Psychoacoustically Optimal Sigma Delta Modulation

Chris Dunn and Mark Sandler

Department of Electronic and Electrical Engineering
King's College, London

email: chris.dunn@scalatech.co.uk

Abstract

A psychoacoustically-optimal sigma-delta modulator (SDM) possesses a noise floor with a power spectral density that is invariant with input signal characteristics, and which is also minimally-audible. While SDM idle tones and noise modulation can be efficiently eliminated using dither, the noise floor can be made minimally audible by forcing the noise-shaping characteristic to follow the threshold of hearing. Such an action is possible by appropriate control of noise-shaping zero locations, and has the benefit of increasing the perceived resolution of a given modulator design. Alternatively, for a given perceived resolution, psychoacoustically-optimal zero locations allow a reduction in oversampling factor and/or modulator order. In this paper optimal zero locations and associated enhancements in perceived resolution are determined for SDM orders ranging between 1 and 8.

1 INTRODUCTION

Sigma-delta modulation (SDM) is a now commonly-used technique in high-resolution analog-to-digital (AD) and digital-to-analog (DA) conversion, principally due to a favourable insensitivity to component tolerances [1]. Fig. 1 shows an SDM system where a 1-bit quantiser is embedded within a negative feedback loop including loop filter *H*. By heavily oversampling the input signal and appropriate choice of *H* [2], [3], high resolution can be obtained in the baseband despite the coarse 1-bit quantisation.

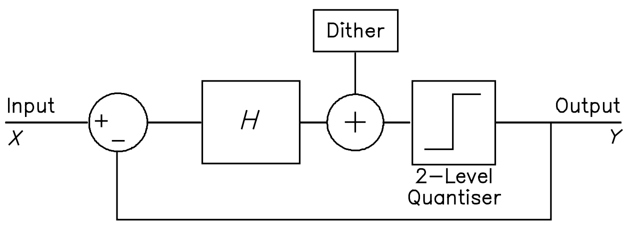

Fig. 1. General sigma-delta structure.

When sigma-delta modulators are used in audio conversion applications, psychoacoustic considerations dictate the characteristics required of an optimal modulator [4]. Firstly, the noise floor due to quantisation should possess a power spectral density within the audio band which is invariant with changes in the input signal. Secondly, the noise floor should be minimally audible. In this paper we discuss how both of these conditions can be efficiently realised.

Although standard sigma-delta modulators typically suffer from low-level idle tones and noise modulation, such artifacts can be eliminated by appropriate use of dither [5] or by making the modulator chaotic [6], [7]. While recent work has indicated that chaos may be an inefficient technique for linearising SDMs [3], [8], [9], dither can eliminate unwanted artifacts with relatively small dynamic range penalties [3], [10]. The most recent work by the authors [10] has indicated that, for modulator orders ranging between 2 and 5, signal-to-noise ratio (SNR) penalties associated with multi-level dither are approximately constant at 5 dB.

Interestingly, such a reduction in SNR is approximately equal to that paid for correctly dithering a pulse-code modulation (PCM) quantiser [11]. Our simulations further show that, for higher-order modulators, single-bit dither can be used with typically no dynamic-range penalty compared to multi-level dither signals [10]. This overcomes implementation difficulties when dithering SDM ADCs, since an analog single-bit dither source can be simply implemented as a single switched capacitor controlled from the digital domain. In this paper we extend the study of single-bit dither to SDM orders 6 and 7.

Given that appropriate steps have been taken to eliminate unwanted nonlinear artifacts, the quantisation noise floor of a sigma-delta modulator can be considered relatively benign in a psychoacoustic sense. The noise floor can be made even less intrusive by forcing the noise-shaping characteristic of the modulator to follow the threshold of hearing - that is, minimally-audible noise shaping. This technique has been explored for PCM quantisation [12], [13], and briefly considered by Jantzi et al. for sigma-delta modulation [14]. Minimally-audible noise shaping can be implemented in SDMs by controlling the zero locations in the noise-shaping response, and, in general, optimal zero locations will differ from those found in a modulator designed for maximum unweighted SNR. Thus although use of psychoacoustically-optimal zero locations will tend to improve the perceived SNR compared to a modulator designed for maximum unweighted SNR, such a move will tend to deteriorate the unweighted SNR. We determine optimal zero locations for modulator orders ranging from 1 to 8, the associated improvements in perceived resolution, and also the unweighted SNR penalties. Finally we investigate how, for a given perceived resolution, minimally-audible noise-shaping allows a reduction in SDM oversampling ratio and/or modulator order.

1 PSYCHOACOUSTICALLY-OPTIMAL NOISE SHAPING

The benefits of noise shaping PCM quantisation according to the ears sensitivity to low-level noise have been investigated in detail by Lipshitz and co-workers [11], [12], [13]. Fig. 2 illustrates a noise shaping system based around a PCM quantiser Q, where the noise-shaping loop filter *G* can be chosen to redistribute quantisation noise to regions of the audio band where it is less audible - that is, remove noise from frequencies where the ear exhibits the highest sensitivity. For a non-oversampled system and an F-weighting sensitivity curve (see below), Wannamaker [13] has determined that an optimal 9th-order FIR noise shaping filter can increase the weighted signal-to-noise ratio by 17 dB - an enhancement in perceived resolution of approximately 3 bits. Such an action incurs the cost of increased unweighted noise by 18 dB, which is acceptable for audio systems when the noise shaping is performed around a 16-bit quantiser. It should be noted that, even for higher-order noise-shaped PCM quantisers, dither is required in order to completely eliminate quantisation distortion and noise modulation from the truncated output signal.

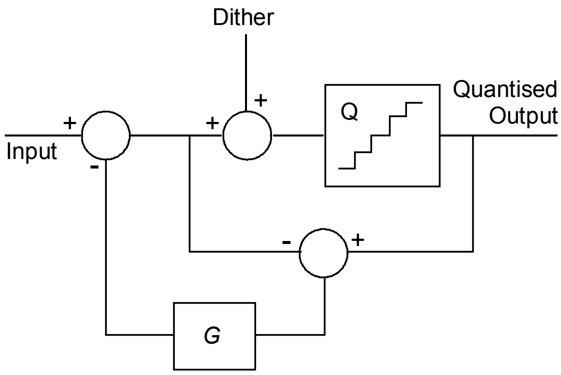

Fig. 2. PCM noise shaper.

Spectral noise shaping which follows psychoacoustic sensitivity contours is also possible with sigma-delta modulators. Referring to Fig. 1, and assuming that the quantiser can be modelled as an additive noise source *E*, the noise transfer function (NTF) relates the quantisation noise component in the output signal

Y to quantisation noise *E*,

$$NTF(z) = \left.\frac{Y(z)}{E(z)}\right|_{X(z)=0} . \tag{1}$$

This expression can also be written in terms of noise-shaping poles p_i and zeros z_i, where for an *n*th-order modulator,

$$NTF(z) = \frac{\prod_{i=1}^{n}(z - z_i)}{\prod_{i=1}^{n}(z - p_i)} = \frac{1}{1 + H(z)} . \tag{2}$$

In general p_i and z_i are selected to effect a highpass noise-shaping response such that quantisation noise is attenuated in the baseband. A popular noise-shaping characteristic is the Butterworth highpass response, where the poles are arranged in a Butterworth configuration in order to control high-frequency quantisation noise and preserve stability [15], while zeros are set to lie on the unit circle at dc frequency in the *z*-domain. We shall term such an arrangement a *DC-zero set* - examples of practical modulators designed with this approach can be found in [16], [17].

