

A Comparison of Dithered and Chaotic Sigma-Delta Modulators

Chris Dunn and Mark Sandler

Department of Electronic and Electrical Engineering
King's College London

email: chris.dunn@scalatech.co.uk

ABSTRACT

Recent work has shown that higher-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.

1 INTRODUCTION

A perfect analogue-to-digital (AD) or digital-to-analogue (DA) conversion process would convert a signal from one domain to the other without introducing either linear errors (that is, magnitude and phase-response aberrations within a specified passband) or nonlinear errors (noise and distortion). While linear errors in modern converters can be made to be vanishingly small, quantisation theory tells us that any conversion process must contribute some degree of nonlinear error, hence a fundamental issue which must be addressed when assessing the quality of a conversion process is deciding what represents an acceptable degree of nonlinearity. As far as the psychoacoustic significance of conversion errors in audio-system AD and DA processes is concerned, a hierarchy of different forms of nonlinearity can be made [1]. Beginning with the least-objectionable form of nonlinearity, ADC and DAC nonlinear errors can be broadly classified within the following groups:

- A constant noise floor which is invariant with input signal level and spectral content. This type of error is relatively benign.
- Noise modulation, where the noise floor of the test device changes as the input signal changes. About 1 dB of noise modulation can be detected under critical listening conditions, if the noise floor of the device is above the threshold of audibility at any audio frequency [1].
- Distortion due to a nonlinear input-output transfer function - that is, the type of error which manifests as distortion harmonics of a sinusoidal stimulus. This is the most objectionable type of nonlinearity, especially if the distortion is of high order [1]. Also belonging to this class of nonlinearity are error tones which are not harmonically related to the stimulus frequency, such as idle tones in sigma-delta modulators.

For real-world conversion processes the target performance is a constant noise floor which is invariant with input signal characteristics (up to the point of overload). Hence noise modulation and distortion are both forms of nonlinearity which are undesirable and, if possible, avoided.

It is well known that quantising a signal represented in a pulse-code modulation (PCM) format results in a quantisation error which is spectrally white if the signal input to the quantiser is complex and has an amplitude much higher than the quantisation step size. However, if the input signal amplitude becomes small, or its frequency has a simple relationship to the system sampling frequency, then quantisation distortion results where both harmonic and anharmonic error tones can be seen at the output of the quantiser. Such nonlinear errors can be extremely disturbing from a psychoacoustic perspective. Fortunately the quantisation error can be made benign by the appropriate use of dither before the quantiser (Fig. 1); addition of triangular-probability-density (TPD) dither of peak amplitude equal to the PCM quantisation step size results in a quantisation error whose spectral power density is completely invariant with input signal characteristics [2], [3], [4]. The correctly-dithered quantisation process therefore

does not suffer from either noise modulation or distortion. Fig. 2(a) shows the frequency-domain output of a dithered 16-bit PCM quantiser for a 1 kHz sinusoidal input of peak amplitude -12 dB with respect to full scale (that is, -12 dBFS), indicating a complete absence of distortion tones. This simulation result was obtained by averaging 100 4096-point frames. It can be shown that the noise-floor penalty for introducing TPD dither is 4.8 dB, while for a quantiser with more than a few bits the introduction of dither does not significantly reduce the overload point of the quantiser. A further benefit of dithering a PCM quantiser lies in the ability to recover, without distortion, signals with amplitudes well below the noise floor of the conversion process. This is not possible with an undithered quantiser, where low-level input signals are not linearly coded. An example of resolution below the noise floor in a correctly-dithered PCM quantiser is shown in Fig. 2(b) where a -100 dBFS sinusoid input to the quantiser is clearly resolved in the output spectrum without distortion, even though the input amplitude is well below the (total) noise floor of the quantisation process at -93.3 dBFS.




Fig. 1. Dithered PCM quantiser.

An "ideal" PCM quantisation process, where the quantisation intervals are equally spaced and the correct dither signal is used, therefore introduces a nonlinear error which can be considered relatively benign from a psychoacoustic perspective. Ideal PCM quantisation can be implemented with relative ease in the digital domain - for example, requantisation operations in a digital signal processor (DSP). However, practical multi-bit PCM quantisers which interface with the analogue domain suffer from the effects of finite component tolerances, which tend to introduce nonlinearity to converter transfer characteristics. Such nonlinearity manifests as distortion and noise modulation at the converter output, even if the conversion is correctly dithered [5]. Practical multi-bit PCM-conversion linearity thus deviates from the ideal of a noise floor which is invariant with input-signal characteristics.

One method of reducing the component matching requirements of multi-bit converters is to oversample the signal input to the quantiser, which for a given converter resolution allows fewer quantiser bits. The number of bits can be further reduced by noise-shaping the quantisation error away from the signal band. Oversampling and noise shaping can be viewed as techniques which achieve high resolution by trading quantiser accuracy (lower number of bits) for circuit speed (higher clock frequencies). Ultimately, the processes of oversampling and noise shaping can result in high-resolution converters using only a single-bit quantiser, known as sigma-delta modulators (SDMs). An excellent introduction to sigma-delta techniques is provided by Candy and Temes in [6], while the following discussion must be necessarily brief. Fig. 3(a) shows a block diagram of an SDM, where the single-bit quantiser is embedded in a negative feedback loop including loop filter H .

If the single-bit quantiser is modelled as an additive quantisation noise source E [Fig. 3(b)], then the output can be written in the z -domain as

$$Y(z) = \frac{H(z)}{1 + H(z)} X(z) + \frac{1}{1 + H(z)} E(z). \quad (1)$$




Fig. 2. Output spectra from correctly-dithered ideal PCM quantiser with 1 kHz input at (a) -12 dBFS, indicating complete absence of distortion tones, and (b) -100 dBFS, showing resolution below the noise floor of the quantisation process.

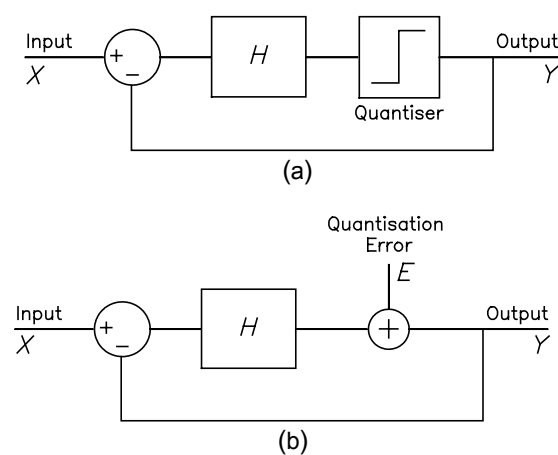


Fig. 3. (a) General sigma-delta modulator. (b) Equivalent quantiser model.

The output has an in-band quantisation noise component which is determined by the noise-shaping function of the modulator,

$$NS(z) = \left. \frac{Y(z)}{E(z)} \right|_{X(z)=0} = \frac{1}{1 + H(z)}. \quad (2)$$

High resolution in the signal band can be achieved for a given oversampling ratio by using a high-order loop filter H . Design strategies for determining H are discussed in Sec. 2.

The advantages of moving to a single-bit quantiser in sigma-delta modulators are that finite component tolerances in constructing a practical quantiser result in gain error and dc offset and *not* differential nonlinearity. Practical high-order SDMs thus achieve high resolution with very low distortion, at least for high-level sinusoidal inputs. Fig. 4(a) shows a simulated output spectrum for a fourth-order modulator with a 1 kHz sinusoidal input at -12 dBFS. As for all of the simulation results presented in this paper, the oversampling factor is set to 64, and the sampling frequency $f_s = 2.8224$ MHz - this yields a baseband of 0 kHz to $f_B = 22.05$ kHz, equivalent to consumer digital audio systems with a sampling frequency of 44.1 kHz. The signal-to-noise ratio (SNR) in this example is 106 dB, and the output spectrum can be seen to be almost completely free of distortion tones. Fig. 4(c) shows that high resolution is maintained below the noise floor when a -125 dBFS signal is applied to the converter. Results like these have led many to believe that high-order SDMs behave as virtually ideal converters. However, recent research has indicated that at moderate input amplitudes, sigma-delta converters can be corrupted by idle tones and noise modulation, and thus deviate from our notionally perfect converter. Fig. 4(b) shows the output spectrum from the same modulator with a 1 kHz input at -47 dBFS, where low-level tones can be seen rising above the noise floor at several frequencies.

