Linearising Sigma-Delta Modulators using Dither and Chaos

Chris Dunn and Mark Sandler
King's College London
email: chris.dunn@scalatech.co.uk

Abstract - Recent work has shown that high-order single-bit sigma-delta modulators suffer from low-level artifacts such as idle tones and noise modulation. Techniques that have been proposed to reduce or eliminate these errors include the application of dither inside the one-bit quantiser loop, and selecting a loop filter which makes the modulator chaotic. This paper compares the efficacy of these two approaches by simulating high-resolution sigma-delta modulators suitable for audio-conversion applications. Dynamic-range penalties for successful linearisation are determined for two types of dither signal and two classes of chaos.

I. INTRODUCTION

Analogue-to-digital converters (ADCs) and digital-to-analogue converters (DACs) used to code audio signals achieve minimal signal degradation in a psychoacoustic sense when the noise floor is invariant with input signal characteristics [1] (up to the point of overload). Hence noise modulation and distortion are both forms of nonlinearity which are undesirable and, if possible, avoided. While sigma-delta modulators (SDMs) can exhibit excellent linearity for large-amplitude signals, idle-tones can corrupt the output from the modulator when lower-amplitude signals are coded. Techniques for eliminating idle tones in SDMs include the use of dither, and making the converter chaotic. In this paper we focus upon the dynamic-range penalties associated with successfully implementing these approaches. The study commences in Section 2 with a discussion of SDM idle-tone phenomena, while Sections 3 and 4 examine dithered and chaotic sigma-delta modulators respectively.

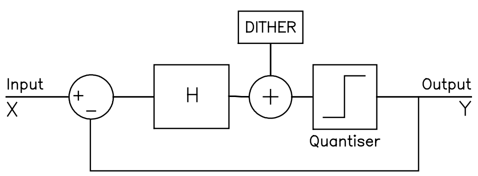

Fig. 1.General sigma-delta modulator.

II. SDM IDLE TONES

Fig. 1 shows a block diagram of an SDM, where a single-bit quantiser and an optional dither noise-source are embedded in a negative feedback loop including loop filter *H*. The signal-to-noise ratio (SNR) achieved in the signal band is dependent upon the noise-shaping function of the modulator, which can be described in terms of *z*-domain poles p_i and zeros z_i. For an *n*th-order modulator,

$$NS(z) = \frac{\prod_{i=1}^{n}(z - z_i)}{\prod_{i=1}^{n}(z - p_i)} = \frac{1}{1 + H(z)} . \qquad (1)$$

A popular choice for the noise-shaping pole locations is to arrange them in a Butterworth configuration, while high SNRs are achieved if noise-shaping zeros are distributed across the baseband in conjugate pairs [2]. The resolution increases as the Butterworth cutoff frequency is increased, although such an action will also tend to reduce stability. In general there will be an optimal cutoff frequency which yields the highest dynamic range obtainable from the modulator. Schreier [3] describes a technique for empirically determining this value by simulation, and a similar optimisation procedure - described in detail in [4] - was implemented for all of the modulators examined in this study.

Although it has been known for some time that low-order SDMs suffer from noise modulation [5], until recently it was generally believed that higher-order systems (n > 2) did not suffer from such nonlinear artifacts [2]. Indeed, it is easy to show that at high sinusoidal-input amplitudes, the output spectrum from a higher-order SDM is virtually free of distortion tones, and results such as these have led many to believe that higher-order systems behave as virtually ideal converters. However, consider Fig. 2 which shows a simulated output spectrum for a fourth-order modulator with a 1 kHz sinusoidal input at -47 dBFS, where the magnitude response has been averaged over 100 frames;

low-level tones can be seen rising above the noise floor at several frequencies. As for all of the simulation results presented in this paper, the oversampling factor *OSF* is set to 64, and the sampling frequency f_s = 2.8 MHz - this yields a baseband of 0 kHz to f_B = 22 kHz.

