Audio Engineering Society

Convention Paper

Presented at the 116th Convention
2004 May 8–11 Berlin, Germany

The harmonic content of a limit cycle in a DSD bitstream

Joshua D. Reiss and Mark B. Sandler

[1] Queen Mary, University of London, Mile End Road, London E14NS, United Kingdom
josh.reiss,mark.sandler@elec.qmul.ac.uk

ABSTRACT

This work explores the effects of limit cycles on the frequency content in the DSD bitstream. We show how any periodic bitstream can be expressed as a sum of square waves of various phases with width equal to the sampling period. A Fourier expansion may be used to exactly determine the phases and amplitudes of all spectral content. We thus determine all harmonics that appear in the output, and thus are able to distinguish limit cycles from idle tones. These results are put into the context of recent advances in the theory of limit cycles and idle tones in sigma delta modulators.

1. INTRODUCTION

A variety of different definitions have been proposed for limit cycles and idle tones (see for example [1-3]). It is important to make a distinction between these two phenomena. Limit cycles refer to a behaviour that is easily seen in the time domain, whereas idle tones represent a frequency domain phenomenon. Limit cycles actually have a formal mathematical definition, whereas the meaning of idle tones can be derived from the dictionary definitions of 'idle' and 'tone.'

According to dynamical systems theory, limit cycles refer to isolated periodic orbits in a dynamical system. In a sigma delta modulator, the existence of limit cycles refers to a repeating sequence of integrator states and quantization output that can occur for a given input. [4] puts a description of limit cycles in sigma delta modulators within the proper mathematical context. In practice however, only the quantizer output is monitored. So a repeating output sequence is commonly referred to as a limit cycle.

Idle tones, on the other hand, represent a strong periodicity that is observed in the frequency domain. Specifically, it is the occurrence of peaks in the power spectrum, imposed on a noisy background, that can occur with constant input, i.e., while the system is 'idling.' Idle tones need not be due to the existence of a limit cycle. For instance, idle tones may appear and disappear- behaviour that is not explained by simple periodic motion.

This distinction between idle tones and limit cycles raises the question of how idle tones might be related to limit cycles. For instance, one may ask if a limit cycle (time-domain periodicity) may produce an idle tone (discrete peak in a power spectrum) or idle tone-like behaviour.

For this reason we examine the harmonic content of limit cycles that can occur in the DSD bitstream. We investigate if its possible for a limit cycle to produce a single (or very few) peaks in the spectrum. Furthermore, we calculate how many different output spectra may occur for a given period limit cycle. We are thus able to show that observed idle tones are not the direct product of limit cycle behaviour. Finally, we find a very high period limit cycle that exists in a 5th order sigma delta modulator, of the type that may be used to encode DSD data. We investigate its frequency content and show that it may produce audible artifacts.

2. COUNTING LIMIT CYCLES

In order to count limit cycles, we note that a limit cycle of period P may appear in one of P cyclically shifted forms, e.g., the limit cycle $\overline{001}$ may appear as $\overline{001}, \overline{010}$ or $\overline{100}$. Furthermore, a limit cycle observed over a P bit sequence may actually be a repetition of a shorter period limit cycle, e.g. the apparent limit cycle $\overline{001001}$ is actually two repetitions of the limit cycle $\overline{001}$. This distinction is important since it prevents us from counting limit cycles multiple times, and it guarantees that each limit cycle I associated with the correct frequency. Thus, we use the following notation:

$BSL_P(N)$ represents the number of bit sequences of length P containing N 1s, whereas $BSP_P(N)$ represents the number of bit sequences with period P containing N 1s. Similarly, $LCL_P(N)$ represents the number of limit cycles observed over a length P containing N 1s, whereas $LCP_P(N)$ represents the number of limit cycles with period P containing N 1s. We drop the (N) to represent all bit sequences or limit cycles for a given P.

