The benefits of multibit chaotic sigma delta modulation

J. D. Reiss[a] and M. B. Sandler
Department of Electronic Engineering, King's College, University of London, Strand, London WC2R 2LS, United Kingdom

(Received 4 October 2000; accepted 14 March 2001; published 7 May 2001)

Sigma delta modulation is a popular technique for high-resolution analog-to-digital conversion and digital-to-analog conversion. We investigate chaotic phenomena in multibit first-order sigma–delta modulators. Particular attention is placed on the occurrence of periodic orbits or limit cycles. These may result in idle tones audible to the listener when sigma-delta modulation is used for audio signal processing. One suggested method of eliminating idle tones is the operation of a sigma delta modulator in the chaotic regime. Unfortunately, chaotic modulation of a first order sigma delta modulator is a poor system for signal processing. We show that minor variations on a traditional first order sigma–delta modulator, together with a multibit implementation, may be used to produce an effective, stable chaotic modulator that accurately encodes the input and helps remove the presence of idle tones. © *2001 American Institute of Physics.* [DOI: 10.1063/1.1371284]

Conventional analog to digital and digital to analog converters are based on the linear, multibit pulse code modulation (PCM) format. They require high-precision analog circuits and they are vulnerable to noise and interference. In recent years, the consumer audio industry has moved towards oversampled nonlinear converters for many applications. An important oversampling A-D or D-A conversion strategy now employed is the sigma delta modulator. In sigma delta converters the signal is sampled at a high-sampling frequency and converted to a low-bit binary output. They are cheaper to manufacture than PCM converters, consume less power, and operate well at the voltage range used in battery-powered audio equipment. Thus sigma delta modulators are used in the digital to analog converters of many compact disc players and in the audio processing of many wireless communication systems, such as cellular phone technology. In addition, the sigma–delta bitstream format is under consideration for the mastering and archiving of audio recordings. Unfortunately, sigma delta modulators are susceptible to limit cycle oscillations which may produce audible idle tones that are not present in the input signal. Operation of sigma delta modulators in the chaotic regime has been proposed as a method of avoiding limit cycle oscillations. However, chaotic modulation may be unstable and can result in inaccurate output. We look at a variety of different ways to operate a sigma delta modulator chaotically. By investigating the nonlinear dynamics of these systems, we are able to show that a new form of chaotic sigma delta modulation may be used which is accurate, stable, and avoids the presence of idle tones.

INTRODUCTION

Sigma delta (or delta sigma) modulation is a popular method for high-resolution A-D and D-A converters. It is frequently used in audio processing and has a wide range of applications. Sigma–delta modulators operate using a tradeoff between oversampling and low resolution quantization. That is, a signal is sampled at much higher than the Nyquist frequency, typically with one bit quantization, so that the signal may be effectively quantized with a resolution on the order of 14–20 bits.[1] Recent work has concentrated on tone suppression,[2–4] multibit modulation,[5] and chaotic modulation.[6–9] In this paper, we investigate the behavior of chaotic modulators.

The simplest, first-order sigma-delta modulator consists of a 1-bit quantizer embedded in a negative feedback loop which also contains a discrete-time integrator, as depicted in Fig. 1(a). The input to the modulator is sampled at a frequency higher than the Nyquist frequency and is converted into a binary output. The system may be represented by the map[10]

$$U_n = \alpha U_{n-1} + X_{n-1} - Q(U_{n-1}), \tag{1}$$

where X represents the input signal, bounded by -1 and $+1$, and Q is the quantizer

$$Q(u) = \begin{cases} 1 & \text{if} \quad u \geqslant 0 \\ -1 & \text{if} \quad u < 0 \end{cases}. \tag{2}$$

In this representation, the output $Q(U_n)$ represents the quantization of input X_{n-1}. The initial conditions, X_0 and U_0 are typically set to 0. On average, the quantized output will be approximately equal to the input. If $\alpha = 1$, then this system works by quantizing the difference between the input and the accumulated error. The operation of such a first-order sigma–delta modulator is depicted in Fig. 2. The input is a 1.5 kHz sine wave with amplitude 0.8 sampled at a frequency of 256 kHz (these parameters were chosen to accentuate the behavior of the modulator). Typically, the integra-

[a]Electronic mail: josh.reiss@kcl.ac.uk

 © 2001 American Institute of Physics

Downloaded 10 Jan 2002 to 192.58.150.40. Redistribution subject to AIP license or copyright, see http://ojps.aip.org/chaos/chocr.jsp

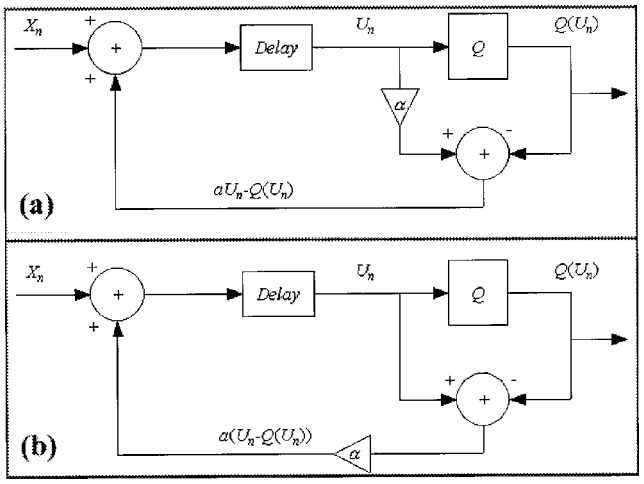

FIG. 1. Block diagrams for the two systems. In (a), gain is applied just to the integrator output. In (b), gain is applied to the quantizer error, that is, the difference between the integrator output and the quantizer output.