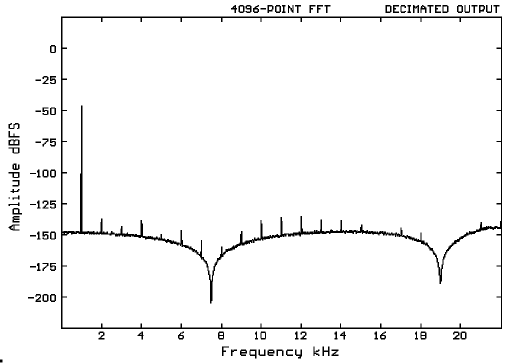

Fig. 2. Output from 4th-order SDM with 1 kHz -47 dBFS input.

Ledzius and Irwin [6] have shown that, for constant inputs, idle-tone frequencies are related to the modulator input amplitude x_{DC}. For an SDM with quantiser output levels ±1, an important idle-tone occurs at a frequency given by

$$f_{IT} = x_{DC}\, f_S = 2\, x_{DC}\, f_B\, OSF\,. \qquad (2)$$

Fig. 3(a) shows how the noise floor of a fourth-order modulator changes as a dc input sweeps across the range 0 -> 1/256, clearly indicating idle tones with frequencies f_{IT} and $2f_{IT}$. In the following sections we examine the techniques of dither and chaos as strategies for linearising sigma-delta modulation.

III. DITHERED SIGMA-DELTA MODULATORS

In a previous study of dithered SDMs, Norsworthy and Rich [7] describe how a pseudo-random dither signal added to the input of the single-bit quantiser (Fig. 1) can break up idle-tones in the quantisation noise floor, and is conveniently noise-shaped by the sigma-delta loop. Nevertheless, the introduction of dither reduces the dynamic range available from the modulator, partly because of the increase in total noise power within the loop, but also because the addition of dither tends to reduce modulator stability. For the present study, an extensive series of simulations was performed to determine the peak-SNR penalty for dithered modulators of orders ranging from 1 to 4, and two types of dither signal: rectangular probability distribution (RPD), and triangular probability distribution (TPD). For each class of dither signal, and for a range of dither amplitudes in approximate steps of 3 dB, modulators were optimised to yield maximum dynamic range. A dithered modulator was considered to be "linearised" if no idle tones were visible in the dc-input sweep plot, and noise modulation was less than 1 dB for sinusoidal input signal amplitudes within the dynamic range of the modulator. An example of a linearised modulator is shown in Fig. 3(b) where a dc input sweep is applied to a fourth-order system dithered using ± 0.3 RPD pseudo-random noise. No idle tones are apparent, although, compared to the undithered modulator, this linearisation has been achieved at the expense of a decrease in dynamic range of 6.7 dB.

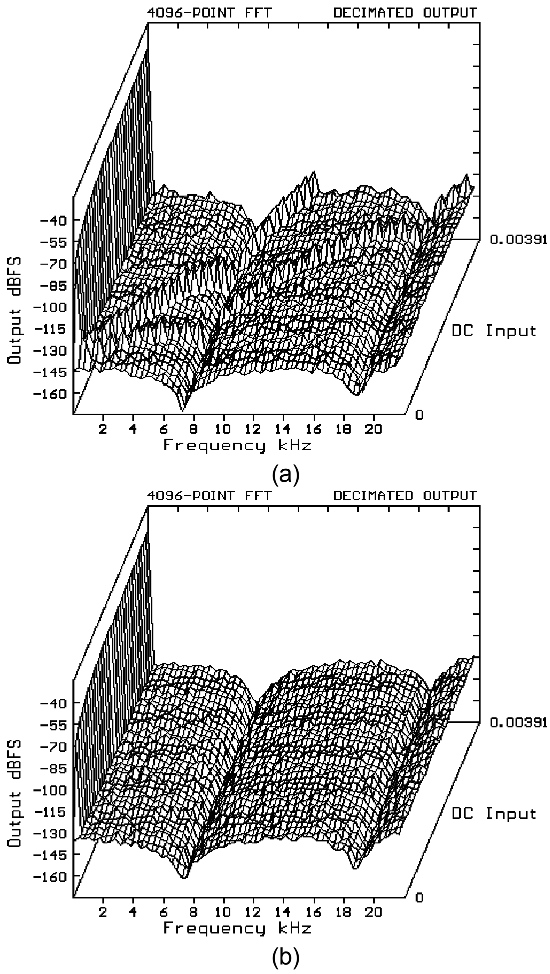

Fig. 3. DC-input sweep plots for 4th-order modulators. (a) Standard. (b) Dithered.