Since the quantiser only has two levels,

$$BSL_P = 2^P \tag{1}$$

To compute the total number of bit sequences of a given period, we have to take into account that certain sequences are actually of lower period. Over a length P, a possible bit sequence is P/R repetitions of a sequence of length R, where R divides P. So, the bit combinations over a given length can be written as a sum of bit combinations of certain periods[4],

$$BSL_P = \sum_{R \uparrow P} BSP_R \tag{2}$$

Therefore,

$$BSP_P = 2^P - \sum_{R \uparrow P,\, R \neq P} BSP_R \tag{3}$$

and this formula can be computed iteratively.

A limit cycle of period P accounts for P sequences,

$$LCP_P = BSP_P / P \tag{4}$$

Furthermore, all limit cycles that can exist over a given length P includes all limit cycles with a period which is a divisor of P.

$$LCL_P = \sum_{R \uparrow P} LCP_R \tag{5}$$

The number of allowable limit cycles is further reduced if the input is known. For a sigma delta modulator with atleast one pole at DC, the average output over 1 period must equal the input. Thus we can determine allowable sequences and limit cycles for a given input. $BSL_P(N)$, the number of bit sequences of length P containing N 1s, is the number of ways of choosing N out of P objects,

$$BSL_P(N) = \binom{P}{N} \tag{6}$$

This is confirmed since,

$$\sum_{N=0}^{P} BSL_P(N) = \sum_{N=0}^{P} \binom{P}{N} = 2^P = BSL_P \tag{7}$$

Now consider a bit sequence with period R, where R divides P. In order for this bit sequence to be included in $BSL_P(N)$, Eq. (6), it must have NR/P 1s, such that over P bits, it has N 1s. Therefore, we can write

$$\binom{P}{N} = BSL_P(N) = \sum_{R \uparrow P \uparrow NR} BSP_R(NR/P) \tag{8}$$

and hence

$$BSP_P(N) = \binom{P}{N} - \sum_{R\uparrow P\uparrow NR, R\neq P} BSP_R(NR/P) \quad (9)$$

Finally, the number of limit cycles with N 1s can be easily determined from Eq. (9),

$$LCP_P(N) = BSP_P(N)/P \quad (10)$$

and

$$LCL_P(N) = \sum_{R\uparrow P\uparrow NR} LCP_R(NR/P) \quad (11)$$

Clearly, the number of sequences over a length P bitstream is dominated by those sequences which are of period *P*. Thus we have

$$LCP_P \sim 2^P/P \sim LCL_P \quad (12)$$

and, for $0 \ll N \ll P$,

$$LCP_P(N) \sim \binom{P}{N}/P \sim LCL_P(N) \quad (13)$$

P	LCL_P	$2^P/P$	LCP_P	$LCL_P(P/2)$	$\binom{P}{P/2}/P$	$LCP_P(P/2)$
2	3	2	1	1	1	1
4	6	4	3	2	1.5	1
6	14	10.7	9	4	3.3	3
8	36	32	30	10	8.7	8
10	108	102.4	99	26	25.2	25
12	352	341.3	335	80	77	75
14	1182	1170	1161	246	245.1	245
16	4116	4096	4080	810	804.4	800
18	14602	14563.6	14532	2704	2701.1	2700
20	52488	52428.8	52377	9252	9237.8	9225

Table 1. The total number of limit cycles occurring over a length P (LCL_P), of period P (LCP_P), and limit cycles over length P or of period P containing P/2 1s, along with estimates from Eq.s (12) and (13).

From Eq.s (12) and (13), we see that even though the number of limit cycles which exist for a given DC input is an ever decreasing proportion of the total number of allowable outputs, it still increases at an exponential rate. Thus, an examination of all limit cycles which might exist for a given input becomes very computationally intensive for low frequency limit cycles within the audible range, i.e., period exceeding 128 bits.

These results are confirmed in Table 1, where we calculate the number of limit cycles existing under various circumstances. It confirms the validity of Eq.s (12) and (13) as approximations for the number of limit cycles. Furthermore, it demonstrates the exponential growth in limit cycles as a function of period.