Techniques for eliminating idle tones in SDMs include the use of dither, and making the converter chaotic. In this paper we examine the efficacy of these techniques, and focus upon the associated dynamic-range penalties involved in successfully linearising SDMs. The study commences in Section 2 with a discussion of optimisation techniques that can be used to maximise resolution in SDMs of a given order and oversampling ratio. Section 3 examines SDM idle-tone phenomena in greater detail, and reviews techniques for revealing their presence in high-resolution converters. Sections 4 and 5 examine dithered and chaotic sigma-delta modulators respectively, and use an optimisation technique to determine dynamic-range penalties for SDM linearisation.

2 OPTIMAL SIGMA-DELTA MODULATORS

High-resolution sigma-delta modulators require as much as possible of the quantisation noise to be removed from the baseband - a requirement which implies that the noise-shaping function NS be highpass, and hence the loop filter H be lowpass. If the noise-shaping response is described in terms of z -domain poles p_i and zeros z_i , then for an n th-order modulator,

$$NS(z) = \frac{\prod_{i=1}^n (z - z_i)}{\prod_{i=1}^n (z - p_i)}. \quad (3)$$

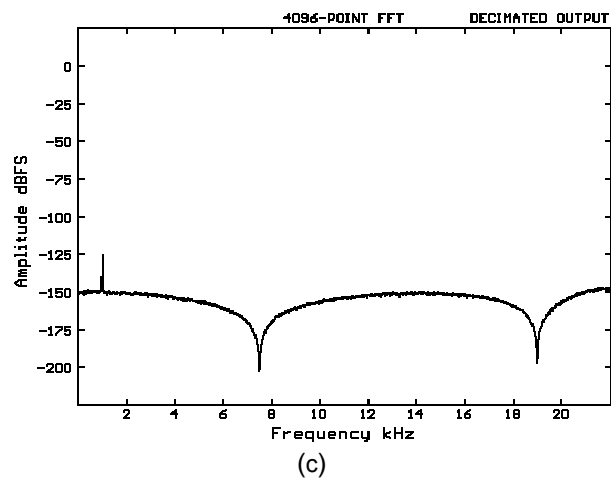




Fig. 4. Output from fourth-order SDM with 1 kHz input at (a) -12 dBFS, (b) -47 dBFS, and (c) -125 dBFS.

Chao et al. [7] developed a loop filter structure suitable for high-order SDMs, where flexibility in specifying the noise-shaping response of the modulator is achieved by embedding cascaded integrators within a recursive feedback loop. Fig. 5 shows a modified version of Chao's structure, where the noise-shaping response is determined by feedforward coefficients a_i feeding the single-bit quantiser, and feedback coefficients b_i . This loop-filter structure has been used in several commercially-available SDM ADCs and DACs [8], [9], and was used to generate all of the simulation results presented in this paper. It is easy to show that

$$H(z) = \frac{a_1(z-1)^{n-1} + a_2(z-1)^{n-2} + \dots + a_{n-1}(z-1) + a_n}{(z-1)^n + b_1(z-1)^{n-1} + b_2(z-1)^{n-2} + \dots + b_{n-1}(z-1) + b_n}. \quad (4)$$

Combining Eqs. (2) and (4), the noise-shaping function can be determined in terms of a_i and b_i ,

$$NS(z) = \frac{(z-1)^n + b_1(z-1)^{n-1} + b_2(z-1)^{n-2} + \dots + b_n}{(z-1)^n + (a_1 + b_1)(z-1)^{n-1} + (a_2 + b_2)(z-1)^{n-2} + \dots + (a_n + b_n)}. \quad (5)$$

Comparing Eqs. (3) and (5), it is clear that noise-shaping zero locations z_i are solely dependent upon feedback coefficients b_i , while pole locations p_i depend upon both a_i and b_i .

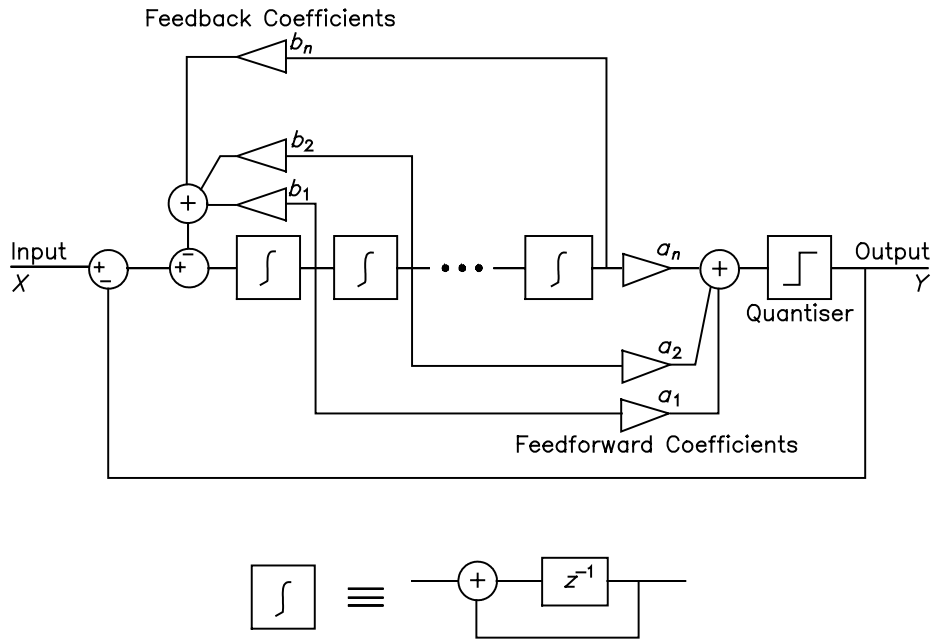


Fig. 5. Sigma-delta modulator with cascaded-integrator loop-filter structure.

Several workers have investigated the effect that the noise-shaping characteristic has upon the dynamic range, stability and overload characteristics of sigma-delta modulators [7], [10], [11]. A popular choice for the noise-shaping pole locations is a Butterworth configuration, such that, for a given oversampling factor, p_i are determined by a single variable - the Butterworth -3 dB cutoff frequency, f_c . If the zeros z_i all reside at dc then the noise-shaping response is Butterworth highpass (Fig. 6). For dc zeros, $z_i = 1$, which for the cascaded-integrator filter structure corresponds to the feedback coefficients b_i being set to zero, while the feedforward coefficients a_i are determined by f_c . Fig. 7 shows the highpass-Butterworth noise-shaping

magnitude response obtained when $f_c = f_s/25$ for a fourth-order modulator. Note that the high-frequency gain of the noise-shaping response is greater than 0 dB - that is, the high-frequency quantisation noise seen at the output of the system is amplified by the noise-shaping action. This high-frequency gain effect increases as the cutoff frequency of the Butterworth filter increases, and tends to reduce the stability of the system.




Fig. 6. Noise-shaping pole and zero positions for fourth-order SDM with highpass-Butterworth noise-shaping response.

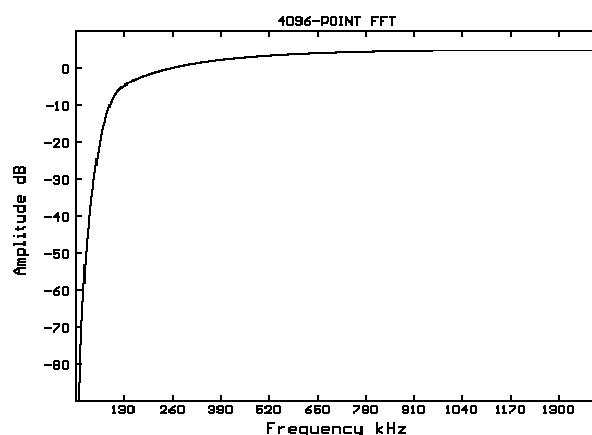


Fig. 7. $NS(f)$ for fourth-order modulator with highpass-Butterworth noise-shaping response.

Fig. 8(a) shows the noise-shaping magnitude response across the baseband, clearly indicating the presence of noise-shaping zeros at dc. A technique that can be used to increase the SNR of the modulator is to move the noise-shaping zeros away from dc and distribute them across the baseband in conjugate pairs [7]. This corresponds to non-zero feedback coefficients b_i in the cascaded-integrator loop-filter structure. Schreier [11] has determined the zero locations which yield optimal signal-to-noise ratios at the output of the modulator; Table 1 reproduces these results for modulator orders 1 to 4. Fig. 8(b) shows the noise-shaping response for the fourth-order modulator with Butterworth poles ($f_c = f_s/25$) when the zeros are moved from dc to these optimal locations. Fig. 9 shows the new zero positions on the unit circle in the z-domain. This modulator is used for several simulations below, where it is referred to as the "standard" fourth-order modulator.




Fig. 8. $|NS(f)|$ over baseband for fourth-order SDM with Butterworth noise-shaping poles. (a) Thin trace - noise-shaping zeros at dc. (b) Bold trace - distributed noise shaping zeros minimise baseband noise.