The simulation results collated in Table 1 represent, for each loop-filter order and dither class, the lowest-amplitude dither signal which successfully linearises the modulator. Several results can be obtained from this data:

- The peak SNRs for modulators optimised without dither agree quite well with Schreier's results [3], and hence verify the accuracy of the optimisation process used in the simulations.
- With the exception of the first-order system, RPD dither spanning ± 0.5 linearised all modulator orders - a result which agrees with Norsworthy's findings [7]. However, the peak SNR penalty of 6.3 dB for dithering the second-order modulator is significantly higher than reported by Norsworthy.
- Included in the Table for each dither setting is the dither power relative to ±0.5 TPD, data which indicates that the dither power required for linearisation decreases as *n* increases.
- For second-, third-, and fourth-order systems, it is seen that the dynamic-range penalty for implementing dither remains approximately constant at 6 dB. Risbo has reported that the SNR penalty for dithering a sixth-order system is as high as 27 dB [8].

Table 1. Dynamic-range penalties for dithered SDMs.

Order n	SNR_{max} standard modulator (dB)	Dither	Dither Power	SNR_{max} dithered modulator (dB)	Dynamic range penalty (dB)
1	58.6	RPD±2.0	8	42.1	16.5
		TPD±1.0	4	47.4	11.2
2	75.8	RPD±0.5	2	69.5	6.3
		TPD±0.7	2	68.6	7.2
3	93.4	RPD±0.3	0.72	88.1	5.3
		TPD±0.5	1	87.4	6.0
4	109.0	RPD±0.3	0.72	102.3	6.7
		TPD±0.5	1	99.9	9.1

IV. CHAOTIC SIGMA-DELTA MODULATORS

An alterative to using dither to eliminate low-level SDM artifacts is to make the modulator *chaotic*, where noise-shaping zeros are moved outside the unit circle in the *z*-domain. Schreier [9] notes that this arrangement is equivalent to making $H(z)$ open-loop unstable, and will tend to disrupt limit cycles. Although the output of a chaotic system is generally non-periodic, with a continuous spectrum, this condition does not preclude the *combination* of tones and noise in the modulator quantisation error [9]. Hein [10] recently described how introduction of chaos to a standard second-order modulator (noise-shaping zeros at dc) can break up idle tones occurring for dc input signals. Our simulations suggest that such behaviour does indeed occur for some dc input levels, but for other input levels chaos is less successful in breaking up tonal components of the output signal. The standard (non-chaotic) second-order system was simulated with x_{DC} = 1/256, revealing strongly-periodic quantisation noise with 7 baseband idle tones - including f_{IT} = 11 kHz. Both noise-shaping zeros were then moved outside the unit circle by setting the noise-shaping zero radii r_z = 1.01, and the system resimulated for the same input signal. The introduction of chaos makes the output spectrum less tonal (Fig. 4), although the amplitude of the f_{IT} idle tone remains approximately unchanged. Idle-tone suppression relative to the random noise-floor component increases as the noise-shaping zeros are moved further outside the unit circle, although such an action also tends to reduce the baseband suppression of quantisation noise, with a consequent reduction in SNR. In fact the introduction of moderate degrees of chaos (r_z ~ 1.01) can have the surprising effect of *increasing* the absolute amplitude of baseband idle tones. Further simulation examples presented in [4] indicate that moderate degrees of chaos are also only partially successful at linearising lower-order SDMs at low signal amplitudes.

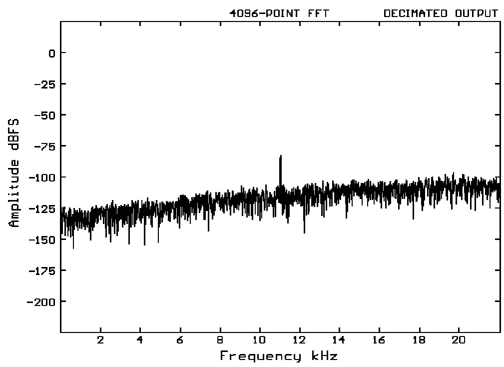

Fig. 4. Output spectrum from chaotic second-order modulator (r_z = 1.01) with dc input of 1/256.