Although DC-zero sets have the advantage that modulators can be designed with simple loop filter structures [3], [18], the in-band noise can be further suppressed by setting z_i to lie on the unit circle in the *z*-domain with frequencies distributed as conjugate pairs within the baseband [18], and commercially-available coders have been designed using such an approach [15], [19]. Schreier [2] derived the following analytical expression which approximately relates the in-band (unweighted) noise power N_u^2 to zero frequencies f_i by

$$N_u^2 \approx k \int_0^{f_B} \prod_{i=1}^{n} (f - f_i)^2 \, df$$

where k = constant , (3)

f_B = upper baseband frequency .

It is a relatively straightforward task to optimise f_i such that the in-band noise power according to Eq. (3) is minimised; these optimal zero frequencies scale proportionally with the passband bandwidth f_B, but are approximately constant with changes in the oversampling factor *OSF*. We shall term a noise-shaping zero set optimised to yield minimum unweighted in-band noise power a *UOPT-zero set*; Schreier has determined these zero locations analytically for $n \le 5$, and numerically for $6 \le n \le 8$. By evaluating the integral in Eq. (3) it is possible to determine the improvement in unweighted signal-to-noise ratio over DC-zero-set modulators ($\Delta SNRU_{DC}$); Table 1 lists f_i and $\Delta SNRU_{DC}$ for modulator orders 1 to 8 when f_B = 22.05 kHz (commensurate with the compact disc sampling frequency of 44.1 kHz).

Fig. 3 shows two examples of baseband NTFs with UOPT zero sets; these responses are obtained by assuming a noise floor which follows Eq. (3) - that is, that noise-shaping poles have negligible influence in the baseband - and have been normalised for unity gain with a white noise source. Note how f_i are distributed to yield NTFs which are relatively flat in the baseband.

Table 1. UOPT noise-shaping zero locations f_i for maximum unweighted SNR; $\Delta SNRU_{DC}$ and $\Delta SNRW_{DC}$ give the unweighted and F-weighted SNR advantages respectively over a DC zero set.

Order n	f_i for f_B = 22.05 kHz, kHz	$\Delta SNRU_{DC}$ dB	$\Delta SNRW_{DC}$ dB
1	0	0.0	0.0
2	±12.7	3.5	-15.1
3	0, ±17.1	8.0	-11.5
4	±7.5, ±19.0	12.8	-12.9
5	0, ±11.8, ±20.0	17.9	-11.6
6	±5.3, ±14.6, ±20.6	23.2	-7.1
7	0, ±8.9, ±16.4, ±20.9	28.6	-9.4
8	±4.0, ±11.6, ±17.6, ±21.2	34.0	-1.1

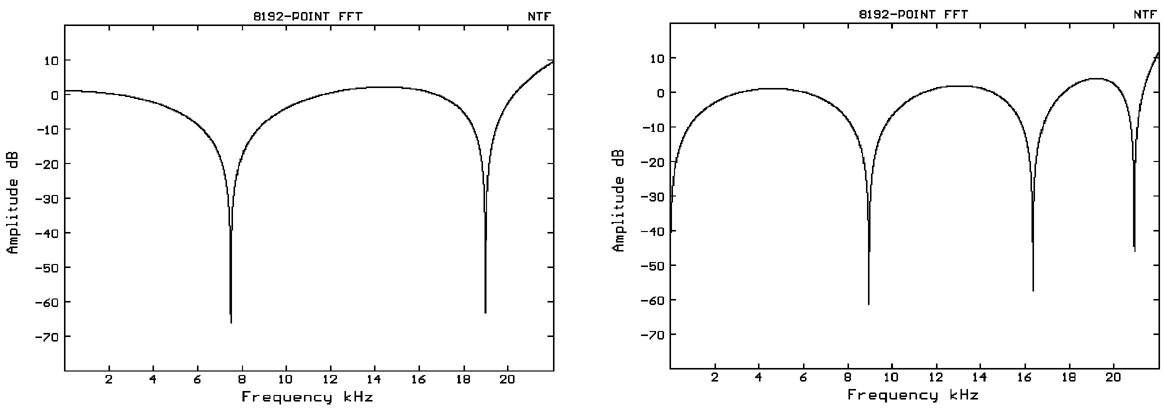

(a) (b)

Fig. 3. Unweighted NTFs for UOPT zero sets, order (a) 4 (b) 7.

So far we have considered zero locations which yield maximum unweighted SNR, but such a measure does not take into account how the ears sensitivity to low-level noise changes with frequency. In order to make an SNR measurement more psychoacoustically relevant it is possible to accommodate this nonuniform sensitivity by skewing the noise floor with a weighting function W, which yields the following analytical expression for in-band weighted noise power:

$$N_w^2 \approx k \int_0^{f_B} W(f) \prod_{i=1}^{n} (f - f_i)^2 \, df . \qquad (4)$$

Choosing an appropriate weighting function which accurately models the behaviour of an average adult hearing system is a complex undertaking. For this study we elected to use the F-weighting curve derived by Wannamaker [13], so that results could be directly compared to existing data for psychoacoustically noise-shaped PCM quantisers, although we note that there is much evidence to suggest that alternative weighting measures may be more appropriate for high-resolution audio converters [12]. The F-weighting curve (Fig. 4) is obtained by modelling ISO 15-phon equal loudness data for noise, with a free-to-diffuse field correction applied. When the weighting curve is used to evaluate modulators with UOPT zero sets (Fig. 5), it can be seen that noise in the 4 kHz region (where the ear has its highest sensitivity) dominates the weighted in-band noise power. The data in Table 1, where $\Delta SNRW_{DC}$ is the weighted SNR compared to a DC zero set of the same order, indicates that DC-zero sets yield lower weighted noise powers than UOPT sets. Thus moving zero locations away from dc to positions which minimise the unweighted noise has the effect of increasing the audibility of the noise floor.