Table 1. Optimal noise-shaping zero locations.

| Modulator Order n | Zero frequencies, normalised to f_B | Zero frequencies for $f_B = 22.05$ kHz (kHz) |
|---------------------|--|--|
| 1 | 0 | 0 |
| 2 | $\pm \sqrt{1/3}$ | ± 12.7 |
| 3 | $0, \pm \sqrt{3/5}$ | $0, \pm 17.1$ |
| 4 | $\pm \sqrt{[(3/7) \pm \sqrt{(3/7)^2 - (3/35)}]}$ | $\pm 7.5, \pm 19.0$ |

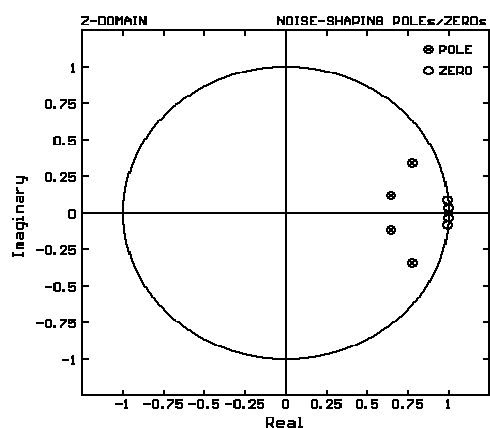


Fig. 9. Noise-shaping poles and zeros plotted in z-domain for fourth-order SDM where zeros have been optimally distributed across the baseband.

Another way to increase the resolution of an SDM with Butterworth poles is to increase the cutoff frequency f_c [11], although such an action will also tend to compromise stability since the high-frequency gain of NS increases, with a corresponding increase in the total noise power within the loop. A decrease

in the stability of an SDM system with Butterworth poles results in a lower value of x_{opt} , the input amplitude which yields the peak SNR from the modulator before the onset of large-amplitude oscillations. In general there will be an optimal value of f_c which yields the peak SNR obtainable from a modulator with a given oversampling ratio and Butterworth pole locations. Schreier [11] describes a technique for empirically determining this value by simulation, where x_{opt} is found by applying a dc input to the modulator and noting the value at which the modulator becomes unstable. Instability is detected by observing the peak value at the input to the quantiser and noting the input amplitude x_{max} at which it exceeds a predefined value. The instability criterion adopted in this paper, for modulators with normalised feedforward coefficients¹ and 1-bit quantisers whose outputs can take the values ± 1 , is the smallest dc input which causes the peak quantiser input to exceed 5 for any sample within 10^5 oversampled time steps. Fig. 10 shows such a simulation for the standard fourth-order modulator, where $x_{\text{max}} = 0.71$.




Fig. 10. Determining maximum stable input amplitude for standard fourth-order modulator. (a) Bold trace - peak quantiser input as a function of dc modulator-input amplitude. (b) Thin trace - SNR as a function of peak sinusoidal input amplitude.

Once x_{max} has been determined, the peak signal-to-noise ratio SNR_{max} for the modulator is found by, as in Schreier's study [11], applying a 1 kHz sinewave with peak amplitude $x_{\text{opt}} = x_{\text{max}} - 1$ dB to the modulator, and comparing the fundamental energy at the modulator output to the quantisation-error energy in the frequency domain. To obtain adequate accuracy in the SNR calculation with relatively short simulation times, the oversampled modulator output is decimated by a factor of 64 before transforming to the frequency domain. Fig. 11 shows the decimation structure used in the simulations, where a 7-stage comb filter precedes two half-band filters; this structure is relatively complex, but has a stopband rejection performance which ensures that, in the most critical case, aliased out-of-band quantisation noise does not increase the baseband noise-floor of the modulators by more than 0.05 dB. Fig. 10(b) shows how the SNR of the standard fourth-order modulator varies with peak sinusoidal input amplitude, and indicates that the SNR peaks at approximately $x_{\text{opt}} = 0.64$. This result along with simulations performed on different modulators vindicates the $x_{\text{opt}} = x_{\text{max}} - 1$ dB input-amplitude criterion in determining SNR_{max} . Note that determining x_{opt} by observing input quantiser levels for dc modulator inputs requires far less simulation time than directly calculating the signal-to-noise ratio as a function of input amplitude, where for the latter approach a time-consuming decimation routine must be implemented for each SNR calculation.

¹ Since single-bit quantisers are effectively sign-detectors, then SDM operation is invariant with scalings of the feedforward coefficients a_i . For all of the simulations presented in this paper, a_i are scaled such that $a_1 = 1$. This scaling procedure is important to observe when determining a_{max} by detecting peak values at the quantiser input, or in setting dither amplitudes (see Sec. 4).




Fig. 11. Decimation structure used in simulations.

This procedure, shown as a flow diagram in Fig. 12, allows the peak SNR ratio for a modulator with Butterworth poles to be calculated as a function of the Butterworth cutoff frequency f_c . It is a fairly straightforward optimisation task to find $f_{c_{opt}}$, the Butterworth cutoff frequency which maximises the SNR obtainable from the modulator. For the purposes of this paper, we shall term the modulator resulting from such an optimisation an "optimal sigma-delta modulator." An optimisation routine was written to determine $f_{c_{opt}}$ using a 1-dimensional pattern search. The results of an optimisation for the fourth order modulator are collated in Table 2, indicating that the peak SNR obtained from the optimal fourth-order SDM is approximately 7 dB greater than the standard modulator with $f_c = f_s/25$. Note that the zero locations remain fixed at their optimal locations as determined by Schreier [11], and hence the optimisation process does not alter the values of the feedback coefficients b_i . We use the optimisation process described in this section to investigate characteristics of optimal dithered and chaotic SDMs in Secs. 4 and 5 respectively.

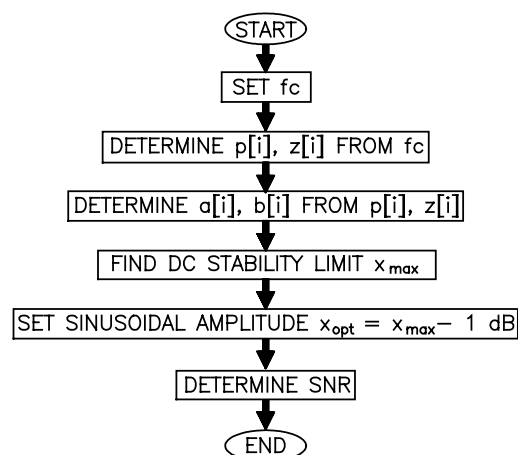


Fig. 12. Flow diagram for determining peak SNR from SDM with Butterworth noise-shaping poles.

Table 2. Result of fourth-order modulator optimisation.

| Modulator | Standard | Optimal |
|------------------|----------|---------|
| f_c/f_s | 0.04 | 0.0642 |
| a_{opt} , dBFS | -3.9 | -8.4 |
| SNR_{max} , dB | 102.0 | 109.3 |
| a_2 | 0.313 | 0.492 |
| a_3 | 0.055 | 0.131 |
| a_4 | 0.0045 | 0.0157 |
| b_1 | 2.07e-3 | 2.07e-3 |
| b_2 | 2.07e-3 | 2.07e-3 |
| b_3 | 9.95e-3 | 9.95e-3 |
| b_4 | 4.98e-7 | 4.98e-7 |

3 IDLE TONES AND NOISE MODULATION IN SDMS

It has been known for some time that low-order sigma-delta modulators suffer from noise modulation [12, 13], threshold effects [14], and, for certain input signals, low-level tones [15]. In many ways such non-idealities are not surprising, since sigma-delta modulation, like undithered PCM quantisation, is a deterministic mapping operation [16]. Low-order SDMs ($n \leq 2$) have only a small number of state variables, which for constant inputs tends to result in repeating bit patterns at the modulator output, and a high degree of correlation between the input signal and quantisation noise is to be expected. However, until recently it was generally believed that higher-order ($n > 2$) sigma-delta modulators did not suffer from such nonlinear artifacts. For example, Chao et al. argue in [7] that the greater number of state variables present in higher-order modulators lends a higher degree of randomisation to quantisation noise. Nevertheless, in the same paper results are presented which show that a fourth-order modulator suffers from at least 5 dB of noise modulation. We will now consider some of the techniques that exist for quantifying noise-modulation effects and idle-tone behaviour in sigma-delta modulators.