3. HARMONIC CONTENT OF A LIMIT CYCLE

The preceding section provides a mechanism for counting all allowable limit cycles of a given period for a constant input. In the context of distinguishing idle tones from limit cycles, this is useful because we can now determine the proportion of limit cycles which produce idle tone behaviour, if any.

However, this requires that we determine the frequency content of a limit cycle. Most importantly, it is necessary to determine when the frequency content might yield a single or select few peaks in the power spectrum. For this, we will begin by performing a standard Fourier series expansion of a pulse train. Then we will show how that may be used to determine the Fourier series for any periodic bitstream, and use this to determine when harmonic cancellation might result in an idle-tone.

3.1. Periodic Pulse Harmonics

Consider an arbitrary square wave like the one shown below

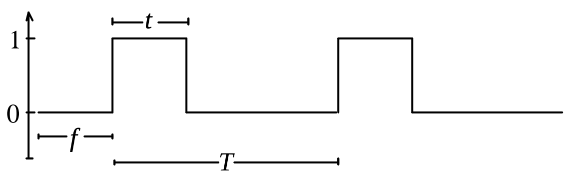

Here, T is the period or pulse rate (w=2p/T), t is the pulse width and f is the phase. T, t and f are all multiples of the sampling rate, and hence may be normalised to integer values. The DC approximation is given by the duty cycle, $d=t/T$. A square wave is thus given by $X(t)=$ square$(f(t-f))$, where

$$\text{square}(n) = \begin{cases} 1 & n - \lfloor n \rfloor < d \\ 0 & n - \lfloor n \rfloor \geq d \end{cases} \quad (14)$$

A Fourier series expansion results in .

$$x(t) = \sum_{n=-\infty}^{\infty} c_n e^{jnwt} \tag{15}$$

where

$$c_n = d\,\mathrm{sinc}(npd)e^{-jnw(t/2+f)}, \quad n = 0, \pm 1, \pm 2\ldots \tag{16}$$

And the $n=0$ term represents the DC component,

$$c_0 = \frac{1}{T}\int_0^T x(t)\,dt = d \tag{17}$$

If we set $f=0$ and $d=1/2$ ($t=T/2$), this reduces to the familiar expression for a Fourier series expansion of a square wave,

$$x(t) = \frac{1}{2} + \frac{2\sin(wt)}{p} + \frac{2\sin(3wt)}{3p} + \frac{2\sin(5wt)}{5p}\ldots \tag{18}$$

It says, among other things, that a square wave with a 1/2 duty cycle does not have even harmonics and that the strength of the remaining odd harmonics is inversely proportional to harmonic number.

The $n=1$ component has the same frequency as the periodic signal, and is the fundamental frequency. Further harmonics arrive at a frequency spacing of w. A closer examination of Eq. (16) reveals that the sinc(npd) term is responsible for suppressing harmonics. To suppress the n^{th} harmonic, d must be a multiple of $1/n$. Thus, for a square wave with $d=1/2$, all even harmonics are suppressed.

If we want to suppress the n^{th} harmonic, the pulse width can be m/n times the period, where $m<n$. Conversely, the amplitude first becomes zero at frequency $nw = 2p/t$.

When T is increased there is no change in the position of the first point at which the amplitude spectrum hits 0. The general form of the spectrum remains the same, as given by sinc x. The number of harmonics up to the first zero amplitude harmonic is increased.

3.2. Two Pulse Harmonic Cancellation

We are concerned with how a periodic bitstream devoid of certain frequencies may be devised. From the above section, it is clear that only even harmonics are removed from a periodic bitstream with a single pulse. Thus we introduce additional pulses in order to achieve harmonic cancellation.