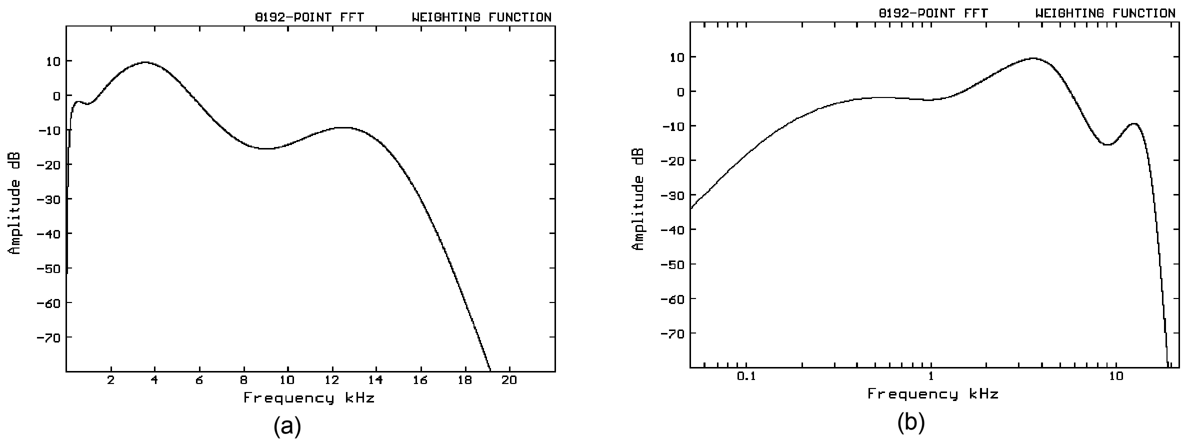

Fig. 4. F-weighting curve, indicating increased sensitivity in 4 kHz region.
(a) Linear frequency scale. (b) Logarithmic frequency scale.

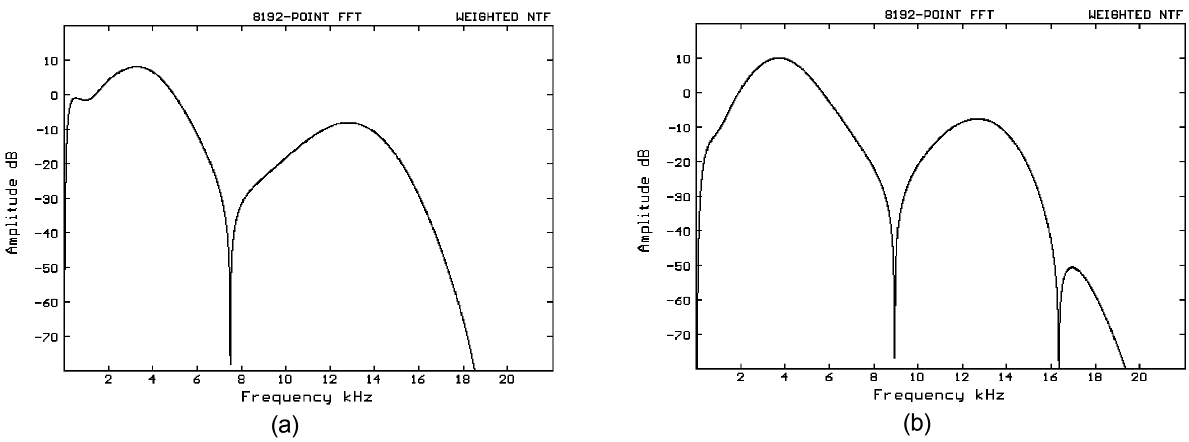

Fig. 5. F-weighted NTFs for UOPT zero sets, order (a) 4 (b) 7.

We can use Eq. (4) to yield zero sets which minimise *weighted* noise power, and hence by our adopted measure of noise audibility are psychoacoustically optimal. We term such zero sets *POPT-zero sets*. Table 2 lists POPT zero frequencies f_i for modulator orders ranging between 1 and 8; this data was obtained with f_B = 22.05 kHz [1]. Figs. 6 and 7 show unweighted and weighted NTFs for two POPT zero sets; note how these systems exhibit *weighted* NTFs which are approximately evenly distributed across the baseband.

Table 2 also lists POPT unweighted and weighted resolution relative to DC-zero sets, $\Delta SNRU_{DC}$ and $\Delta SNRW_{DC}$ respectively. It is clear that the psychoacoustically-optimal sets improve upon the DC-zero sets by both measures. The improvements in weighted SNR increase with order, and, recalling the 17 dB improvement in weighted SNR for a noise-shaped PCM converter [13], are relatively large for higher-order modulators. The sizeable improvements in weighted SNR that occur when moving from UOPT to POPT zeros suggest either a higher perceived resolution or, for a given perceived resolution, the possibility of reduced oversampling ratios and/or modulator order.

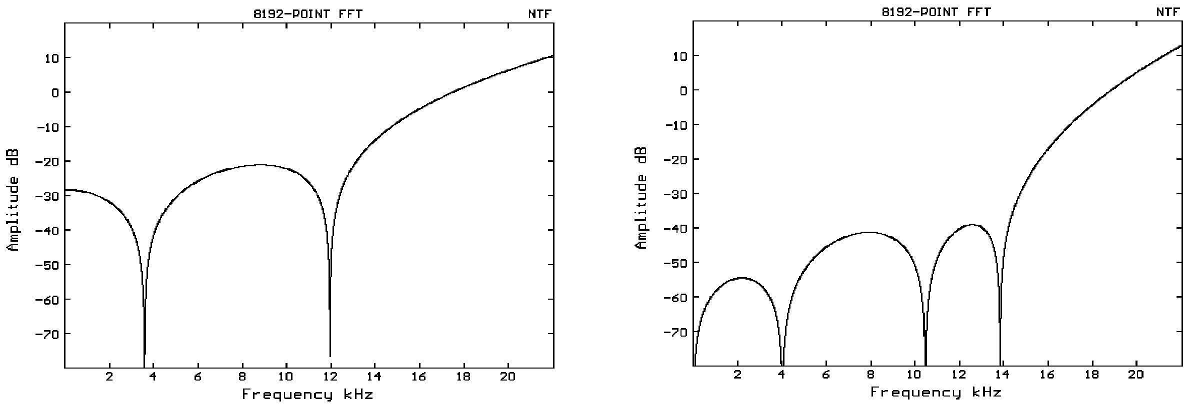

(a) (b)

Fig. 6. Unweighted NTFs for POPT zero sets, order (a) 4 (b) 7.

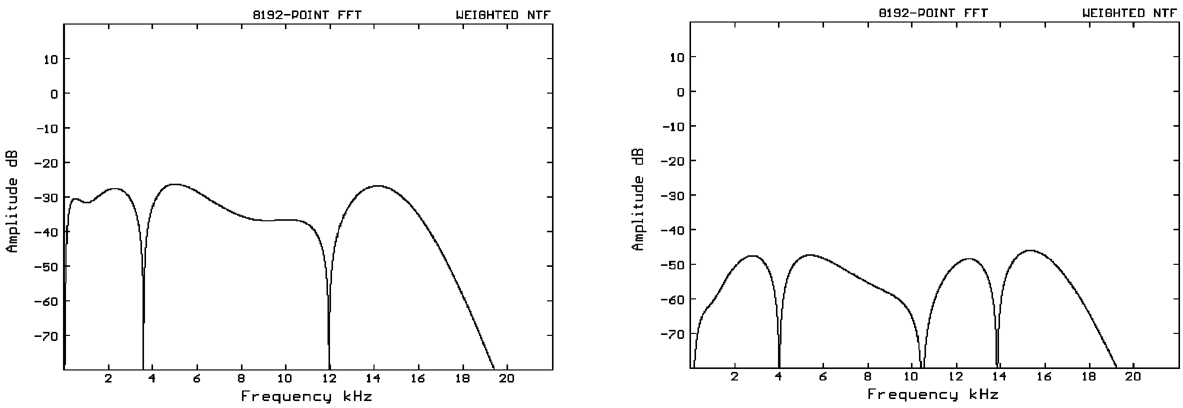

(a) (b)

Fig. 7. F-weighted NTFs for POPT zero sets, order (a) 4 (b) 7.