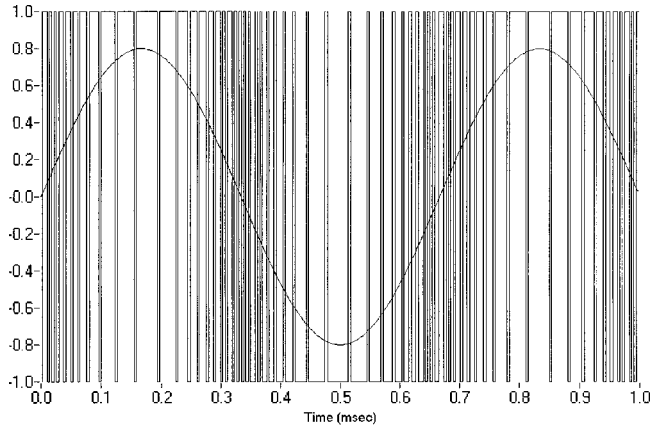

FIG. 2. A 1.5 kHz sine wave with an amplitude of 0.8 is sampled at a frequency of 256 kHz. The input sine wave and the quantized output of the sigma delta modulator are depicted.

tor leaks due to finite operational amplifier gain, which is represented by $\alpha<1$. If $\alpha>1$, then the modulator may behave chaotically for constant input. Thus when the error grows sufficiently large, the quantizer will flip in order to reduce the error.

If a gain is added to the quantization error (difference between integrator output and quantized output), as opposed to the integrator output, then the difference equation describing this modified sigma delta modulator takes the form

$$U_n = X_{n-1} + \alpha(U_{n-1} - Q(U_{n-1})). \tag{3}$$

This system is depicted in Fig. 1(b). It is relatively simple to implement in a circuit, and still accomplishes the goals of sigma delta modulation.

An alternative representation of (3) is found by defining $V_n = U_n/\alpha$ and $Y_n = X_n/\alpha$. Hence

$$V_n = \alpha V_{n-1} + Y_{n-1} - Q(\alpha V_{n-1}). \tag{4}$$

This allows the gain to be applied only to the integrator output and to the input signal. (This representation was suggested by Dr. Orla Feely of the University College Dublin.) Thus no gain needs to be applied directly to the quantization. In the case of a single-bit quantizer, (4) has the same functional form as (1).

In this work, we consider chaotic modulators where a gain term multiplies either the integrator output (1) or the error term (3). In particular, we consider whether either form of chaotic modulation is an effective means of idle tone prevention. We demonstrate that for the case of gain applied to integrator output, although an implementation of a chaotic multibit modulator may lead to idle tone suppression, it may not be practical. This is because in many cases, the output of a chaotic modulator does not effectively approximate the input.

THE SYSTEM

A multibit implementation of either (1) and (3) may offer increased resolution in the quantization. Rather than quantizing the output into -1 and 1, the output can instead assume a range of discrete values. For an n bit first-order

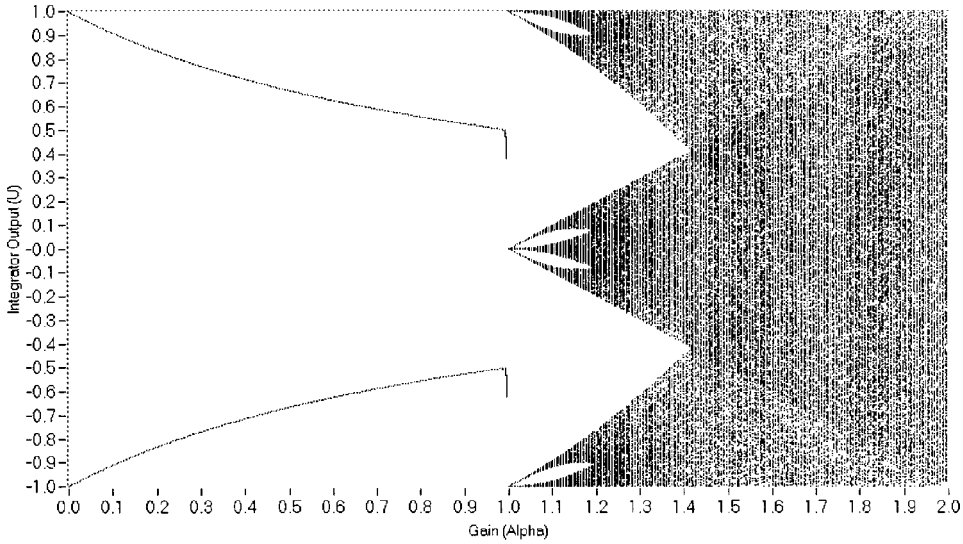

FIG. 3. Bifurcation diagram of a first order, one bit sigma delta modulator with 0 input and gain applied to the integrator output (System 1).