In a study of idle-tone behaviour in low-order modulators, Ledzius and Irwin [17] show that, for constant inputs, idle-tone frequencies are related to the modulator input amplitude x_{DC} . Although the highest-amplitude idle tones occur at frequencies close to half the sampling frequency, tones of significant amplitude can also occur within the baseband. For an SDM with quantiser output levels ± 1 , an important idle-tone occurs at a frequency given by

$$\begin{aligned} f_{IT} &= x_{DC} f_S \\ &= 2 x_{DC} f_B \text{ OSF}, \end{aligned} \quad \text{where} \quad (6)$$

f_B = upper baseband frequency,

OSF = oversampling factor.

Consider the standard second-order modulator shown in Fig. 13 ($\alpha = \beta = 1$); this modulator is operationally equivalent to a cascaded-integrator SDM structure with the loop-filter coefficient set $\{a_1 = 1, a_2 = 0.5, b_1 = 0, b_2 = 0\}$. Fig. 14(a) shows 256 samples of the ac component of the time-domain output signal obtained for this modulator with $\text{OSF} = 64$ and $x_{DC} = 1/256$. The quantisation noise can be seen to be strongly periodic rather than random and, according to Eq. (6), should have an idle tone at frequency $f_{IT} = 11.0$ kHz. The frequency-domain representation of the modulator output shown in Fig. 14(b) indicates 7 idle tone frequencies within the baseband, including f_{IT} .




Fig. 13. Second-order sigma-delta modulator. Standard modulator is achieved by setting $\alpha = \beta = 1$, while chaotic modulators require $\alpha > 1$, and/or $\beta > 1$.




Fig. 14. Output of standard second-order modulator for dc input of 1/256. (a) Time domain. (b) Frequency domain.

Given the presence of idle tones in a quantisation noise floor, visual examination of a single frequency-domain noise-floor plot provides only a poor measure of whether the tones will be perceived as tonal or not. Schreier [18] presents a technique for quantifying the tonality of quantisation noise, based upon comparing the energy contained in the highest-amplitude FFT bins of the transformed noise signal to the total error energy, hence deriving a "tonality index" for the noise floor. Alternatively several noise floor plots can be arranged in a 3D format, where the eye is quite sensitive to detecting patterns in the noise floor that remain obscure in single plots. Ledzius and Irwin developed such a method in [17], where SDMs are subjected to dc-input sweeps and noise-floor spectra for several dc input levels are displayed in a single graph. Fig. 15 shows such a simulation for the standard second-order modulator, with a dc input spanning the range $0 \rightarrow 1/256$. Each trace in the plot represents the power spectrum of the quantisation noise obtained from a 4096-point FFT of the (decimated) modulator output signal. Fig. 15 clearly indicates an idle tone with frequency f_{IT} , increasing from dc to 11 kHz as the dc input level increases, and with an amplitude well above the level of the noise floor². Fig. 16 shows a similar simulation for the optimal fourth-order modulator discussed in Sec. 2, where idle tones f_{IT} and $2f_{IT}$ can again be seen above the level of the noise floor (note that the spectral notches observed in the noise floor at 7.5 kHz and 19.0 kHz are due to the optimally-located noise-shaping zeros - see Table 1). This result clearly indicates that higher-order modulators are *not* immune to undesirable idle-tone artifacts. As a point of reference, consider the 3D plot shown in Fig. 17, where a dc-input sweep is applied to a TPD-dithered 16-bit PCM quantiser; other than spectral ripples due to the finite number of FFT bins used in the simulation, the noise-floor can be seen to be invariant with input signal level, with no idle tones visible. This result again indicates the near ideal characteristics of correctly-dithered PCM quantisation, as discussed in Sec. 1.

² Note that inadequate stopband performance of the decimation filter used in these simulations can cause misleading results, since aliased high-frequency tones are indistinguishable from baseband tones. This potential problem is avoided for the simulations presented in this paper by using the high-performance decimation filter structure shown in Fig. 11, where high-frequency idle tones are attenuated by at least 250 dB.

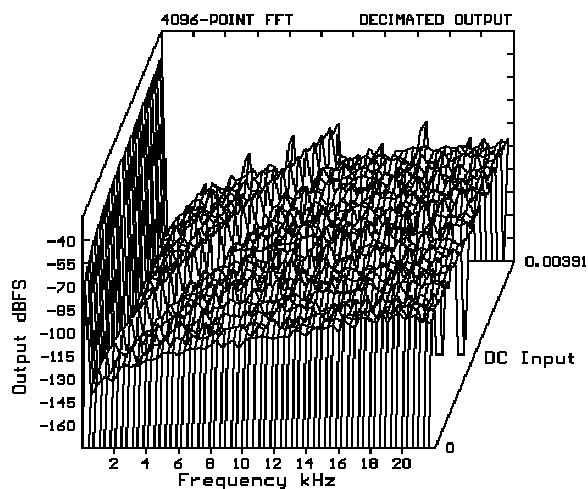


Fig. 15. DC input sweep for standard second-order modulator.




Fig. 16. DC-input sweep for optimal fourth-order sigma-delta modulator.




Fig. 17. DC input sweep for dithered 16-bit PCM quantiser.

In studies of dithered sigma-delta modulators, Norsworthy [19] and later Norsworthy and Rich [20] use an autocorrelation technique to detect the presence of idle tones in SDM quantisation noise, and show that the existence of idle tones corresponds to noise modulation when a modulator is subjected to a dc or sinusoidal signal of varying amplitude. Hence a modulator which is free from idle tones should also be free from noise modulation. The relationship between idle-tone artifacts and noise modulation has also been observed by Risbo [21], and is further supported by simulation results presented below in Secs. 4 and 5. Fig. 18 shows how the total quantisation noise power of the optimised fourth-order modulator changes as the amplitude of a 1 kHz sinusoid applied to the modulator input is ramped from -120 dBFS to 0 dBFS in steps of 1 dB. It can be seen that the quantisation noise energy of this modulator changes by approximately 3.8 dB over the range of input amplitudes; following the brief discussion of noise-modulation audibility in Section 1, this degree of noise modulation would be deemed to be audible under critical listening conditions. Fig. 19 shows a reference plot for the dithered 16-bit PCM quantiser, which is essentially free of noise modulation³; the small degree (less than 0.5 dB) of variation in the noise-floor energy observed in this diagram is due to the finite number of FFT bins used in the measurement.

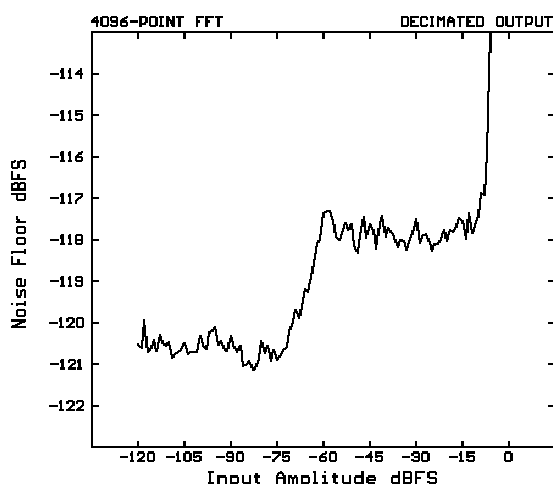


Fig. 18. Noise floor plot for optimal fourth-order SDM with 1 kHz sinusoidal input.

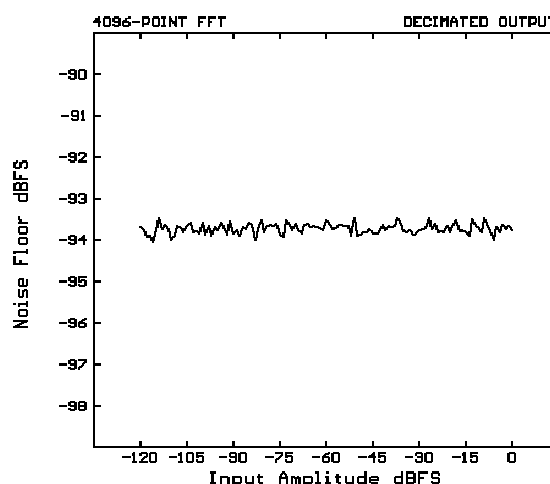


Fig. 19. Noise-floor plot for dithered 16-bit PCM quantiser with 1 kHz sinusoidal input.

Given the presence of undesirable idle tones and noise modulation in sigma-delta modulators, even for higher-order loop filters, there are several techniques that can be used to reduce or eliminate the problem:

- Increase the number of *linearly-spaced* quantiser steps.
- Increase the number of *logarithmically-spaced* quantiser steps.
- Introduce dither to the sigma-delta modulator.
- Make the sigma-delta modulator chaotic.