[1] In general the optimal zero locations will change as f_B changes; however, as for UOPT zeros, POPT zero frequencies remain constant with changes in oversampling ratio.

Table 2. Unweighted and weighted SNR advantages $\Delta SNRU_{DC}$ and $\Delta SNRW_{DC}$ of POPT zero locations relative to DC-zero sets.

Order n	f_i for f_B = 22.05 kHz, kHz	$\Delta SNRU_{DC}$ dB	$\Delta SNRW_{DC}$ dB
1	0	0.0	0.0
2	±4.014	0.5	2.1
3	0, ±6.443	1.1	2.0
4	±3.590, ±11.954	4.3	9.2
5	0, ±4.308, ±12.959	5.0	14.2
6	±3.325, ±7.078, ±13.389	6.2	17.9
7	0, ±4.017, ±10.471, ±13.842	8.0	22.3
8	±2.933, ±5.167, ±12.012, ±14.381	9.8	28.0

Up to this point our theoretical analysis has been based upon approximate expressions for in-band noise powers [Eqs. (3) and (4)], assuming large oversampling ratios. In an attempt to quantify the changes in SNR that occur with different zero sets in real modulators, a number of discrete-time systems were simulated with both UOPT and POPT zero frequencies. Two different oversampling ratios were examined - *OSF* = 16 and *OSF* = 64, the latter typical of contemporary higher-order SDM design [15], [19]. Each modulator simulated with a given zero set was optimised for maximum SNR following the procedure detailed in [3]. No dither signals were used for these experiments.

The simulation results are presented in Table 3, where $\Delta SNRU_U$ and $\Delta SNRW_U$ detail the unweighted and weighted SNRs of the POPT systems relative to UOPT modulators; since UOPT zero sets are optimised to yield maximum unweighted SNR, $\Delta SNRU_U$ will always be negative. It can be seen that the theoretical improvements in weighted SNR are largely preserved in the simulated systems at the higher oversampling factor. However, some shortfall in $\Delta SNRW_U$ is seen at *OSF* = 16, particularly for higher-order modulators. Further simulations (not detailed here) suggest that this behaviour is due to a worsening of stability experienced with POPT zero locations at the lower oversampling factor, which in turn may be due to the proximity of noise-shaping pole frequencies to the baseband in these systems. The unweighted SNR penalties in moving from POPT to UOPT sets also agree fairly well with the theoretical figures; again, there is a small additional penalty at the lower oversampling ratio.

2 SDM DITHER

Early work on lower-order sigma-delta systems identified the presence of low-level nonlinear artifacts which can be disturbing when these systems are used in audio conversion applications [20], [21], while more recently such nonlinearity has been identified in higher-order systems up to and including 5th order [5]. One way in which sigma-delta nonlinearity manifests is the presence of idle tones in the noise floor for constant (dc) input signals. Ledzius and Irwin [22] have identified a relationship between idle-tone frequencies and the dc input level x_{DC} applied to second-order modulators; for an SDM with quantiser levels ±1, and sampling frequency f_s, a strong idle tone appears at a frequency given by

$$f_{IT} = x_{DC}\, f_S = 2\, x_{DC}\, f_B\, \mathrm{OSF}. \quad (5)$$

Table 3. Theoretical and simulated SNR advantages/penalties for undithered POPT systems relative to UOPT zero sets.

Order n	ΔSNR	Theoretical dB	Simulated OSF = 16 dB	Simulated OSF = 64 dB
2	W_U	17.2	19.4	20.1
	U_U	-3.0	-2.9	-1.6
3	W_U	13.5	14.8	15.0
	U_U	-6.9	-6.5	-6.2
4	W_U	22.1	21.5	21.9
	U_U	-8.5	-9.3	-8.9
5	W_U	25.8	24.3	27.6
	U_U	-12.9	-13.8	-10.4
6	W_U	25.0	21.5	24.2
	U_U	-17.0	-19.9	-16.7
7	W_U	31.7	26.3	30.7
	U_U	-20.6	-24.9	-20.3

By careful simulation it is possible to show that these idle tones also occur in higher-order systems; Fig. 8 shows the noise floor of a 6th-order system with UOPT zeros as a dc input signal is swept from 0 to 1/256. For this 64X-oversampling system, an idle tone can clearly be distinguished from the random component of the noise floor, traversing the plot according to Eq. (5). It is also possible to show that idle tones corrupt the noise floor with input signals other than dc [3], [9].

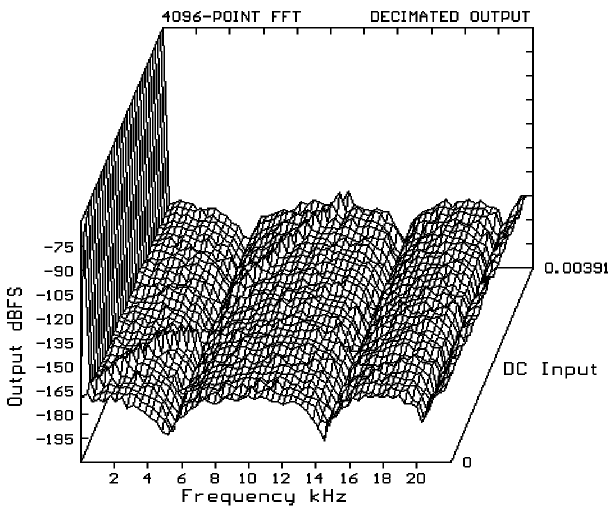

Fig. 8. DC-input sweep for an undithered sixth-order modulator where OSF = 64. An idle tone can clearly be seen traversing the plot.

A problem related to the existence of idle tones is the dependence of baseband noise power upon input signal amplitude - that is, noise modulation - which can also be seen in higher-order systems. Our investigations suggest that SDMs of order up to and including 7th suffer from such artifacts. Fig. 9 charts both noise modulation and idle-tone amplitudes in SDMs with UOPT zero locations as modulator order changes. Noise modulation in each simulated modulator was computed by determining total baseband noise power as the amplitude of a 1 kHz sinewave was ramped across the dynamic range of the modulator, whereas idle tone amplitudes were determined relative to the random component of the noise floor with a dc input signal equal to 1/(4*OSF*). For both graphs it is clear that although increasing the SDM order reduces the magnitude of nonlinear errors, these problems remain to some extent in the higher-order systems.