Downloaded 10 Jan 2002 to 192.58.150.40. Redistribution subject to AIP license or copyright, see http://ojps.aip.org/chaos/chocr.jsp

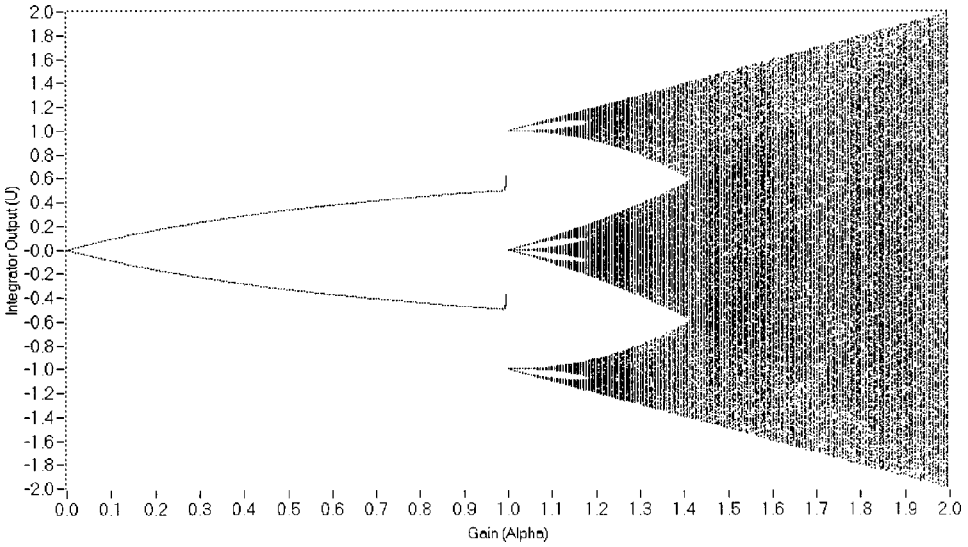

FIG. 4. Bifurcation diagram of a first order, one bit sigma delta modulator with 0 input and gain applied to the error (System 2).

sigma–delta modulator, the quantized output can assume one of $m=2^n$ states. The quantizer would take the form

$$Q(u) = \begin{cases} 2(m-1)/m & \text{if} \quad u \geqslant 2(m-2)/m \\ 2(m-3)/m & \text{if} \quad 2(m-2)/m > u \geqslant 2(m-4)/m \\ 2(m-5)/m & \text{if} \quad 2(m-4)/m > u \geqslant 2(m-6)/m \\ \vdots & \vdots \\ -2(m-1)/m & \text{if} \quad -2(m-2)/m > u \end{cases} \tag{5}$$

Here, for reasons explained in the section on bifurcation, we assume that quantizer input is in the range -2 to 2. Thus, the systems that will be studied are

(1) the first-order, single bit sigma–delta modulator with gain applied to the integrator: (1) and (2);
(2) the first-order, single bit sigma–delta modulator with gain applied to the error: (3) and (2);
(3) the first-order, multibit sigma–delta modulator with gain applied to the integrator: (1) and (5);
(4) the first-order, multibit sigma–delta modulator with gain applied to the error: (3) and (5).

ANALYSIS

Bifurcations

System 1 [Eqs. (1) and (2)], a first order, single bit sigma–delta modulator, is perhaps the most well-known and simplest form of sigma–delta modulation. It exhibits chaos if the gain is in the range $1<\alpha\leqslant 2$. The bifurcation diagram of this system is depicted in Fig. 3.

System 2 [Eqs. (3) and (2)], has a slightly different bifurcation diagram. It also exhibits chaos if the gain is in the range $1<\alpha\leqslant 2$. The bifurcation diagram of this system is depicted in Fig. 4. Notably, the dynamics here are somewhat different. For instance, the integrator output does not immediately reach the extremes as α is increased past 1. The full range of integrator output is between -2 and 2, and for $\alpha\geqslant 1$, the range of output extends from $-\alpha$ to α. This is a direct consequence of the fact that the bifurcation diagram measures possible values of $U_n=\alpha V_n$ from Eq. (4). It may seem problematic at first, since the expected input, X, is between -1 and 1. However, as shall be seen later, as long as the average integrator output sufficiently approximates the input, then this is not a difficulty. We simply require that the input signal be bounded by ± 1, even though the quantizer can accept input bounded by ± 2.

Stability

One difficulty with operating a sigma–delta modulator with greater than unity gain is that, for nonzero input, the modulator may become unstable. That is, $U_n \to \pm\infty$ as n