Unfortunately, increasing the number of quantiser levels mitigates against one of the fundamental benefits of sigma-delta technology, that of achieving high resolution without the need for accurate component matching. Although multi-bit SDMs with linearly-spaced quantiser levels can be designed such that finite component tolerances do not result in a high degree of correlation between the quantisation noise and the input signal [22], such designs are extremely sensitive to clock-timing errors, or "jitter" [23]. An alternative approach to multi-bit SDM design is to space the quantiser levels logarithmically, as recently described by Ledzius and Irwin [17]. Although such designs appear to be immune to idle-tone phenomena and are also relatively insensitive to component mismatches, they suffer from severe noise modulation.

³ As little as 1 dB noise modulation in any *single* critical bandwidth can be audible under critical listening conditions. Hence plots such as Fig. 19 which show the *total* error energy to be invariant with input signal characteristics are a necessary but not sufficient requirement for noise modulation to be inaudible.

In the remaining sections of this paper we will examine dither and chaos as strategies for linearising sigma-delta modulation. The 3D dc-input sweep and noise-floor plot techniques described in this section will be used to quantify the success of designs which aim to decorrelate quantisation error from input signal characteristics.

4 DITHERED SIGMA-DELTA MODULATORS

In Sec. 1 we briefly discussed the use of dither signals to linearise PCM quantisation, where it was noted that addition of the correct level of TPD dither increases the noise floor of the quantiser by 4.8 dB. Dither can also be used to reduce noise modulation and idle-tone phenomena in single-bit sigma-delta converters, where the dither action can be thought of as a phase-randomisation of idle tones such that they are indistinguishable from the random component of the quantisation error. In this section we investigate how effective dither is in eliminating idle tones in sigma-delta modulators, and quantify dynamic-range penalties for implementing dither.

Although high-frequency squarewaves can be used to dither SDMs [24], [25], the most complete elimination of idle-tone phenomena appears to occur when the dither signal is a random or pseudo-random signal. Generation of such a dither signal is a simple matter in sigma-delta DACs, although in ADCs it may significantly increase the complexity of the conversion system if the dither signal has many levels. Norsworthy [19] has identified the primary design parameters that require attention in dithered sigma-delta modulators:

- The optimal position at which to apply dither within the sigma-delta loop.
- The relative dither power required to decorrelate quantisation noise from input-signal characteristics.
- The dynamic-range penalty for implementing SDM dither.
- The best probability distribution for the dither signal.




Fig. 20. General SDM structure where dither is added to input of 1-bit quantiser.

Norsworthy [19], and later Norsworthy and Rich [20], describe how a pseudo-random dither signal added to the input of the single-bit quantiser (Fig. 20) is conveniently noise-shaped by the noise-shaping action of the sigma-delta loop, such that the reduction in baseband dynamic range due to the addition of dither is minimised. Norsworthy has also addressed the question of the most effective dither amplitude to use. Measurement results are presented in [20] for a prototype second-order modulator which showed how addition of rectangular-probability-distribution (RPD) dither spanning ± 0.5 - that is, half of the one-bit quantiser output range - eliminated idle tones, at the expense of increasing the noise floor by 2 dB. Simulations were performed in an attempt to confirm this result. The standard second-order modulator described in Sec. 3 was dithered using a ± 0.5 RPD noise, and stimulated with a dc input of $1/256$. The results shown in Figs. 21 and 22 indicate that idle tones are almost eliminated (these diagrams can be directly compared against results for the undithered second-order modulator shown in Figs. 14 and 15). When the modulator was stimulated using a 1 kHz sinusoidal signal, the noise-floor power increase due to the addition of dither was found to range between 2 dB and 8 dB, depending upon the input amplitude. The *peak* SNR obtainable from this standard second-order modulator was 6 dB lower with the addition of dither.




Fig. 21. Standard second-order modulator with ± 0.5 RPD dither and dc input of $1/256$.
(a) Time domain. (b) Frequency domain.




Fig. 22. DC-input fade for standard second-order modulator with ± 0.5 RPD dither.

As well as eliminating idle tones and noise modulation, correctly-implemented dither also has the ability to eliminate low-level dead zones in lower-order modulators, so that low-amplitude input signals below the level of the quantisation noise can be accurately resolved. Fig. 23(a) shows the output spectrum for the standard undithered second-order SDM for a 1 kHz sinusoidal input signal at -115 dBFS; the quantisation error is extremely tonal, and for input amplitudes below -140 dBFS the input signal is not coded at all (for a discussion of dead-zone phenomena, see Naus and Dijkmans [14]). Fig. 23(b) shows the output spectrum for the same input signal when the modulator is dithered with ± 0.5 RPD noise; the sinusoidal input is now resolved at the correct level without tonal distortion.

Generally, the higher the dither amplitude the more successfully idle-tones are removed from the quantisation noise floor. Fig. 24 shows noise-floor spectra for the standard fourth-order modulator with various levels of RPD dither and a dc-input of $1/256$ (the traces have been vertically offset to aid clarity). The amplitude of the f_{IT} idle tone at 11 kHz clearly decreases relative to the random noise-floor component as the dither amplitude increases, and is finally eliminated when the dither spans a range of ± 0.375 .




Fig. 23. Output spectra for standard second-order modulator with 1 kHz input at -115 dBFS. (a) Undithered. (b) Dithered.

For higher-order modulators ($n > 2$), the dynamic range penalty suffered when dither is implemented can be greater than 6 dB. A simulated example of the SNR penalty to be paid for dithering a sixth-order system is presented by Risbo in [26], where introduction of ± 0.5 RPD dither reduces the peak SNR by 27 dB. Such large SNR reductions in dithered higher-order systems are partly due to the increase in noise within the sigma-delta loop, but also because the addition of dither tends to reduce system stability, requiring a reduction in loop gain and corresponding increase in baseband quantisation noise [21]. The reduction in stability is indicated in Fig. 25, which shows the maximum amplitude occurring at the quantiser input as a function of dc input level for the standard fourth-order modulator. The two traces in this diagram correspond to undithered and ± 0.5 RPD dithered modulators, and it is clear that the addition of dither causes instability to be approached at lower input levels. Thus modulators whose loop-filter characteristics are optimised for maximum dynamic range without dither often suffer from poor stability when dither signals are introduced within the loop.




Fig. 24. Effect of RPD dither amplitude on standard fourth-order SDM with dc input = $1/256$. Modulator output spectra for peak dither amplitudes (a) 0, (b) 0.125, (c) 0.25, and (d) 0.375.

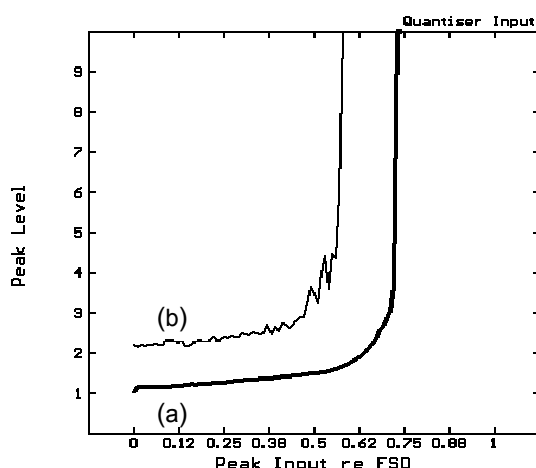


Fig. 25. Peak amplitude at quantiser input as a function of dc input amplitude for fourth-order modulator. (a) Bold trace: no dither. (b) Thin trace: ± 0.5 RPD dither.

In general, *optimal* dithered sigma-delta modulators - that is, dithered SDMs with loop filter characteristics optimised for maximum dynamic range using the routine described in Sec. 2 - will suffer the lowest-possible SNR penalties for dithering. Simulations were performed to determine minimum SNR penalties for optimal RPD-dithered modulators of orders ranging from 1 to 5, where the dither amplitude was increased in approximate steps of 1 dB until modulator linearisation was achieved; a modulator was considered to be "linearised" if no idle tone artifacts were visible in 3D dc-input sweep plots, and noise modulation was less than 1 dB for input signal amplitudes within the dynamic range of the modulator.