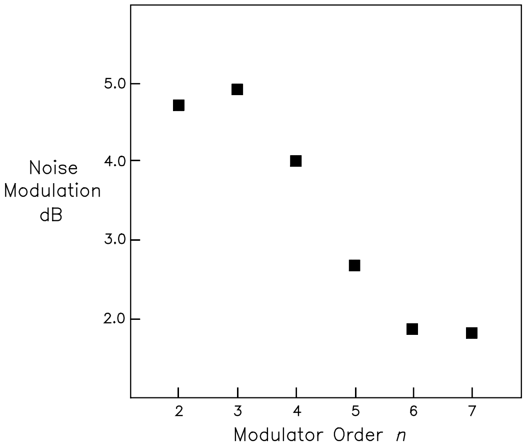

(a)

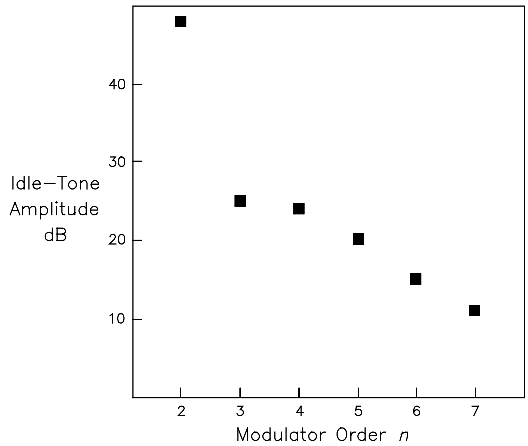

(b)

Fig. 9. Variation of nonlinearity with modulator order. (a) Noise modulation. (b) Idle-tone amplitude relative to noise floor.

For high-quality audio applications the noise floor of a conversion process is ideally invariant with input signal characteristics [4]. In practical terms this implies that noise modulation should be less than 1 dB across the dynamic range of the system, with no idle-tone components appearing in the noise floor. When these conditions are satisfied then the noise floor can be considered psychoacoustically benign, and the modulator "linearised". For sigma-delta systems linearisation can be achieved by dithering the modulator with random noise injected at the input to the quantiser (Fig. 1). Earlier work by the authors determined multi-level dither amplitudes required for linearisation, and associated dynamic-range penalties, for UOPT modulators up to and including fourth order [3], [9]. More recently we have investigated the use of single-bit (that is, two-level) dither as an efficient signal for linearising SDMs [10]. In this section we extend the work on single-bit dither to modulator orders 6 and 7, and to modulators with POPT zero sets.

Table 4 lists the lowest (optimal [2]) single-bit dither amplitude sufficient to linearise modulators with UOPT and POPT zero sets at two oversampling factors. For the higher-order UOPT systems the optimal dither amplitudes remain approximately constant at 0.13 (for a 1-bit quantiser with output levels ±1). For lower-order systems and modulators with POPT zeros the optimal dither amplitude increases, suggesting that these systems are more difficult to linearise. This finding is supported when examining the simulation data for SNR penalties with optimal dithering (Table 5), where the cost in dynamic range is lowest for modulators of moderate order with UOPT zero locations, but rises at lower and higher orders and for psychoacoustically-optimal noise-shaped systems. In general the dynamic-range penalties for optimal dithering are of the same order of magnitude as that paid for correctly dithering PCM quantisers (4.8 dB).

[2] Optimal in the sense that linearisation is achieved with the smallest SNR penalty relative to an undithered modulator.

Table 4. Single-bit dither amplitudes required for linearisation.

n	OSF = 16		OSF = 64	
	UOPT	POPT	UOPT	POPT
2	0.56	0.79	0.71	0.79
3	0.14	0.28	0.14	0.35
4	0.14	0.18	0.14	0.25
5	0.13	0.13	0.13	0.20
6	0.13	0.13	0.14	0.18
7	0.13	0.14	0.13	0.16

Table 5. SNR penalties for linearisation using single-bit dither.

n	OSF = 16		OSF = 64	
	UOPT	POPT	UOPT	POPT
2	10.1	13.7	10.1	11.3
3	3.2	8.8	3.3	9.7
4	3.9	5.8	5.4	9.3
5	4.5	7.7	3.3	9.2
6	5.3	5.9	5.9	9.2
7	7.4	7.7	6.7	9.7

Higher SNR penalties in the POPT systems will tend to reduce some of the perceived dynamic range benefit for selecting psychoacoustically-appropriate noise transfer functions, and also increase the unweighted SNR penalty associated with such a move. Table 6 compares the weighted and unweighted resolution of optimally-dithered systems with POPT zero locations relative to UOPT systems. It is clear that the advantage of psychoacoustically-appropriate noise shaping is reduced somewhat for the dithered systems compared to the undithered systems (Table 3) - for example, for the 16X-oversampled 3rd-order system, $\Delta SNRW_U$ is reduced from 14.8 dB to 7.9 dB, while the unweighted SNR penalty increases from 6.5 dB to 12.6 dB. Nevertheless, significant advantages remain for adopting POPT zeros with dithered higher-order systems - for example, $\Delta SNRW_U$ = 27.8 dB for the 7th-order 64X-oversampled system.

Table 6. Theoretical and simulated SNR advantages/penalties for optimally-dithered POPT modulators relative to optimally-dithered UOPT systems.

n	ΔSNR	Theoretical dB	Simulated OSF = 16 dB	Simulated OSF = 64 dB
2	W_U	17.2	14.5	19.2
	U_U	-3.0	-7.1	-3.2
3	W_U	13.5	7.9	8.1
	U_U	-6.9	-12.6	-12.6
4	W_U	22.1	19.5	17.8
	U_U	-8.5	-11.3	-12.5
5	W_U	25.8	20.9	21.8
	U_U	-12.9	-17.2	-16.0
6	W_U	25.0	20.8	21.0
	U_U	-17.0	-20.6	-20.2
7	W_U	31.7	26.1	27.8
	U_U	-20.6	-25.0	-23.4

Table 7. Comparison of gain in weighted SNR obtained by moving from UOPT to POPT zero sets, against gain in SNR obtained by doubling *OSF*.

Order n	Theoretical $\Delta SNRW_U$ dB	Typical ΔSNR obtained for doubling *OSF* dB
2	17.2	14
3	13.5	21
4	22.1	26
5	25.8	30
6	25.0	36
7	31.7	40

Table 8. Comparison of gain in weighted SNR obtained by moving from UOPT to POPT zero sets, against gain in SNR obtained by increasing modulator order at *OSF* = 64.

Order n	Theoretical $\Delta SNRW_U$ dB	ΔSNR for $\Delta n = 1$ dB	ΔSNR for $\Delta n = 2$ dB
2	17.2	21.5	36.1
3	13.5	14.6	26.0
4	22.1	11.5	23.5
5	25.8	12.0	22.0
6	25.0	10.0	-

Finally in this section we briefly discuss the possibilities of reduced oversampling ratio and/or modulator order allowed by using psychoacoustically-optimal zero locations. For a given modulator resolution, reductions in oversampling ratio can yield significant improvements in practical converter performance, where the reduction in clock speed will tend to reduce artifacts due to unwanted coupling mechanisms [23]. Table 7 compares the theoretical gains in weighted SNR for POPT zero locations relative to UOPT zero sets, against the gain in resolution that typically occurs when the oversampling factor is doubled. [3] Although the SNR gained when doubling *OSF* is in most cases greater than that achieved in moving to POPT zeros, the results are similar in magnitude - for example, 22 dB for a fourth-order modulator with POPT zeros compared to 26 dB for doubling *OSF*.