We suppose the period T, is divided into M equal portions and that there are two pulses of width T/M, with discrete phases that are multiples of $1/M$. The phase of a pulse, f, depends on its quantized position. The Fourier series for a pulse in the m^{th} position is thus

$$c_{n,m} = \mathrm{sinc}(np/M)e^{-jnp(1+2m)/M}/M \tag{19}$$

Magnitude at a given harmonic depends only on the harmonic number n, and not on the pulse position, m. Phase depends on both harmonic and pulse position.

Now consider how two pulses, m_1 and m_2, interact with one another.

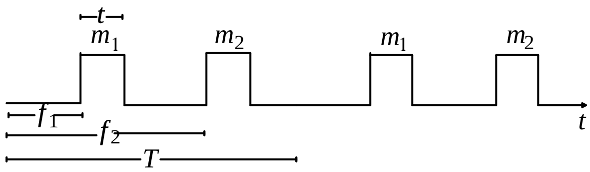

Figure 1 An arbitrary two pulse periodic bitstream.

To obtain the Fourier series of the two-pulse pulse train we only need to add together the Fourier series obtained from the pulses when they're taken individually, $c_n = c_{n,m_1} + c_{n,m_2} \quad n = \pm 1, \pm 2, \pm 3,\ldots$

And the DC component is

$$c_0 = 2t/T$$

So the n^{th} Fourier term of the composite wave is the sum of the n^{th} terms of the m_1 and m_2 pulses.

$$c_n = \frac{\mathrm{sinc}(np/M)[e^{-jn2pm_1/M} + e^{-jn2pm_2/M}]}{Me^{jnp/M}} \tag{20}$$

The term outside of the brackets in Eq. (20) depends only on the harmonic number and not position or phase of the pulse. As seen earlier, this can be made zero only by varying the pulse width. However the pulse width has been fixed to T/M. Consequently we need to study the term in the brackets if we wish to investigate harmonic cancellation.

In general, we can induce phase cancellation when the term in the brackets in Eq. (20) is zero. We assume an ordering to the two pulses, so that $m_1 < m_2$. So,

$$n(m_2 - m_1) = \frac{M}{2}, \frac{3M}{2}, \frac{5M}{2} \ldots \frac{(2n-1)M}{2} \qquad (21)$$

The first thing to notice is that the left hand side of (21) is an integer. So there will never be phase cancellation unless M is even. Odd periods will not produce phase cancellation (and hence remove harmonics) from any two pulses.

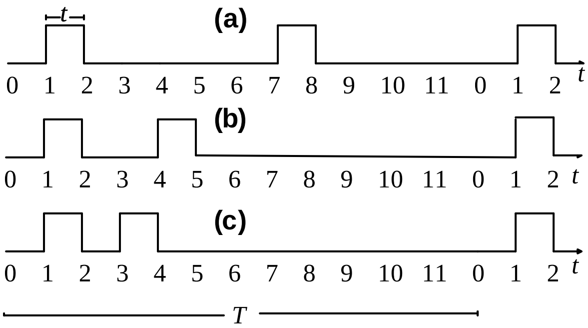

Figure 2 Harmonic cancellation with two pulses in a twelve pulse periodic bitstream. (a) Pulses at positions 1 and 7, (b) pulses at positions 1 and 4, (c) pulses at 1 and 3.

Harmonic cancellation with two pulses, for M=12, is depicted in Figure 2. Pulses at positions 1 and 7 cancels the fundamental. This is identical to creating a new square wave with twice the frequency of the first. Pulses at positions 1 and 4 suppresses the 2^{nd} harmonic. The fundamental and third harmonics are $\sqrt{2}$ times larger than they are for just a single pulse. Pulses at positions 1 and 3 suppresses the 3^{rd} harmonic.

This can be generalised: the n^{th} harmonic can only be cancelled if M is a multiple of $2n$. In general, the n^{th} harmonic is cancelled if the phase difference, $f_2 - f_1$, is a multiple of $T/2n$. Furthermore, any phase difference that cancels the n^{th} harmonic, also cancels the $(3n)^{th}$, $(5n)^{th}$… harmonics. This is because

$$\frac{M}{2n} = \frac{3M}{2(3n)} = \frac{5M}{2(5n)} \ldots \qquad (22)$$

The same does not hold in reverse, of course.