Table 3. Dynamic-range penalties for optimal dithered SDMs.

| Order n | Minimum RPD amplitude required for linearisation | SNR_{\max} for optimal modulator (dB) | | Dynamic range penalty (dB) |
|--------------|--|---|----------|-------------------------------|
| | | undithered | dithered | |
| 2 | 0.4 | 75.8 | 71.3 | 4.5 |
| 3 | 0.28 | 93.6 | 89.2 | 4.4 |
| 4 | 0.25 | 109.3 | 104.0 | 5.3 |
| 5 | 0.2 | 121.6 | 117.4 | 4.2 |

Table 3 shows the minimum RPD dither amplitudes required to linearise the simulated systems with orders 2 to 5⁴; note that the peak SNRs obtained for modulators optimised without dither agree quite well with Schreier's results in [11], and hence verify the accuracy of the optimisation process described in Sec. 2. The results show that as the modulator order n increases, the dither amplitude required for linearisation decreases - this trend has also been noted by Norsworthy [19], [20]. SNR penalties for linearisation remain approximately constant with n , at 4 to 5 dB. Interestingly, these figures are close to the SNR penalty paid for dithering a PCM quantiser (4.8 dB). Another study by the authors extending the results to higher orders has shown that similar SNR penalties are paid for linearising 6th- and 7th-order systems [27]. Note that such results are significantly lower than the figure of 27 dB obtained by Risbo in [26] for a 6th-order modulator. This discrepancy can be explained by the authors' use of optimisation techniques and the use of many dither amplitudes to determine the minimum dither amplitude required for

⁴ Note that the 1st-order systems were not completely linearised for any of the dither levels tested.

linearisation. Preliminary investigations indicate that SNR penalties for linearisation using dither do not greatly change as the oversampling factor changes [27]. However, note that results presented in Table 3 are specific to SDMs with Butterworth noise-shaping poles and optimally-located zeros; results may differ for modulators with alternative noise-shaping responses.

Experiments were also performed with alternative dither amplitude distributions, with some interesting results:

- TPD-dithered systems tend to suffer slightly higher SNR penalties compared to RPD-dithered systems - by approximately 1 dB. Thus while a triangular amplitude distribution is the correct dither signal to use with PCM quantisation, RPD is to be preferred in sigma-delta modulation.
- High-pass TPD (HPTPD) dither can easily be formed by subtracting consecutive samples from an RPD noise source [4], and, with a power spectral density which is skewed towards high frequencies away from the signal band, appears an attractive candidate for SDM dither. However, experimentally TPD and HPTPD dithers gave very similar performance in terms of their ability to linearise SDMs and the resultant SNR penalties. Thus it would appear that the dither spectrum is not a critical factor in dithering SDMs, perhaps because a one-bit quantiser represents a gross nonlinearity and thus normal rules of superposition [4] are no longer valid. More important factors appear to be the randomness of the dither signal, the dither amplitude, and the dither amplitude distribution.
- For higher-order systems ($n \geq 3$), single-bit (BPD) dither can be successfully used with little or no additional SNR penalty compared to RPD-dithered systems [28]. This is an important result, since single-bit dither can easily be implemented in sigma-delta ADCs as a single switched capacitor at the quantiser input, controlled from the digital domain.

As an example of the higher-order dithered modulators obtained following optimisation, Figs. 26 and 27 show simulation results for the optimal fourth-order SDM with ± 0.25 RPD dither. No idle tones are apparent in the dc-input fade (Fig. 26), while noise modulation is less than 1 dB (Fig. 27). The application of dither has increased the noise floor of the modulator by between 5 dB and 8 dB (compare Fig. 27 with Fig. 18), while the peak SNR obtainable from the modulator has decreased by 5.3 dB.




Fig. 26. DC-input fade for optimal fourth-order modulator with ± 0.25 RPD dither.

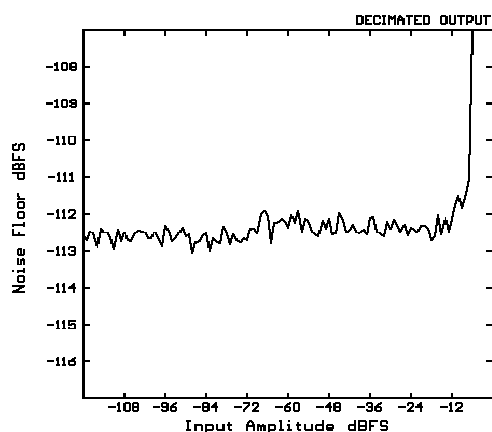


Fig. 27. Noise floor of optimal fourth-order modulator with ± 0.25 RPD dither. Input signal is 1 kHz sinewave.

As well as linearising the in-band quantisation noise floor, an appropriate level of dither also has the beneficial effect of reducing high-frequency ($\gg f_B$) idle tones. This can be extremely important in practical implementations of high-order sigma-delta modulators, where aliasing of high-frequency (HF) tones due to circuit nonlinearities and coupling mechanisms can dominate baseband performance [29], [30]. Risbo [21] has speculated that the presence of HF idle tones is intimately linked with baseband idle tones and noise modulation - that is, suppression of HF tones results in a modulator free from undesirable baseband artifacts. Fig. 28 shows simulation results obtained with the standard fourth-order modulator with various amounts of RPD dither, where the HF idle tone at 1340 kHz (due to a dc input of 0.05) is eliminated with the application of dither spanning ± 1.0 .

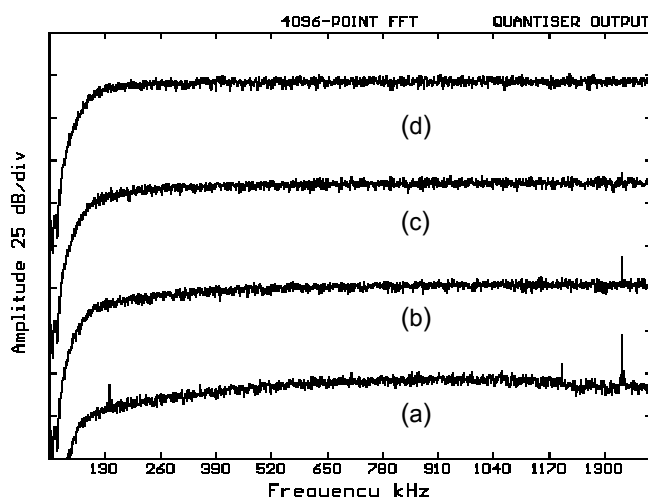


Fig. 28. High-frequency idle-tone behaviour for various levels of dither in fourth-order modulator with dc input of 0.05. RPD dither range is (a) 0, (b) ± 0.5 , (c) ± 0.75 , and (d) ± 1.0 .

5 CHAOTIC SIGMA-DELTA MODULATORS

An alternative approach to using dither to eliminate low-level artifacts in sigma-delta converters is to make the modulator *chaotic*, where noise-shaping zeros are moved outside the unit circle in the z -domain. Such an approach is simply implemented by appropriate choice of loop-filter coefficients, and has the advantage over dithered sigma-delta modulators that no random-noise generator circuitry is required - hence making the technique especially attractive for linearising sigma-delta ADCs.

An explanation of why chaotic SDM systems are likely to be less prone to idle tones and limit cycles was recently provided by Schreier [18]. With reference to Fig. 3 and Eq. (2), a modulator with noise-shaping zeros outside the unit circle is equivalent to the poles of the loop filter $H(z)$ also being outside the unit circle - that is, $H(z)$ is open-loop unstable. Such a system exhibits extreme sensitivity to initial conditions and small changes in the condition of the modulator, hence any repetitive pattern that would have occurred at the output of a non-chaotic modulator (with noise-shaping zeros on the unit circle) will tend to be broken up. The speed at which limit cycles are disrupted is determined by how unstable $H(z)$ is - that is, how far outside the unit circle the noise-shaping zeros are situated. However, although the output of a chaotic system is generally non-periodic, with a continuous spectrum, Risbo points out in [21] that this condition does not preclude the *combination* of tones and noise in the modulator quantisation error. Indeed, Schreier's simulation results for first- and second-order modulators show that introduction of chaos can reduce the amplitude of but not eliminate tonal quantisation components [11]. Similarly Motamed et al. [31] note that although the possibility of quantisation noise with stable periodic components is avoided in chaotic systems, tones may well be observed in short-term noise spectra (such as a measurement over a finite number of samples). In this section we investigate how successful a certain class of chaotic modulator is in eliminating tones and noise modulation in sigma-delta conversion.

With the restriction that at least one noise-shaping zero must lie outside the unit circle, Risbo [26] identifies three classes of chaotic sigma-delta modulator:

- Take the loop filter $H(z)$ for a standard sigma-delta modulator and scale the poles (that is, the noise-shaping zeros) to positions outside the unit circle.
- Design $NS(z)$ to be minimum phase and reflect zeros to reciprocal locations outside the unit circle.
- Design part of $NS(z)$ with a lower order than the order of the modulator n and with zeros on the unit circle, and arrange the remaining poles and zeros as allpass sections (with the allpass zeros outside the unit circle).

In this study we focus on the first of these options, where all noise-shaping zeros are outside the unit circle, and by the same amount. For example, Fig. 29 shows z -domain noise-shaping pole and zero placement for a fourth-order chaotic modulator with poles in a Butterworth configuration and zeros with radii $r_z = 1.05$; zero frequencies are equal to the values found in an optimal zero configuration (Table 1).