In order to determine whether higher levels of chaos can successfully linearise SDMs, an extensive series of simulations was performed on modulators with orders ranging from 1 to 4. Two classes of chaotic system were investigated:

- *Scaled-zero*, where z_i for a standard non-chaotic system are all moved to lie outside the unit circle, by scaling r_z from unity while zero frequencies remain unchanged. Fig. 5(a) shows noise-shaping pole and zero placement for a fourth-order system with $r_z = 1.1$.
- *Allpass*, where $NS(z)$ is designed with a single allpass section located at $f_s/2$, while remaining noise-shaping poles are arranged in a Butterworth configuration and zeros set to optimal positions [8]. Fig. 5(b) shows a fourth-order modulator with the allpass zero radius $r_z = 1.1$.

The degree of chaos ~ (r_z-1) ranged from 0.0005 to 0.5 in approximate steps of 3 dB. As for the introduction of dither, implementing chaos tends to deteriorate the stability of sigma-delta modulators [4], hence each chaotic system investigated was optimised for maximum dynamic range. Using the linearisation criteria adopted in the dither experiments, the results shown in Table 2 represent, for each modulator order, the minimum degree of chaos required to successfully eliminate unwanted artifacts. None of the first-order chaotic modulators investigated were completely free of idle tones. For second- and higher-order scaled-zero systems, complete linearisation was only achieved with noise-shaping zeros well outside the unit circle, r_z >= 1.1, and hence suffer from large dynamic-range penalties - a finding which agrees with Risbo's study of sixth-order systems [8]. Although the value of r_z required for linearisation is higher for the allpass structure - probably because of the smaller number of zeros located outside the unit circle - SNR penalties associated with allpass systems tend to be lower than scaled-zero modulators. Nevertheless, the loss in dynamic range suffered by chaotic systems tends to be much higher than for equivalent dithered systems. For example, the fourth-order allpass chaotic modulator has a peak SNR which, compared to the optimal non-chaotic system, is reduced by 38 dB. This can be compared to a dynamic-range penalty of only 6 dB for successful implementation of dither.

V. CONCLUSIONS

Simulations of dithered SDMs have shown that the dither power required for linearisation tends to reduce as the order of the modulator increases, and for the modulator orders considered, the associated dynamic-range penalties are relatively modest. As well as linearising the in-band quantisation noise floor, an appropriate level of dither also has the effect of reducing high-frequency (>> f_B) idle tones [7], which can be beneficial in practical implementations of high-order SDMs [11]. Finally, it is worth noting that correctly-implemented dither has the ability to eliminate low-level dead zones, so that low-amplitude input signals can be accurately resolved [4]. These factors make dither an attractive prospect in sigma-delta DACs, where the technique is easily implemented.

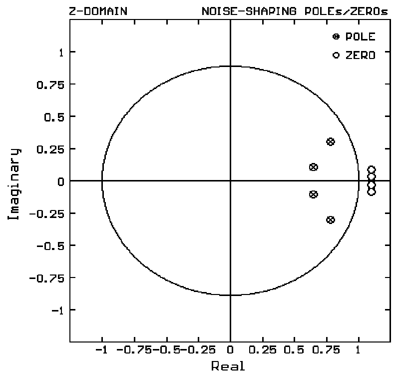

(a)

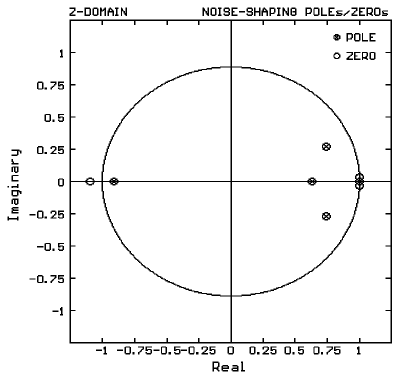

(b)

Fig. 5. Chaotic noise-shaping functions with $r_z = 1.1$. (a) Scaled-zero. (b) Allpass.