Table 8 compares the gains in resolution achieved by moving to POPT zeros against that obtained by increasing the modulator order by 1 and 2, this information acquired by simulating 64X-oversampling modulators with UOPT-zero locations and averaging results from dithered and undithered systems. [4] For

[3] These results were obtained by simulating UOPT systems with oversampling ratios ranging between 16 and 64.

[4] Note that the gain in SNR obtained by incrementing the modulator order will increase as the oversampling ratio is increased.

$n \geq 4$ the advantage in perceived resolution from moving from UOPT to POPT zero locations is greater than that achieved by increasing the order by 1 with UOPT zeros, while for $n \geq 5$ the advantage in implementing psychoacoustically-optimal zero locations is greater than incrementing n by 2.

3 RESULTS

In this section we give some examples of the performance available from modulators designed using psychoacoustically-appropriate noise-shaping functions, commencing with some simulation results obtained from higher-order systems operating at low oversampling factors.

While early SDM designs achieved high resolution by using high oversampling ratios with lower-order noise shaping - for example, Naus et al. used 256X oversampling with a second-order noise-shaping characteristic [16] - more recently higher-order systems have been described in the literature which allow reductions in oversampling ratio whilst maintaining high resolution. Examples are provided by Sooch et al. [24], who obtained 120 dB SNR from a 5th-order 64X-oversampling modulator, and more recently Risbo [25], who described an experimental 8th-order 32X-oversampling system which achieved a maximum simulated SNR of 97 dB. Psychoacoustically-optimal noise shaping allows this trend to be continued, where for a given perceived resolution further reductions in oversampling factor are possible.

Fig. 10 shows simulation results from an undithered 7th-order 16X-oversampling system with UOPT zeros. 62.6 dB peak unweighted SNR is achieved, while the peak weighted SNR is 61.6 dB - approximately equivalent to 10.7-bit dithered PCM resolution. Note that the weighted SNR figure is dominated by noise in the 4 kHz region [Fig. 10(b)], while idle tones are apparent in the dc-input-sweep plot of this undithered modulator [Fig. 10(c)]. This performance can be contrasted with a 7th-order dithered system operating at the same oversampling ratio, where POPT zero locations effect psychoacoustically-optimal noise shaping. Fig. 11(a) shows the output spectrum with a full-scale 1 kHz input signal, the effective suppression of noise in the 4 kHz region achieving a weighted SNR of 80.2 dB - equivalent to 13.8-bit PCM resolution. It is quite possible that alternative weighting measures [12] would yield closer to 15-bit perceived resolution for this converter - an interesting result given the very low 16X-oversampling factor. The absence of idle tones in the output spectrum for a dc-input sweep [Fig. 11(c)] indicates the effectiveness of the dithering implemented in this system. Of course use of POPT zeros results in an unweighted SNR penalty relative to a UOPT design, and this is seen in the low unweighted SNR figure of only 30.3 dB for this POPT system; in fact, such a high unweighted noise level would be unacceptable for an audio converter, but this example does serve to illustrate the potential benefits of psychoacoustically-optimal noise shaping.

In Sec. 2 we studied the possibility of psychoacoustically-optimal zero locations allowing a reduction in modulator order whilst maintaining perceived resolution. A brief example is now presented to illustrate this principle. Fig. 12 shows unweighted and F-weighted output spectra when a full-scale 1 kHz signal is input to an undithered 4th-order SDM with UOPT zeros operating at 64X-oversampling; the unweighted and weighted SNRs are 109.1 dB and 108.6 dB respectively, equivalent to 18.6-bit PCM resolution - and suitable for high-resolution audio converter applications. Fig. 13 shows simulation results for an alternative system operating at the same oversampling factor but using a POPT zero set; the modulator order has been reduced by 1, but the use of psychoacoustically-optimal zero locations yields a peak weighted SNR (109.7 dB) which exceeds that of the UOPT example. The use of POPT zeros carries the penalty of decreasing the unweighted SNR to 87 dB, although this would not represent a problem in audio applications.

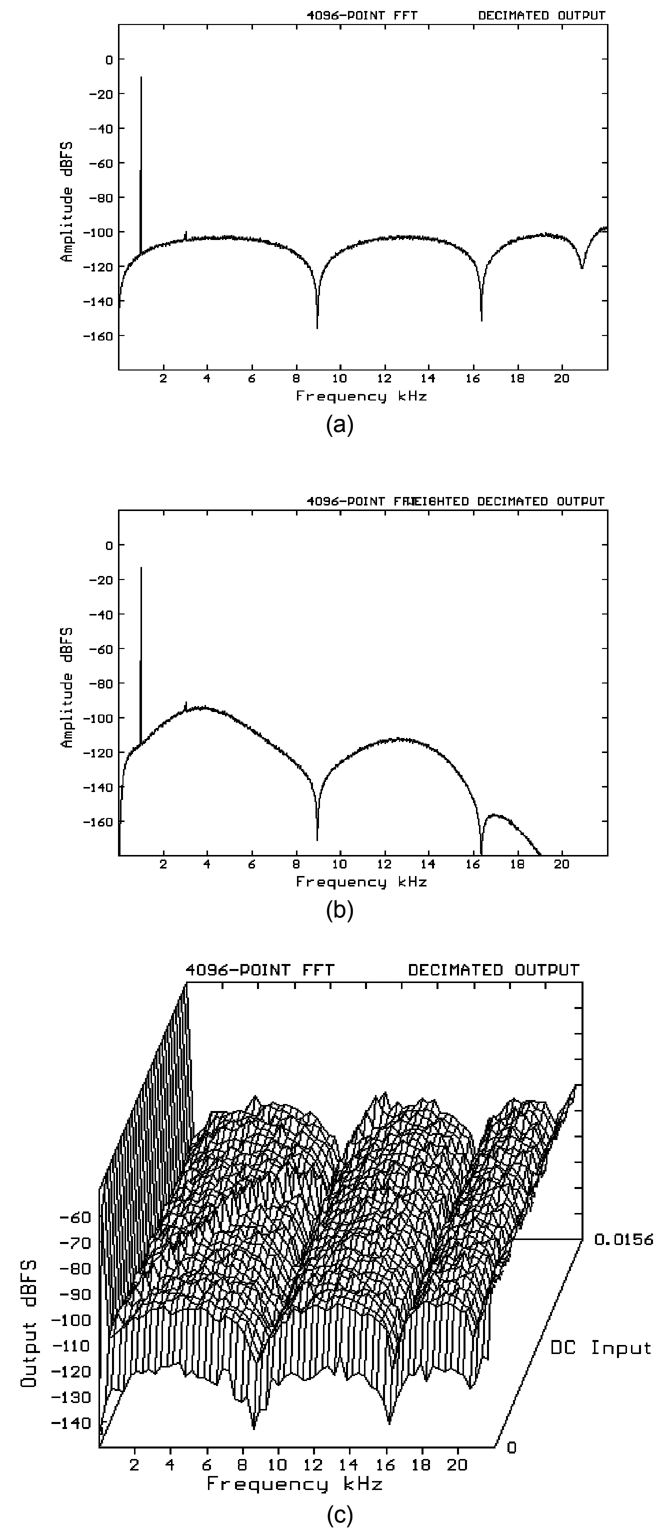

Fig. 10. 7th-order 16X-oversampling undithered modulator with UOPT zero locations. (a) Unweighted output spectrum at full-scale input. (b) F-weighted output spectrum at full-scale input. (c) Unweighted output spectrum for DC input sweep.