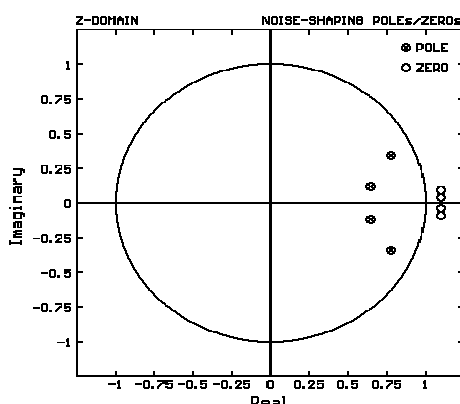


Fig. 29. Noise-shaping pole/zero placement in z -domain for chaotic fourth-order modulator with $r_z = 1.05$.

Hein [32] recently described a simulated comparison of a standard second-order modulator with a chaotic system where all noise-shaping zeros were moved outside the unit circle in the z -plane. Introducing chaos to the system was shown to break up an idle tone in the standard modulator caused by a dc input signal. Fig. 13 shows the standard second-order SDM architecture, where integrator-pole (and hence loop-filter pole, and also noise-shaping zero) locations are set by α and β . Noise-shaping zeros are moved outside the unit circle by setting either or both α , β to be greater than unity. Consider Fig. 30, which shows time- and frequency-domain plots of the standard (nonchaotic) second-order modulator output for a dc input of $1/128$ (the oversampling factor was again set to 64). The periodic nature of the output signal is clearly evident, with three tones present within the baseband. Both noise-shaping zeros were then moved outside the unit circle by setting $\alpha = \beta = r_z = 1.01$, and the system resimulated for the same input signal. The results shown in Fig. 31 indicate that the introduction of chaos has completely broken up any periodicity in the modulator output signal, the frequency-domain representation of the output appearing as random noise in the baseband. Although using a different oversampling factor and dc-input level, such a result agrees with Hein's simulation results published in [32].




Fig. 30. AC component of output from non-chaotic second-order modulator with $r_z = 1.0$ and dc input = $1/128$. (a) Time domain. (b) Frequency domain.




Fig. 31. AC component of output signal from chaotic second-order modulator with $r_z = 1.01$ and dc input = $1/128$. (a) Time domain. (b) Frequency domain.

However, further simulations with different input signals indicate that this chaotic second-order system does not completely break up all tonal components of the quantisation noise. Consider the baseband

output spectra shown in Fig. 32(a) and Fig. 32(b), corresponding to the standard and chaotic modulators respectively, for a dc input equal to $1/256$. Although the introduction of chaos clearly makes the modulator output spectrum less tonal, the amplitude of the f_{IT} idle tone remains approximately unchanged.




Fig. 32. Output spectra of second-order modulators with dc input = $1/256$.
(a) Non-chaotic modulator, $r_z = 1.0$, and (b) chaotic modulator, $r_z = 1.01$.

In an attempt to determine whether higher levels of chaos (that is, $r_z > 1.01$) can completely remove idle-tone artifacts, fourth-order modulators with noise-shaping poles arranged in a Butterworth configuration ($f_c = f_s/25$) and zeros at various positions outside the unit circle were each simulated with a dc input equal to $1/256$. Fig. 33 shows baseband output spectra for r_z ranging between 1.0 and 1.07, where each plot has been vertically offset to aid clarity. Relative to the level of the output noise floor, the amplitude of the $f_{IT} = 11$ kHz idle-tone is reduced as the noise-shaping zeros move further outside the unit circle, until it is eliminated in the upper plot for $r_z = 1.07$. However, moving the noise-shaping zeros outside the unit circle tends to reduce suppression of baseband quantisation noise, hence as well as reducing SNR the introduction of chaos can have the surprising effect of *increasing* the absolute amplitude of baseband idle tones [33]. For example, Fig. 34 compares baseband noise spectra for two fourth-order modulators - one with noise-shaping zeros on the unit circle ($r_z = 1.0$), and the other with chaotic zero placement ($r_z = 1.01$). It is clear that introducing chaos has suppressed the idle tone at 11 kHz relative to the noise floor, but increased it in absolute terms. Fig. 35 shows a set of simulations for fourth-order modulators with a dc input of 0.05, which show how the high-frequency idle tone at 1.34 MHz is affected by the degree of chaos. The idle tone is almost eliminated in the upper plot where $r_z = 1.1$.

These results suggest that chaotic sigma-delta modulators are not completely linearised until the noise-shaping zero radii are well outside the unit circle, $r_z > 1.05$, a finding which agrees with Risbo's study of sixth-order chaotic modulators [26]. In order to determine SNR penalties associated with such high levels of chaos, an extensive series of simulations was performed on modulators with orders ranging from 1 to 5, with r_z ranging from 1.0005 to 1.15, where the degree of chaos ($r_z - 1$) was incremented in approximate steps of 1 dB. The noise-shaping poles of each modulator were, as for the dither simulations, arranged in a Butterworth configuration, while noise-shaping zero frequencies were set to their optimal values (Table 2). Several workers, including Motamed et al. [31] and Risbo [21], have noted that, as for the introduction of dither, implementing chaos tends to deteriorate the stability of sigma-delta modulators. Such behaviour is confirmed in Fig. 36 where a chaotic system with $r_z = 1.05$ becomes unstable for lower dc input amplitudes than a nonchaotic system with the same noise-shaping pole locations. Hence for each chaotic system investigated, the Butterworth cutoff frequency f_c was optimised using the procedure outlined in Sec. 2 to yield the maximum-possible SNR. The peak SNR obtained from each chaotic modulator was then compared against that available from the optimal non-chaotic modulator of the same order; SNR penalties associated with minimum chaos levels required for linearisation are collated in Table 4.




Fig. 33. Effect of noise-shaping zero locations on baseband idle tones in chaotic fourth-order SDM with dc input = $1/256$. Modulator output spectra for zero radii r_z = (a) 1.0, (b) 1.002, (c) 1.035, and (d) 1.07.




Fig. 34. Fourth-order modulators with dc input = $1/256$. (a) Non chaotic, $r_z = 1$. (b) $r_z = 1.01$, indicating how moderate degree of chaos can increase absolute idle-tone amplitude.




Fig. 35. Effect of noise-shaping zero locations on high-frequency idle tones in chaotic fourth-order SDM with dc input = 0.05. Modulator output spectra for zero radii r_z = (a) 1.0, (b) 1.025, (c) 1.05, and (d) 1.1.




Fig. 36. Maximum quantiser input amplitude as function of dc input amplitude for fourth-order modulators. (a) Bold trace: standard non-chaotic modulator. (b) Thin trace: chaotic modulator, $r_z = 1.05$.

Table 4. Reduction in SNR for chaotic SDMs.

| Order n | Noise-shaping zero radii, r_z | SNR reduction compared to non-chaotic modulator (dB) |
|--------------|------------------------------------|---|
| 2 | 1.15 | 36.5 |
| 3 | 1.1 | 43.5 |
| 4 | 1.089 | 65.7 |
| 5 | 1.056 | 60.7 |

These results indicate that, while chaos can reduce the tonal nature of sigma-delta modulation, the technique struggles to completely linearise modulators, and in comparison to dithered modulators is expensive in terms of SNR reduction. This study has focused upon a class of chaotic modulator where noise-shaping zeros have positions in the z -domain which are simply scaled from non-chaotic noise-shaping functions. Risbo describes in [26] how chaotic modulators with allpass-sections reduce idle tones and noise modulation with a smaller dynamic-range penalty than scaled-zero chaotic systems, especially if the allpass zeros are close to $f_s/2$ (where they tend to suppress high-frequency idle tones). However, further results presented by the authors in [34] suggest that, although more efficient at linearisation than scaled-zero systems for higher-order modulators, chaotic SDMs with allpass-sections are also less efficient at removing low-level artifacts than dither.

6 CONCLUSIONS

In this paper we have compared the cost in terms of dynamic range for two techniques - dither and chaos - that have been proposed as methods for linearising sigma-delta modulators. In Sec. 1 we discussed the characteristics required of an "ideal" conversion process, where no distortion tones appear in the converter output, and the noise floor of the converter is invariant with changes in the input signal. The error associated with such an ideal conversion is psychoacoustically benign. Ideal TPD-dithered PCM quantisers meet the conditions of an ideal conversion, but at the expense of a reduction in signal-to-noise ratio of 4.8 dB compared to the undithered case. However, practical PCM converters with finite component tolerances suffer from noise modulation and distortion. Conversely, sigma-delta modulators are extremely tolerant of component mismatches, but inherently suffer from noise modulation and idle tones. A study of tone artifacts in SDMs, using 3D plots of spectra obtained with dc-input fades and noise-modulation plots with sinusoidal stimuli, indicated that even high-order systems suffer from these

unwanted errors. It can be shown that idle-tone frequencies are related to the dc-input amplitude applied to the modulator, and that noise modulation in SDMs is intimately related to the presence of idle tones. Three options for eliminating idle tones in SDMs were considered - multibit quantisers, dither, and chaos. The first of these options was rejected on the grounds that multibit quantisers increase modulator complexity, and either suffer from the effects of finite component tolerancing or exhibit high sensitivity to clock jitter.