While chaos precludes the possibility of quantisation noise which is purely periodic, combinations of tones and noise can occur when the degree of chaos is moderate (r_z ~ 1.01). Relatively-high degrees of chaos (r_z > 1.1) were found necessary to completely eliminate idle tones and noise modulation, although the reduction in signal-to-noise ratio associated with implementing chaos is then high. While allpass structures are more efficient at linearising modulators than scaled-zero systems, in general chaos appears less effective at linearisation compared to dither.

Table 2. Dynamic-range penalties for optimal chaotic SDMs.

Order *n*	SNR_{max} standard modulator (dB)	Chaos class	r_z	SNR_{max} chaotic modulator (dB)	Dynamic range penalty (dB)
2	75.8	Scaled-zero	1.15	38.7	37.1
		Allpass	1.5	28.4	47.4
3	93.4	Scaled-zero	1.1	46.5	46.9
		Allpass	1.3	58.9	34.5
4	109.0	Scaled-zero	1.1	34.9	74.1
		Allpass	1.3	70.8	38.2

It should be stressed that the present study was confined to modulators with a specific noise-shaping function and a fixed oversampling ratio of 64. Different loop filter configurations and oversampling ratios may well yield results which differ from those presented in this paper. Investigations into the use of dither and chaos in modulators across a range of oversampling factors form part of an ongoing study by the authors of psychoacoustically- optimal sigma-delta modulation.

ACKNOWLEDGEMENT

This work has been supported by the Science and Engineering Research Council, UK, with a research grant on "Adaptive Sigma-Delta Modulators for ADCs and DACs," grant no. GR/H 78702.

REFERENCES

[1] C. Dunn, "Psychoacoustic Modelling of Nonlinear Errors," in *Measurement of Nonlinear Errors in Audio Electronics*, Ph. D. Thesis, University of Essex (1994).

[2] K. C. H. Chao, S. Nadeem, W. L. Lee, and C. G. Sodini, "A Higher-Order Topology for Interpolative Modulators for Oversampling A/D Converters," *IEEE Trans. Circuits and Systems.,* vol. CAS-37, pp. 309-318 (1990 Mar).

[3] R. Schreier, "An Empirical Study of High-Order Single-Bit Delta-Sigma Modulators," *IEEE Trans. Circuits and Systems - II: Analog and Digital Signal Processing*, vol. 40, pp. 461-466 (1993 Aug.).

[4] C. Dunn and M. Sandler, "A Simulated Comparison of Dithered and Chaotic Sigma-Delta Modulators," presented at the 97th Convention of the Audio Engineering Society, San Francisco (1994).

[5] J. C. Candy, Y. C. Ching, and D. S. Alexander, "Using Triangularly weighted Interpolation to Get 13-Bit PCM from a Sigma-Delta Modulator," *IEEE Trans. Comm.*, vol. COM-24, pp. 1268-1275 (1976 Nov.).

[6] R. C. Ledzius and J. I. Irwin, "The Basis and Architecture for the Reduction of Tones in a Sigma-Delta DAC," *IEEE Trans. Circuits and Systems - II: Analog and Digital Signal Processing*, vol. 40, pp. 429-439 (1993 July).

[7] S. R. Norsworthy and D. A. Rich, "Idle Tones and Dithering in Sigma-Delta Modulators," presented at the 95th Convention of the Audio Engineering Society, *J. Audio Eng. Soc. (Abstracts)*, vol. 41, p. 1053 (1993 Dec.), preprint 3711.

[8] L. Risbo, "On the Design of Tone-Free Sigma-Delta Modulators," in *Sigma-Delta Modulators - Stability Analysis and Optimisation*, Ph. D. Thesis, Technical University of Denmark (1994).

[9] R. Schreier, "Destabilising Limit Cycles in Delta Sigma Modulators with Chaos," *Proc. 1993 IEEE Int. Symp. on Circuits and Systems*, pp. 1369-1372 (1993 May).

[10] S. Hein, "Exploiting Chaos to Suppress Spurious Tones in General Double-Loop Sigma-Delta Modulators," *IEEE Trans. Circuits and Systems - II: Analog and Digital Signal Processing,* vol. 40, pp. 651-659 (1993 Oct.).

[11] S. Harris, "How to Achieve Optimum Performance from Delta-Sigma A/D and D/A Converters," *J. Audio Eng. Soc.*, vol. 41, pp. 782-790 (1993 Oct.).