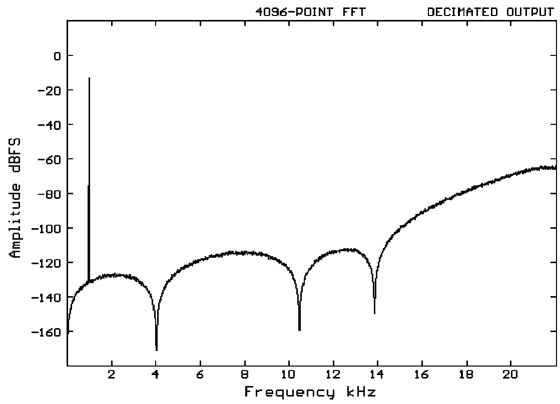

(a)

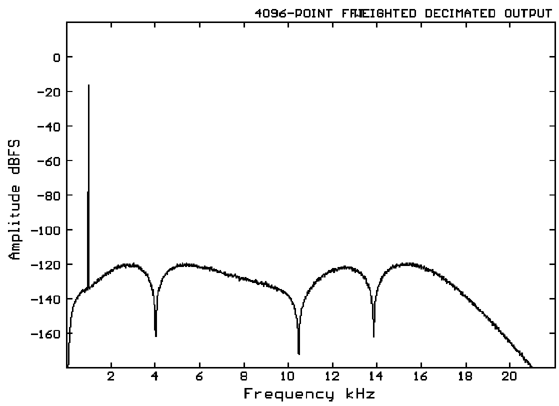

(b)

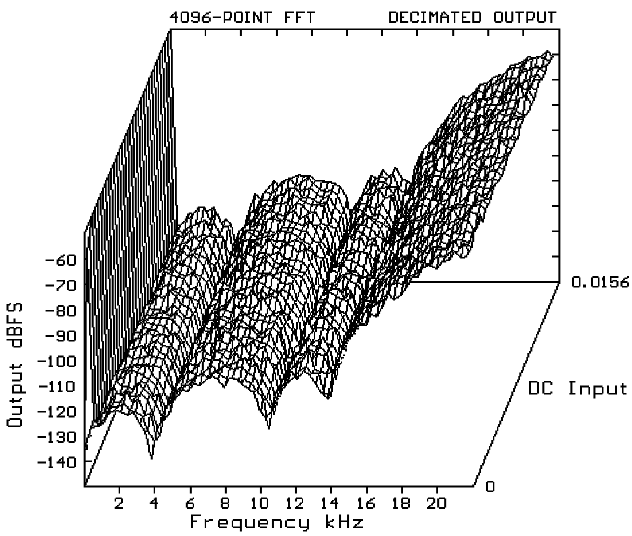

(c)

Fig. 11. 7th-order 16X-oversampling optimally-dithered modulator with POPT zero locations. (a) Unweighted output spectrum at full-scale input. (b) F-weighted output spectrum at full-scale input. (c) Unweighted output spectrum for dc input sweep.

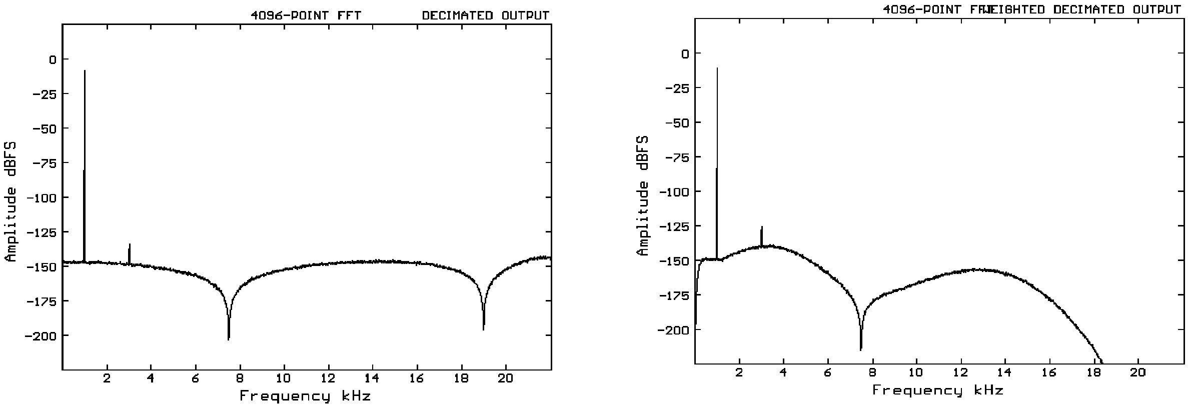

(a) (b)

Fig. 12. 4th-order 64X-oversampling undithered modulator with UOPT zero locations. (a) Unweighted output spectrum at full-scale input. (b) F-weighted output spectrum at full-scale input.

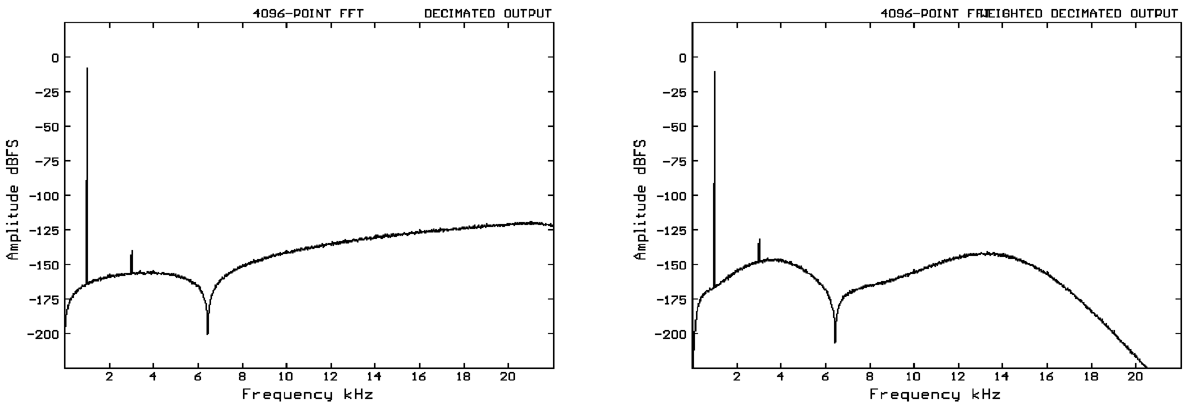

(a) (b)

Fig. 13. 3rd-order 64X-oversampling undithered modulator with POPT zero locations. (a) Unweighted output spectrum at full-scale input. (b) F-weighted output spectrum at full-scale input.

