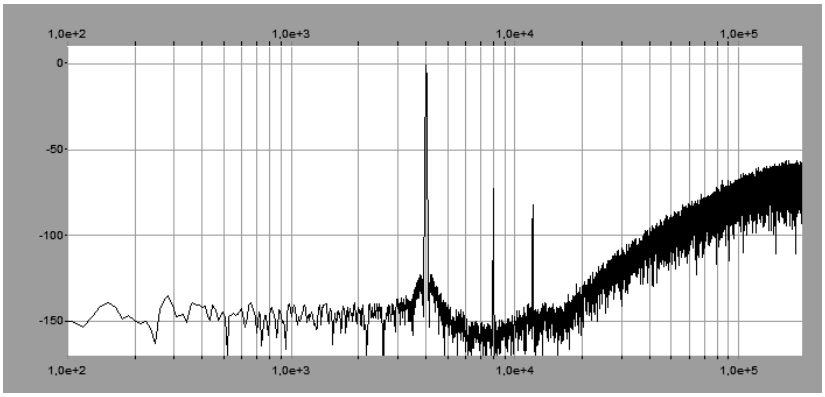

Plot 3: Output spectrum of circuit from figure 1 using a double-sided modulator

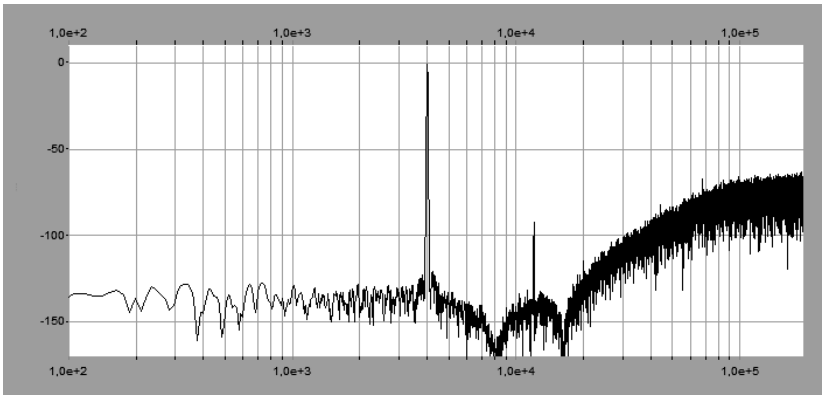

Plot 4: output spectrum from a PWM noise shaper as in figure 4 with a 4kHz -1dB stimulus.

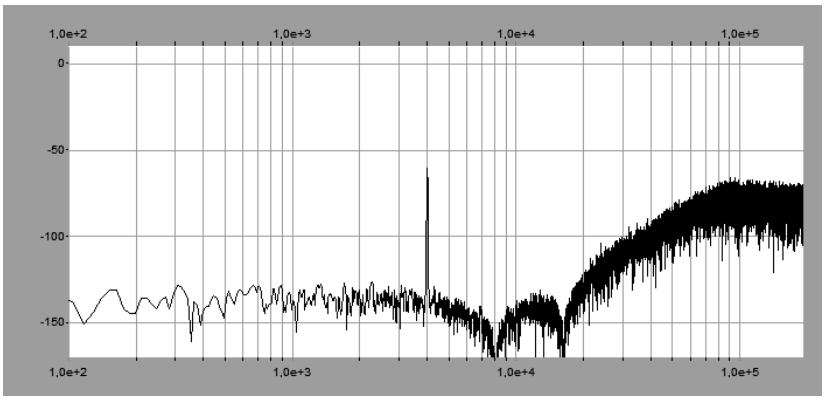

Plot 5: Same as plot 4, but with a -61dB stimulus applied.

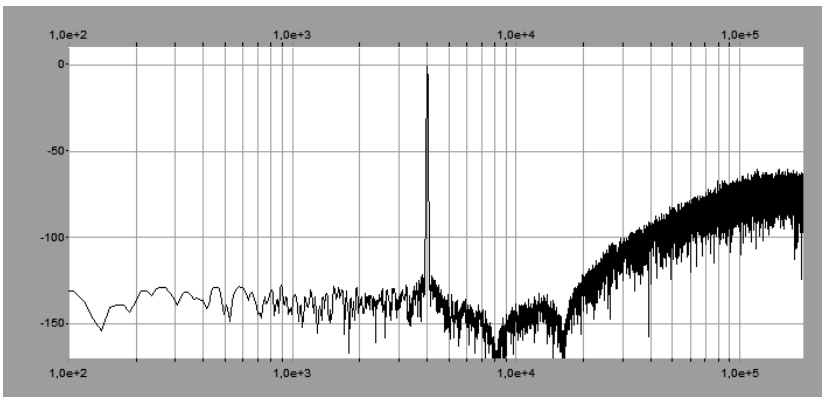

Plot 6: Output spectrum of figure 8 with a -1dB stimulus applied

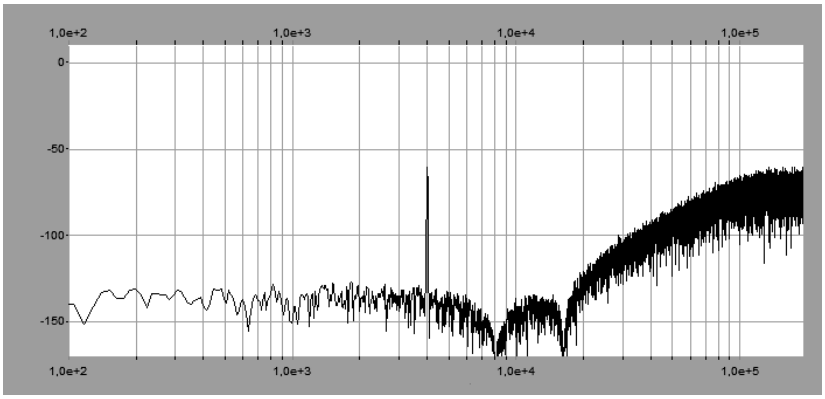

Plot 7: Same as plot 6 with -61dB stimulus.