One other important point to notice is that there is no way to cancel any two of the 2^{nd}, 3^{rd}, 4^{th} and 5^{th} harmonics using just two pulses. An easy way to see this is by looking at all the possible phase differences between the two pulses that will allow phase cancellation for these harmonics.

$$\begin{aligned} n = 2 &: \Delta m = \frac{M}{4}, \frac{3M}{4} \\ n = 3 &: \Delta m = \frac{M}{6}, \frac{3M}{6}, \frac{5M}{6} \\ n = 4 &: \Delta m = \frac{M}{8}, \frac{3M}{8}, \frac{5M}{8}, \frac{7M}{8} \\ n = 5 &: \Delta m = \frac{M}{10}, \frac{3M}{10}, \frac{5M}{10}, \frac{7M}{10}, \frac{9M}{10} \end{aligned} \qquad (23)$$

The only terms that are the same are $3M/6$ for n=3 and $5M/10$ for n=5. But these occur when the phase difference is $M/2$, i.e., for a system with one pulse at twice the frequency. Obviously, doubling the frequency removes all odd harmonics.

3.3. Generalised periodic bitstream

Now suppose, that instead of two pulses, we have a set of pulses, $m_1, m_2, \ldots m_k$ where $k<M$ and the pulses have been ordered $m_1<m_2<\ldots m_k$. This represents a periodic bitstream. For instance, the repeating limit cycle, 00110110 can be represented by M=8, k=4, and m_1=2, m_2=3, m_3=5, m_4=6.

We can now give the Fourier series expansion of the bitstream.

$$c_n = \mathrm{sinc}(np/M)e^{-jnp/M}\sum_{i=1}^{k} e^{-jn2pm_k/M}/M \qquad (24)$$

and the DC component is given by

$$c_0 = k/M \qquad (25)$$

Harmonic cancellation occurs whenever the summation is zero. The rules for harmonic cancellation become far more complicated as the number of 1s in a period is increased. However, the overwhelming majority of harmonics remain. For instance, of the 75 limit cycles of period 12 with 6 1s, only 29 of them experience any harmonic cancellation. Of those 29 limit cycles, only 7 are missing more than 1 of the 5 harmonics below half the sampling frequency (64*44.1/2=1411.2kHz).

If we consider the DSD format, then the sampling frequency is 64*44.1kHz. Thus a period M bitstream has frequency 64*44.1/M kHz, and the harmonics occur at integer multiples of that. Therefore, a limit cycle must have a period of at least 128 to be audible (ignoring the post-filtering effects).

We can also identify some general trends for periodic bitstreams. If we increase the period, but don't add more pulses, then the spectral structure remains the same. Consider the periodic bitstream 11010000, which repeats with period 8. As depicted in Figure 3, all the harmonics of 64*44.1/8=352.8 kHz are present.

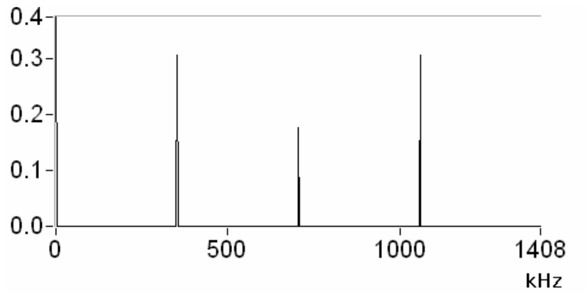

Figure 3 Power spectrum for period 8 bitstream 11010000. Frequency on the *x*-axis is in kHz for a DSD signal (64×44.1 1 samples per millisecond).

If we increase the period to 32, so that the sequence 11010000000000000000000000000000 repeats, then all the harmonics of 64*44.1/32=88.2 kHz are present (Figure 4).