Since the introduction of either dither or chaos can be shown to compromise the stability of sigma-delta modulators, a valid study of the consequences of implementing either technique requires that all modulators compared should have loop filters individually optimised for maximum signal-to-noise ratio. An optimisation process was described in Sec. 2, which involves determining the maximum stable input amplitude for a modulator with noise shaping poles arranged in a Butterworth configuration. Excepting the chaotic modulators, all noise-shaping zeros were set to the optimal locations determined by Schreier [11]. In each case a pattern-search technique was used to find the Butterworth cutoff frequency which yielded maximum signal-to-noise ratio.

In Sec. 4 we showed that dither of sufficient amplitude can completely linearise sigma-delta modulation. The simulation results further indicate that, as well as linearising the baseband quantisation error, dither also eliminates high-frequency tones which can be detrimental to the performance of practical systems. Comparisons of optimised dithered and undithered systems indicated that the dither power required for linearisation tends to reduce as the order of the modulator increases, although the dynamic range penalty for dithering remains approximately constant with n at 4 to 5 dB. Simulations with various dither amplitude distributions suggest that, compared to TPD dither, RPD dither achieves linearisation at marginally-lower SNR cost. No significant performance differences were noted between TPD and HPTPD dither signals. For higher-order systems, linearisation can be achieved with single-bit (BPD) dither at no additional SNR expense compared to RPD dither, and such dither signals can easily be implemented in sigma-delta ADCs.

Finally, characteristics of chaotic sigma-delta modulators were studied. Although there exist several classes of chaotic systems, this study was confined to systems with zero locations scaled from their optimal (non-chaotic) locations to lie outside the unit circle. While chaos precludes the possibility of purely periodic output sequences, combinations of tones and noise can occur when the degree of chaos is moderate ($r_z \approx 1.01$). For all modulator orders relatively high degrees of chaos ($r_z > 1.05$) were found necessary to completely eliminate idle tones and noise modulation, although the reduction in SNR associated with implementing chaos is then high - for example, 66 dB in the case of the fourth-order modulator. The use of chaos thus appears far less efficient than dither at linearising sigma-delta modulators.

7 ACKNOWLEDGEMENT

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

8 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] J. Vanderkooy and S. P. Lipshitz, "Resolution Below the Least-Significant Bit in Digital Systems with Dither," *J. Audio Eng. Soc.*, vol. 32, pp. 106-113 (1984 Mar.).
- [3] J. Vanderkooy and S. P. Lipshitz, "Dither in Digital Audio," *J. Audio Eng. Soc.*, vol. 35, pp. 966-975 (1987 Dec.).

- [4] J. Vanderkooy and S. P. Lipshitz, "Digital Dither: Signal Processing with Resolution Far Below the Least-Significant Bit," *Proc. 7th. AES Int. Conf.: Audio in Digital Times* (Toronto, Ont., Canada, 1989).
- [5] B. A. Blesser, "Digitization of Audio: A Comprehensive Examination of Theory, Implementation, and Current Practice," *J. Audio Eng. Soc.*, vol. 26, pp. 739-771 (1978 Oct.).
- [6] J. C. Candy and G. C. Temes, "Oversampling Methods for A/D and D/A Conversion," in *Oversampling Delta-Sigma Data Converters*, Eds. J. C. Candy and G. C. Temes, IEEE Press (1992).
- [7] 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).
- [8] N. S. Sooch, J. W. Scott, T. Tanaka, T. Sugimoto, and C. Kubomura, "18-Bit Stereo D/A Converter with Integrated Digital and Analog Filters," presented at the 91st Convention of the Audio Engineering Society, *J. Audio Eng. Soc. (Abstracts)*, vol. 39, p. 994 (1991 Dec.), preprint 3113.
- [9] I. Fujimori, K. Hamashita, and E. J. Swanson, "A Fifth-Order Delta-Sigma Modulator with 110 dB Audio-Band Dynamic Range," presented at the 93rd Convention of the Audio Engineering Society, *J. Audio Eng. Soc. (Abstracts)*, vol. 40, p. 1052 (1992 Dec.), preprint 3415.
- [10] R. W. Adams, P. F. Ferguson, A. Ganesan, S. Vincelette, A. Volpe, and R. Libert, "Theory and Implementation of a Fifth-Order Sigma-Delta A/D Converter," *J. Audio Eng. Soc.*, vol. 39, pp. 515-528 (1991 July/Aug.).
- [11] 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.).
- [12] 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.).
- [13] J. C. Candy and O. J. Benjamin, "The Structure of Quantisation Noise from Sigma-Delta Modulation," *IEEE Trans. Comm.*, vol. COM-29, pp. 1316-1323 (1981 Sept.).
- [14] P. J. A. Naus and E. C. Dijkmans, "Low Signal-Level Distortion in Sigma Delta Modulators," presented at the 84th Convention of the Audio Engineering Society, *J. Audio Eng. Soc. (Abstracts)*, vol. 36, p. 382 (1988 May), preprint 2584.
- [15] R. M. Gray, "Oversampled Sigma-Delta Modulation," *IEEE Trans. Comm.*, vol. COM-35, pp. 481-489 (1987 May).
- [16] R. Schreier and W. M. Snelgrove, "Sigma-Delta Modulation is a Mapping," *Proc. 1991 IEEE Int. Symp. on Circuits and Systems*, pp. 2415-2418 (1991 May).
- [17] 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).
- [18] 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).

- [19] S. R. Norsworthy, "Effective Dithering of Sigma-Delta Modulators," *Proc. 1992 IEEE Int. Symp. on Circuits and Systems*, pp. 1304-1307 (1992 May).
- [20] 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.
- [21] L. Risbo, "Designing Tone-Free Modulators," in *Sigma-Delta Modulators - Stability Analysis and Optimisation*, Ph. D. Thesis, Technical University of Denmark (1994).
- [22] M. O. J. Hawksford, "Chaos, Oversampling, and Noise Shaping in Digital-to-Analog Conversion," *J. Audio Eng. Soc.*, vol. 37, pp. 980-1001 (1989 Dec.).
- [23] R. W. Adams, "Comments on 'Chaos, Oversampling, and Noise-Shaping in Digital-to-Analog Conversion'," *J. Audio Eng. Soc. (Letters to the Editor)*, vol. 38, pp. 766-768 (1990 Oct.).
- [24] J. D. Everard, "A Single-Channel PCM Codec," *IEEE J. Solid-State Circuits*, vol. SC-14, pp. 25-37 (1979 Feb.).
- [25] B. H. Leung, R. Neff, P. R. Gray, and R. W. Brodersen, "Area-efficient Multichannel Oversampled PCM Voice-Band Coder," *IEEE J. Solid-State Circuits*, vol. SC-23, pp. 1351-1357 (1988 Dec.).
- [26] 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).
- [27] C. Dunn and M. Sandler, "Psychoacoustically-Optimal Sigma-Delta Modulation," presented at the 99th Convention of the Audio Engineering Society, *J. Audio Eng. Soc. (Abstracts)*, vol. 43, p. 1092 (1995 Dec.).
- [28] C. Dunn and M. Sandler, "Efficient Linearisation of Sigma-Delta Modulators using Single-Bit Dither," *Elec. Let.*, vol. 31, pp. 941-942 (1995).
- [29] 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.).
- [30] R. Adams, "Design Aspects of High-Order Delta-Sigma A/D Converters," *Proc. 1994 IEEE Int. Symp. on Circuits and Systems (Tutorial)* (1994 May).
- [31] M. Motamed, A. Zakhor, and S. Sanders, "Tones, Saturation, and SNR in Double-Loop Sigma-Delta Modulators," *Proc. 1993 IEEE Int. Symp. on Circuits and Systems*, pp. 1345-1348 (1993 May).
- [32] 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.).
- [33] R. Schreier, "On the Use of Chaos to Reduce Idle-Channel Tones in Sigma-Delta Modulators," *IEEE Trans. Circuits and Systems - I: Fundamental Theory and Applications*, vol. 41, pp. 539-547 (1994 Aug.).
- [34] C. Dunn and M. Sandler, "Linearising Sigma-Delta Modulators Using Dither and Chaos," *Proc. 1995 IEEE Int. Symp. on Circuits and Systems*, pp. 625-628 (1995 May).