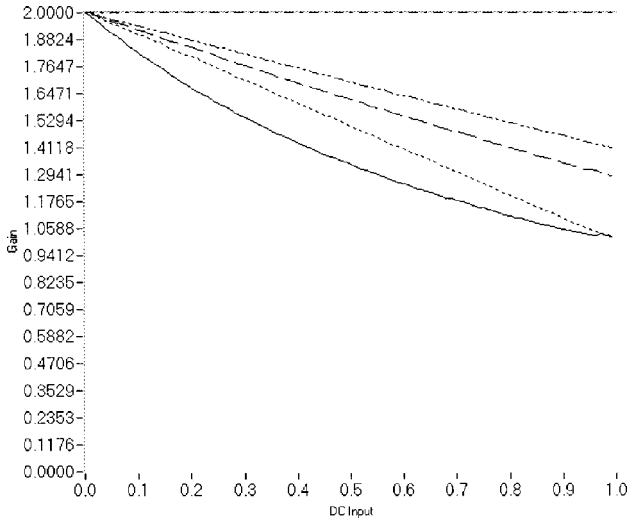

FIG. 5. The stability regime of a sigma delta modulator for various values of gain and constant input in the range 0 to 1. The solid line and below represents the bounded stable regime for a 1 bit modulator with gain applied to the integrator output (System 1). Similarly, the dashed line represents the bounded stable regime for a 2 bit modulator and the dot–dot–dashed line for a 3 bit modulator (System 3). For a modulator with gain applied to the error, the dotted line and below represents the stable regime for the 1 bit case (System 2), and the dot-dashed line and below represents the stable regime for the 2 bit case (System 4).

Downloaded 10 Jan 2002 to 192.58.150.40. Redistribution subject to AIP license or copyright, see http://ojps.aip.org/chaos/chocr.jsp

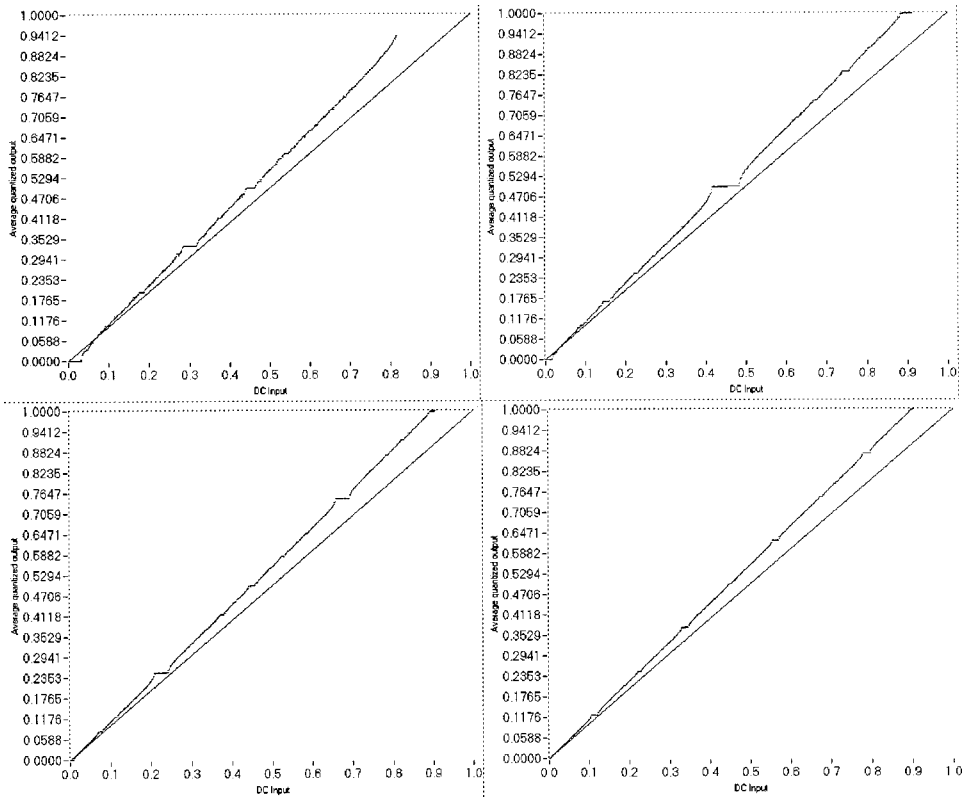

FIG. 6. Clockwise from top-left. The average quantized output as a function of the input for a 1 bit, 2 bit, 3 bit, and 4 bit sigma delta modulator with gain applied to integrator output. The gain is set to 1.1. The 45 degree line represents the ideal average quantization.

$\rightarrow\infty$. This is illustrated in Fig. 5, which depicts the size of the stable regime for input $0\leqslant X\leqslant 1$ (the plot is symmetric for $-1\leqslant X\leqslant 0$) and gain $0\leqslant\alpha\leqslant 2$. Operating a one bit sigma–delta modulator, Eq. (1), in the chaotic regime becomes unworkable for any large input, since the integrator output diverges. Although this can be improved through the use of a multibit quantizer, it is still problematic.