4 CONCLUSIONS

We have investigated the characteristics desired of sigma-delta modulators used for high-resolution audio converters. Two attributes were identified - that the noise floor should be free of idle tones and other low-level artifacts, and that the noise floor should be minimally audible. We have shown that SDMs can be linearised using dither, and that careful choice of noise-shaping zero locations can effect psychoacoustically-appropriate shaping of the modulator noise floor.

Conventional sigma-delta modulators place noise-shaping zeros either at dc (a DC-zero set) or distributed across the baseband to minimise unweighted baseband noise power (a UOPT-zero set). However, a more psychoacoustically-appropriate measure of modulator resolution is to weight the converter error with a function which approximates the ears sensitivity to low-level noise. Using an F-weighting function it is found that UOPT-zero sets actually perform less well than DC zeros - the difference in weighted resolution can be as high as 15 dB, equivalent to 2.5-bits PCM resolution.

An alternative design strategy is to place noise-shaping zeros at frequencies which minimise the weighted noise, and we term such locations POPT-zero sets. Using an approximate measure of baseband noise power, POPT zero frequencies were determined for modulator order *n* ranging between 1 and 8. POPT-zero sets improve both weighted and unweighted resolution over DC-zero sets. Compared to UOPT-zero sets, the weighted resolution advantage of psychoacoustically-optimal zero positions increases with order, ranging between 14 dB (*n* = 2) and 32 dB (*n* = 7); these figures are relatively large when compared to psychoacoustically-optimal PCM noise-shaping. Increases in perceived resolution with POPT modulators are gained at the expense of unweighted SNR, where the penalty increases with modulator order, ranging between 3 dB at *n* = 2, and 24 dB for *n* = 8. Simulation of undithered systems illustrated that theoretical gains in weighted resolution are largely preserved when moving from UOPT to POPT zero locations, although there is some shortfall from the theoretical figures when very low oversampling factors are used.

Results were presented to show that higher-order sigma-delta modulators suffer from nonlinear artifacts such as idle tones and noise modulation, although the magnitude of such problems is reduced compared to lower-order systems. Such unwanted errors can be eliminated using dither, and in this paper we examined the use of single-bit dither which is easily implemented in both SDM DACs and ADCs. Optimal dither amplitudes and associated reductions in SNR were determined for systems with UOPT and POPT zero locations, with *n* ranging between 2 and 7. It was found that the SNR penalty for dithering is lowest for UOPT systems of moderate order, while the reduction in dynamic range increases when dither is used to linearise lower- and higher-order systems, and also those with POPT zero locations. For example, the SNR penalty for dithering a 3rd-order UOPT system is 3 dB, while for a 2nd-order modulator with POPT zero locations the penalty increases to 11 dB. The increased SNR penalty for dithering POPT modulators tends to reduce somewhat the weighted resolution advantage of these devices over systems with UOPT-zero locations, although the advantage remains significant at higher orders - for example, 26 dB for a dithered 7th-order system.

The improvements in weighted SNR achieved with POPT-zero sets allow either an increase in perceived resolution or, for a given perceived resolution, a reduction in oversampling ratio and/or modulator order. The weighted SNR advantage in moving from UOPT to POPT zero sets is similar to that achieved by doubling the oversampling ratio. Similarly, at *OSF* = 64, POPT zero sets allow a reduction in modulator order of 1 or 2. Finally, simulation examples were presented to illustrate the relatively high perceived resolution obtainable from modulators with POPT zero sets operating at very-low oversampling ratios.

5 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.

6 REFERENCES

[1] J. C. Candy and G. C. Temes, "Oversampling Methods for A/D and D/A Conversion," in *Oversampling Delta-Sigma Converters*, Eds. J. C. Candy and G. C. Temes, IEEE Press (1992).

[2] 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.).

[3] 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, *J. Audio Eng. Soc. (Abstracts)*, vol. 42, p. 1066 (1994 Dec.), preprint 3926.

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

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

[6] 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).

[7] 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.).

[8] R. Schreier, "On the Use of Chaos to Reduce Idle-Channel Tones in Delta-Sigma Modulators," *IEEE Trans. Circuits and Systems - I: Fundamental Theory and Applications*, vol. 41, pp. 539-547 (1994 Aug.).

[9] 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).

[10] C. Dunn and M. Sandler, "Efficient Linearisation of Sigma-Delta Modulators Using Single-Bit Dither," *Elec. Let.*, vol. 31, pp. 941-942 (1995 June).

[11] 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).

[12] S. P. Lipshitz, J. Vanderkooy, and R. A. Wannamaker, "Minimally Audible Noise-Shaping," *J. Audio Eng. Soc.*, vol. 39, pp. 836-852 (1991 Nov.).

[13] R. A. Wannamaker, "Psychoacoustically Optimal Noise Shaping," *J. Audio Eng. Soc.*, vol. 40, pp. 611-620 (1992 July/Aug.).

[14] S. Jantzi, C. Ouslis, and S. Sedra, "Transfer Function Design for Sigma-Delta Converters," *Proc. 1994 IEEE Int. Symp. on Circuits and Systems*, pp. 433-436 (1994 May).

[15] R. W. Adams, P. F. Ferguson, A. Ganesan, S. Vincelette, A. Volpe, and R. Libert, "Theory and Practical Implementation of a Fifth-Order Sigma-Delta A/D Converter," *J. Audio Eng. Soc.*, vol. 39, pp. 515-528 (1991 July/Aug.).

[16] P. J. A. Naus, E. C. Dijkmans, E. F. Stikvoort, A. J. McNight, D. A. Holland, and W. Bradinal, "A CMOS Stereo 16-Bit Converter for Digital Audio," *IEEE J. Solid-State Circuits*, vol. SC-22, pp. 390-395 (1987 June).

[17] S. R. Norsworthy, I. G. Post, and H. S. Fetterman, "A 14-Bit 80 kHz Sigma-Delta A/D Converter: Modelling, Design and Performance Evaluation," *IEEE J. Solid-State Circuits*, vol. 24, pp. 256-266 (1989 April).

[18] 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.).

[19] D. R. Welland, B. P. del Signore, E. J. Swanson, T. Tanaka, K. Hamashita, S. Hara, and K. Takasuka, "A Stereo 16-Bit Delta-Sigma A/D Converter for Digital Audio," *J. Audio Eng. Soc*., vol. 37, pp. 476-486 (1989 June).

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

[21] J. C. Candy and O. J. Benjamin, "The Structure of Quantisation Noise from Sigma-Delta Modulation," *IEEE Trans. Commun.*, vol. COM-29, pp. 1316-1323 (1981 Sept.).

[22] 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. 1304-1307 (1992 May).

[23] 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.).

[24] 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. 31, p. 994 (1991 Dec.).

[25] L. Risbo, "Field-Programmable Gate-Array-Based 32-Times Oversampling Eighth-Order Sigma-Delta Audio DAC," presented at the 96th Convention of the Audio Engineering Society, *J. Audio Eng. Soc. (Abstracts)*, vol. 42, p. 398 (1994 May).