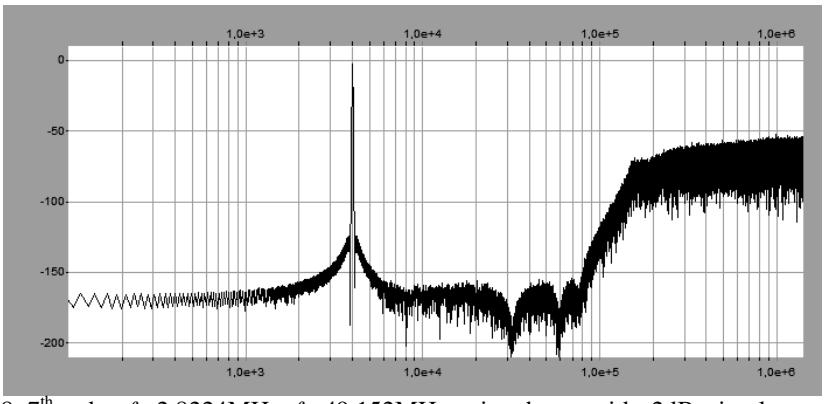

Plot 8: 7th order, f_r=2.8224MHz, f_c=49.152MHz noise shaper with -2dB stimulus applied

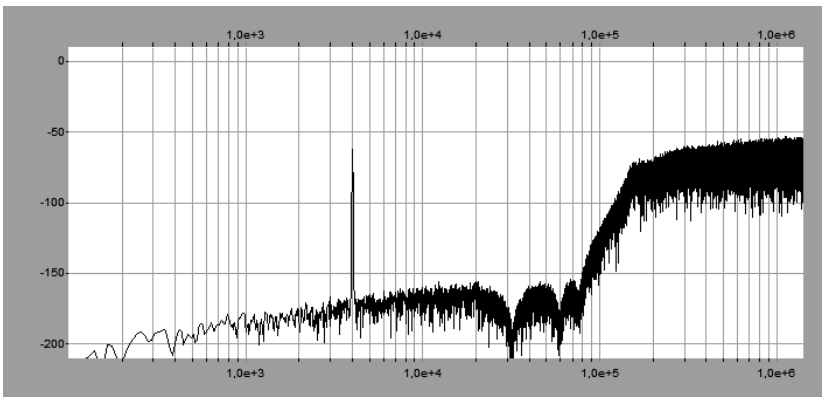

Plot 9: Same as plot 8, -62dB stimulus applied.

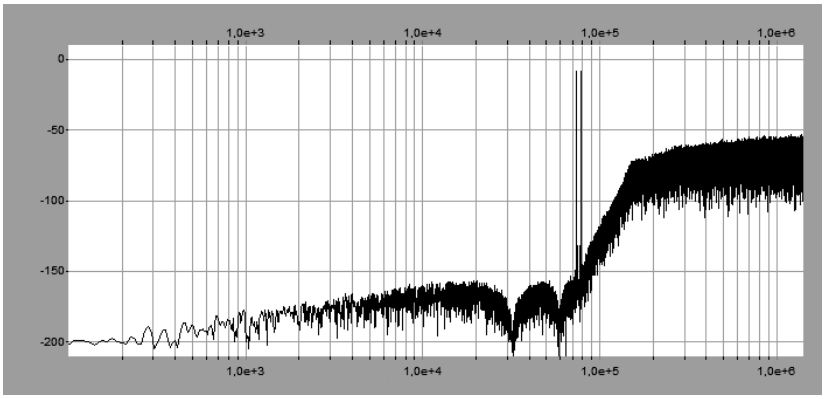

Plot 10: Same as plot 8, 74kHz/-8dB + 78kHz/-8dB stimulus

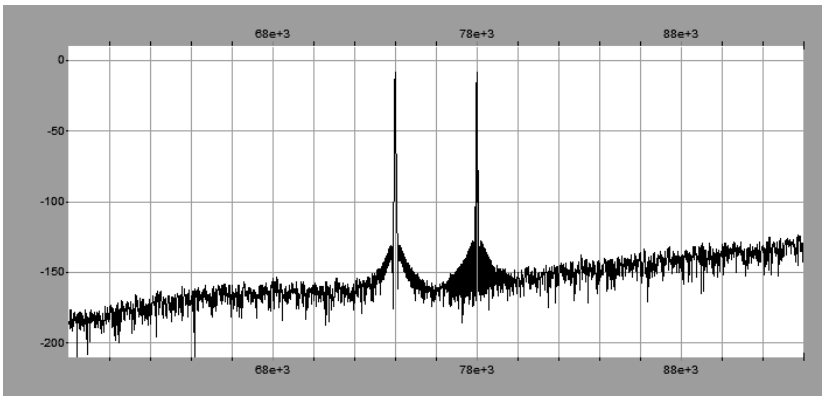

Plot 11: Zoom of plot 10