The stable regime is significantly increased if the gain is

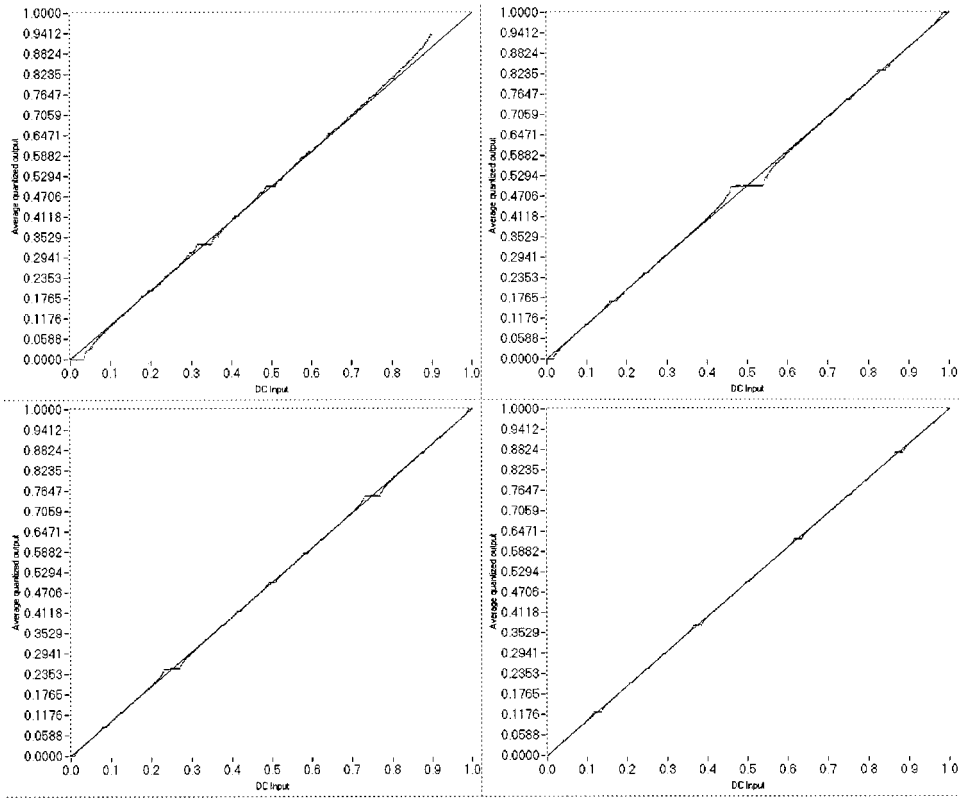

FIG. 7. Clockwise from top-left. The average quantized output as a function of the input for a 1 bit, 2 bit, 3 bit, and 4 bit sigma delta modulator with gain applied to quantization error. The gain is set to 1.1. The 45 degree line represents the ideal average quantization.

Downloaded 10 Jan 2002 to 192.58.150.40. Redistribution subject to AIP license or copyright, see http://ojps.aip.org/chaos/chocr.jsp

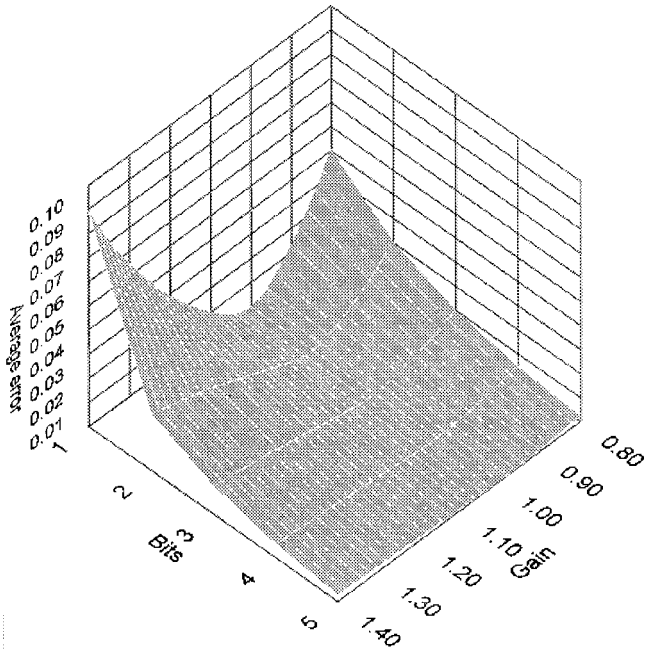

FIG. 8. A three-dimensional plot of average output error vs number of bits vs gain. Input is varied over the range -1 to 1. For each input number of bits, and gain, the absolute value of the input minus the average quantized output is found. This is then averaged over the entire input range, to give an average error dependent on the number of bits in the quantizer and the gain in the modulator [Eq. (3)]. 1 to 5 bit modulators having gain ranging from 0.8 to 1.4 are compared.

applied to the difference between the quantizer output and the integrator output [Eq. (3)]. For a 1 bit quantizer, the stable regime is greater than a 2 bit traditional sigma delta modulator. If we move to a 2 bit quantizer implementation of Eq. (3), then the entirety of the domain has bounded integrator output.

Quantization error

A minimum necessary requirement for sigma delta modulation is that the quantizer output approximate the input signal. That is

$$\lim_{N\to\infty}\frac{1}{N}\sum_{i=1}^{N} Q(U_i)=X, \tag{6}$$

for constant input X. This must hold for any allowable input. All the modulators considered have input in the range $-1 \leqslant X \leqslant 1$. With unity gain ($\alpha=1$), Eq. (6) holds for the single and multibit, first-order sigma–delta modulators. However, this is typically not true for $\alpha\neq 1$. Feely and Chua[11] showed that integrator leak, $\alpha<1$, may cause the average output of the sigma–delta modulator to assume discrete values that misrepresent the input. The resulting structure of average quantized output as a function of the input is known as a devil's staircase. As shown in Fig. 6, this is also the case for a traditional sigma–delta modulator with $\alpha>1$ [Eq. (1)]. In fact, for nonunity gain, the average output is approximately αX. Using a multibit modulator is not sufficient to alleviate this problem. Although it increases the bounded region of the modulator, and minimizes the discrete steps, it does not succeed in making the output more effectively track the input. This is a fundamental problem that is often overlooked in the literature.[12]

However, the modified modulator of (3) behaves quite differently. Figure 7 shows that this modulator, although assuming discrete values, still approximates the input. That is, the average output as a function of input has a slope of 1. In addition, a multibit implementation helps to minimize the length of the stairs in the devil's staircase structure. In other words, for a modulator operating in the chaotic region, as the number of bits used by the quantizer is increased, the average quantized output approaches the average quantized output of an ideal sigma–delta modulator.