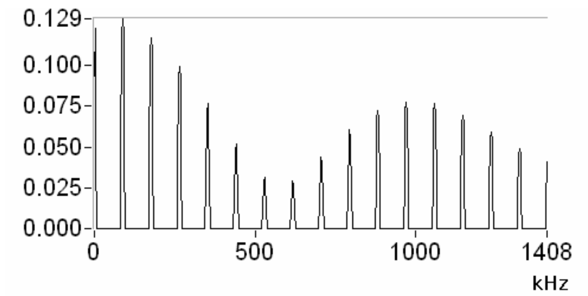

Figure 4 Power spectrum for period 32 bitstream 1101000…000.

Eventually, this approaches a continuous spectrum whose envelope is determined by the sinc term, as in Figure 5. In general, we can suppress any given harmonic by varying the width of the pulse. Notice that unlike with a single pulse, or even two pulses, in this situation, none of the harmonics are ever cancelled.

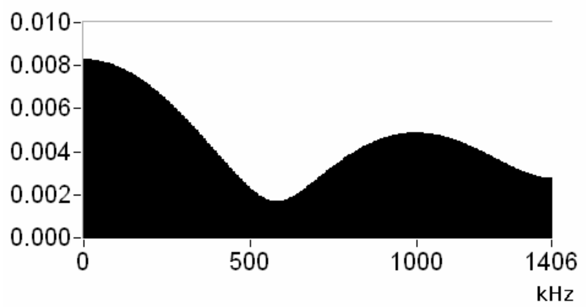

Figure 5. The power spectrum of a continuously sampled periodic square wave with t/T approaching infinity.

4. CONCLUSION

In this work we identified and classified the frequency content of limit cycles that might occur in the DSD bitstream. Equations (1) to (13) can be used to count all limit cycles that might occur for any constant input, and the formulas for their approximation.

We then considered the coefficients of the Fourier expansion of a periodic pulse train, and thus derived an expression for the Fourier expansion of any DSD limit cycle where the phases and amplitudes of each frequency can be determined exactly, Eq.(24).

The first harmonic occurs at the frequency of the limit cycle, and all subsequent harmonics are at integer multiples of that. Therefore a periodic signal must have period greater than 128 for a peak to appear in the audible range (< 22.05 kHz).

Phase cancellation was considered for the one and two pulse situations, and then generalized to any number of pulses (bits set to one) in the limit cycle. It was shown that harmonic cancellation is unlikely, and rarely would it result in many harmonics being removed. In fact, it is easy to devise periodic bitstreams where all the harmonics have a positive amplitude (Figure 5).

Thus, it is wrong to think of a limit cycle as representing a single peak in the spectrum for a sigma delta modulator, since in general, the harmonics will be present. So, the arguments presented here reinforce the concept that idle tones should be treated differently from limit cycles. Recent work on limit cycles in the DSD bitstream does not resolve the issues concerning idle tones, and vice versa.

5. REFERENCES

[1] D. Reefman and E. Janssen, "Signal processing for Direct Stream Digital: A tutorial for digital Sigma Delta modulation and 1-bit digital audio processing," Philips Research, Eindhoven, White Paper 18 December 2002

[2] J. A. S. Angus, "The Effect of Idle Tone Structure on Effective Dither in Delta-Sigma Modulation Systems: Part 2," presented at Audio Engineering Society 113th Convention, Los Angeles, California, 2002.

[3] S. P. Lipshitz and J. Vanderkooy, "Towards a Better Understanding of 1-Bit Sigma-Delta Modulators," presented at Audio Engineering Society 110th Convention, Amsterdam, Holland, 2001.

[4] L. Risbo, "Sigma-Delta Modulators - Stability Analysis and Optimization," *Electronics Institute*. Lyngby: Technical University of Denmark, 1994, pp. 179. http://eivind.imm.dtu.dk/publications/phdthesis.html