To further demonstrate this, a three-dimensional plot of average output error vs number of bits vs gain is depicted in Fig. 8. Notice that quantizer error is greatly reduced as the number of bits in the quantizer is increased. Thus the loss in accuracy as the gain is increased (and chaos is introduced) may be compensated for by adding bits to the quantizer.

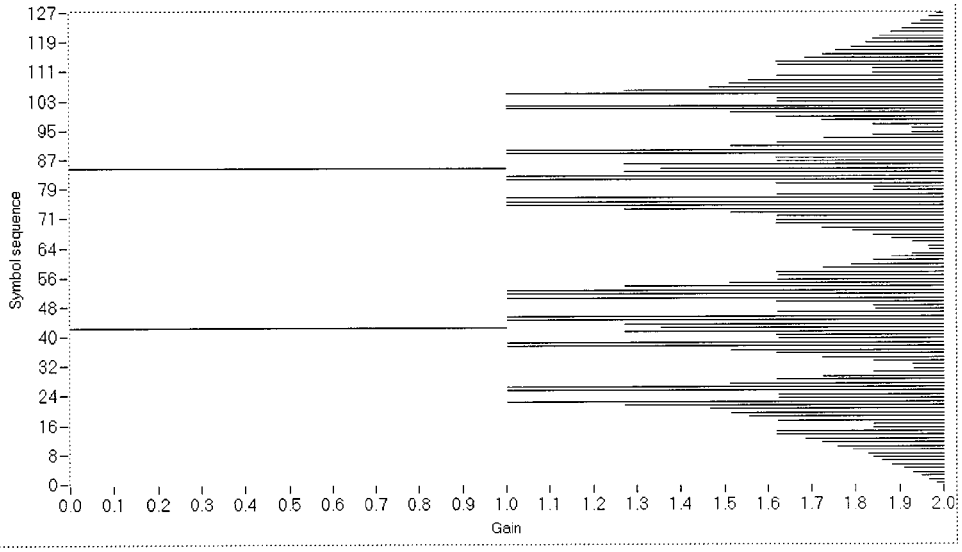

FIG. 9. The permissible 7 bit sequences that can be generated using a first order, single bit sigma delta modulator.

Downloaded 10 Jan 2002 to 192.58.150.40. Redistribution subject to AIP license or copyright, see http://ojps.aip.org/chaos/chocr.jsp

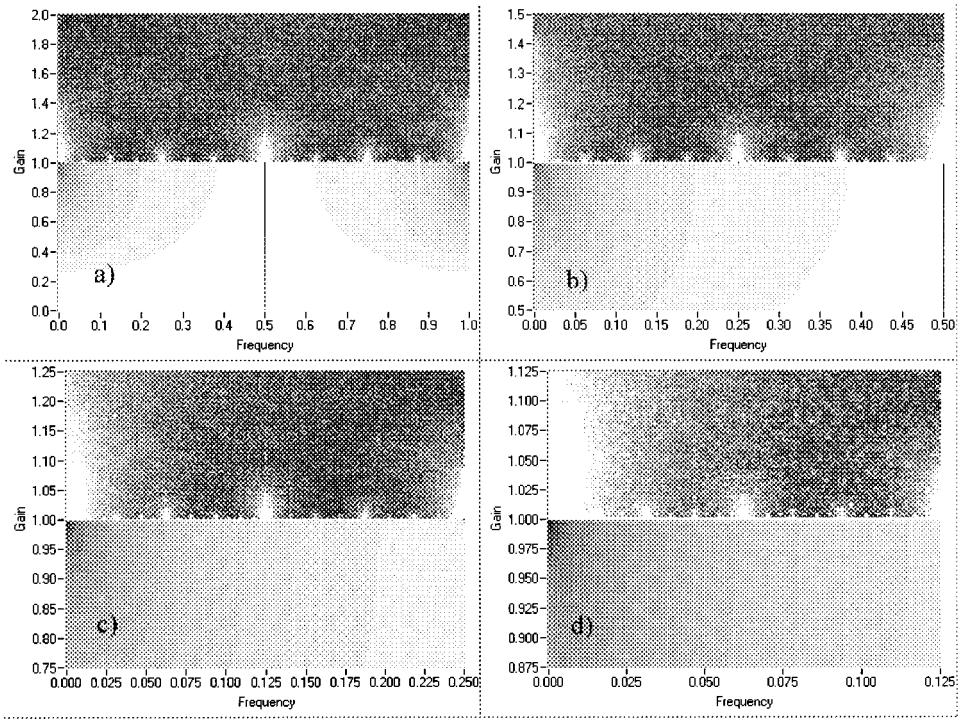

FIG. 10. Intensity plots of power spectra (a) is for a first order sigma delta modulator with constant zero input and gain ranging from 0 to 2. (b), (c), and (d) are successive magnifications (2×, 4×, and 8×) that indicate the fractal self-similar nature of the power spectrum.

Symbol sequences

On a practical level the output prior to quantization is not of primary concern. More importantly, the quantized output must accurately encode the input signal without producing idle tones. That is, tones which are not present in the input signal may appear in the quantized output. For instance, a constant input of 0 with $\alpha=1$ and initial condition $U_0=0$ will produce an output sequence of $1,-1,1,-1,1,$

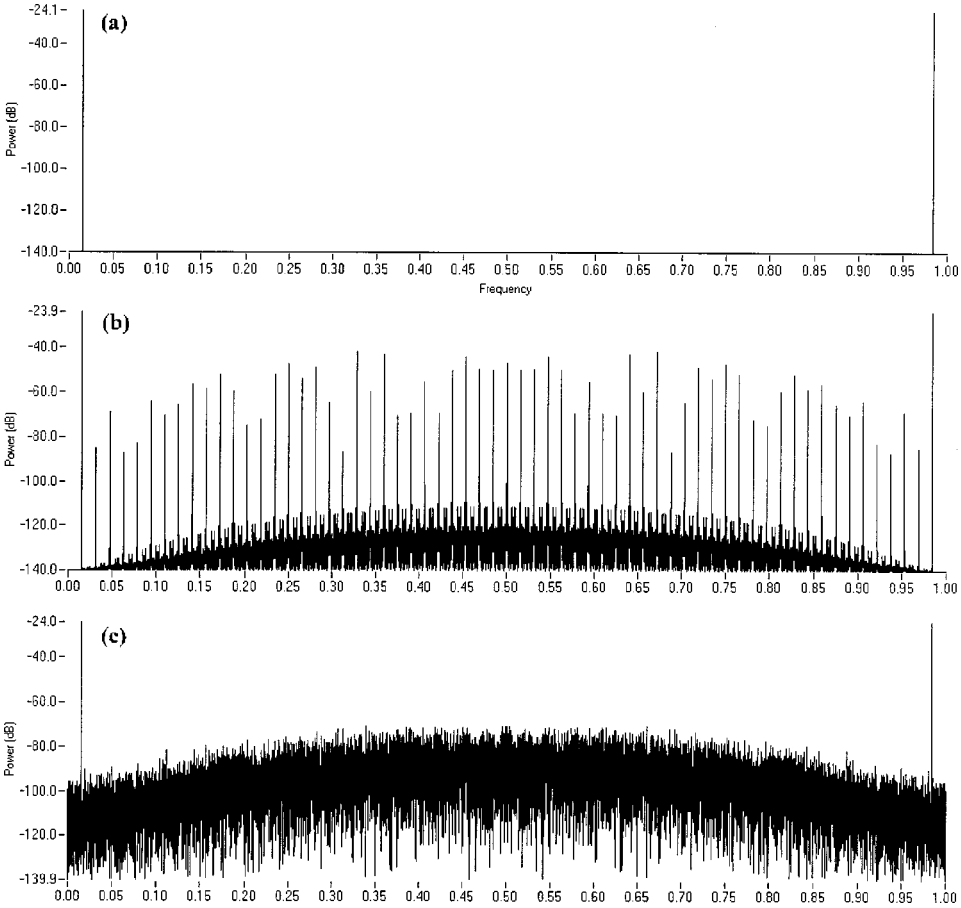

FIG. 11. Power spectra for the input signal $X_n=0.5\cdot\sin(2\pi\cdot n/64)$. (a) is the power spectrum for the input signal, (b) is the power spectrum for the quantized output signal with gain set to 1, and (c) is the power spectrum for the quantized output signal with gain set to 2. The power is assumed to have a base value of 10^{-7} (-140 dB).

Downloaded 10 Jan 2002 to 192.58.150.40. Redistribution subject to AIP license or copyright, see http://ojps.aip.org/chaos/chocr.jsp

-1 . . . . Longer period cycles will produce tones at lower frequencies which may appear audible to the listener.

One proposed method of eliminating these tones is to operate the sigma–delta modulator in the chaotic regime. Although the output will still approximate the input, limit cycles might be eliminated. As an example, a constant input of 0 with $\alpha=1.5$ will produce an output $1,-1,-1,1,-1,1,-1,-1,1,-1,1,-1,1,1$ . . . . This is an endless pattern that never settles into a limit cycle.

Thus the bifurcation diagrams that were produced earlier are not illuminating because they say nothing about the range of quantized dynamics. For these reasons it is important to investigate the symbol sequences that can be generated for various values of α. Figure 9 depicts the seven bit symbol sequences that can be generated for a single bit sigma–delta modulator with zero input (System 2). Seven bits were chosen simply for resolution—the qualitative structure of the resultant plot is the same for various choices of the number of bits. For gain ranging from 0 to 2 in increments of 0.001, 100 000 successive quantizer outputs were calculated. A sliding window of seven bits was applied to produce output sequences in the range 0 000 000–1 111 111 (0–127). A cyclic symbol sequence would be counted as multiple sequences, e.g., 0 101 010 . . . and 1 010 101 . . . are counted as separate allowable symbol sequences. This figure is exactly the same for both Systems 1 and 2.

Power spectra

One of the key reasons to attempt sigma delta modulation in the chaotic regime is to see if it can effectively eliminate idle tones, while at the same time preserving the frequencies in the input signal. For this reason, the power spectrum is an appropriate tool.

In Fig. 10, intensity plots are shown that reveal how the power spectrum is changed for gain from 0 to 2. Figure 10(a) depicts the power spectral intensity over the full range of gain and frequency for a 2 bit sigma–delta modulator (System 4) with zero input. For $\alpha\leqslant 1$, an idle tone exists with a frequency of 0.5. This is due to the quantizer flipping between 1/2 and $-1/2$. This tone is effectively removed in the chaotic regime. (b), (c), and (d) depict the power spectral intensity at $2\times$, $4\times$, and $8\times$ magnification, respectively. They depict the self-similar, fractal nature of the power spectrum for $\alpha>1$. This is another indication of chaos in sigma–delta modulation.

In Fig. 11, power spectra are depicted for an input signal with 32 times oversampling, $X_n=0.5\cdot\sin(2\pi-n/64)$, applied to (System 4). Figure 11(a) depicts the power spectrum for the input. As expected, peaks are seen at frequencies of 1/64 and 63/64. However, for a 2 bit sigma–delta modulator with unity gain, the output power spectrum exhibits additional peaks at all multiples of 1/64 [Fig. 11(b)]. In Fig. 11(c), the modulator is operated at maximum gain [Eq. (3)], $\alpha=2$. The idle tones are completely removed, and replaced by chaotic fluctuations similar to broadband noise. This noise can be filtered, thus leaving only the frequencies that were apparent in the original signal.

CONCLUSIONS

A conventional first-order sigma delta modulator, where gain is applied to the integrator output, does not approximate input for $\alpha\neq 1$. This is true even if multibit quantizers are used. The errors in quantization due to the Devil's staircase structure introduced by chaotic modulation can be compensated for by using a multibit quantizer. However, this does not correct the fact that the output is offset from the input for a traditional chaotic sigma–delta modulator. If instead the gain is applied to the error in quantization, then the sigma–delta modulator may achieve accurate quantization over a far greater range of input. If a multibit quantizer is also used, then the modulator can be made stable over the full range of input. This has the benefit that idle tones can be removed from the quantization process by operating in the chaotic regime.

ACKNOWLEDGMENTS

The authors gratefully acknowledge Orla Feely of University College Dublin and Anthony Davies of King's College London for their insightful comments and criticisms concerning this work.

[1] W. Chou and R. M. Gray, ''Dithering and its effects on sigma delta and multi-stage sigma delta modulation,'' IEEE Trans. Inf. Theory **37**, 500–513 (1991).
[2] A. J. Magrath, ''Algorithms and architectures for high resolution sigma–delta converters,'' Ph.D. thesis, King's College, University of London, 1996.
[3] A. J. Magrath and M. B. Sandler, ''Power digital-to-analogue conversion using a sigma-delta modulator with controlled limit cycles,'' Electron. Lett. **31**, 251–253 (1995).
[4] V. Friedman, ''The Structure of limit cycles in sigma delta modulation,'' IEEE Trans. Commun. **36**, 972–979 (1988).
[5] S. J. Park, R. M. Gray, and W. Chou, ''Analysis of a sigma delta modulator with a multi-level quantizer and a single-bit feedback,'' in Proceedings of the ICASSP 91 (Toronto, Canada, 1991), pp. 1957–1960.
[6] D. T. Hughes, S. McLaughlin, G. Ushaw *et al.*, ''Search for chaotic output from a sigma–delta modulator: An archetypal retarded nonlinear system,'' in *Chaos in Communications*, edited by Louis M. Pecora (SPIE-The International Society for Optical Engineering, Bellingham, WA, 1993), pp. 194–204.
[7] P. A. Littelehales, C. P. Lewis, and S. R. Bishop, ''Safe basins and chaos in sigma delta modulators,'' in Proceedings of the Ninth International Conference on Systems Engineering, Las Vegas, July 1993 (Conference on Systems Engineering, Las Vegas, 1993), pp. 597–601.
[8] O. Feely, ''A tutorial introduction to non-linear dynamics and chaos and their application to sigma-delta modulators,'' Int. J. Circuit Theory and Applications **25**, 347–367 (1997).
[9] O. Feely, ''Nonlinear dynamics of discrete-time electronic systems,'' IEEE Circuits and Systems Society Newsletter **11**, 1–12 (2000).
[10] R. M. Gray, ''Oversampled sigma–delta modulation,'' IEEE Trans. Commun. **COM-35**, 481–487 (1987).
[11] O. Feely and L. O. Chua, ''The effect of integrator leak in S-D modulation,'' IEEE Trans. Circuits Syst. **CAS-38**, 1293–1305 (1991).
[12] O. Feely and L. O. Chua, ''Multilevel and non-ideal quantization in sigma-delta modulation,'' Int. J. of Circ. Theory and Appl. **21**, 61–83 (1993).

Downloaded 10 Jan 2002 to 192.58.150.40. Redistribution subject to AIP license or copyright, see http://ojps.aip.org/chaos/chocr.jsp