Theory and Practical Implementation of a Fifth-Order Sigma–Delta A/D Converter*

R. W. ADAMS, *Member*, P. F. FERGUSON, JR., A. GANESAN, S. VINCELETTE, A. VOLPE, AND R. LIBERT

Analog Devices Semiconductor, Wilmington, MA 01887, USA

One-bit noise-shaping converters with high-order loop filters have long been considered difficult or impossible to stabilize. A new approach is discussed which simplifies the analysis of high-order loops, and several topologies that can be used for the implementation of such loops are presented. An 18-bit analog-to-digital converter integrated circuit has been designed using this approach, which uses a fifth-order loop filter with an unusual loop topology. Implementation of the loop using differential switched-capacitor techniques is described. A one-step decimation filter with 115-dB stopband attenuation is used to remove out-of-band quantization noise. Preliminary measurements indicate a dynamic range of 105 dB, a number that is compatible with the requirements of professional digital audio.

0 INTRODUCTION

The advantages of sigma–delta analog-to-digital (A/D) and digital-to-analog (D/A) converters over conventional successive-approximation converters are by now well known in the audio community. The lack of differential nonlinearity and the resulting low-level distortion products as well as the built-in linear-phase filtering remove most of the objections that audio designers raised in the early years of digital audio.

There are currently two distinct design approaches that can be used to design a 1-bit single-loop sigma–delta converter—single-loop low-order designs and single-loop high-order designs.

0.1 Single-Loop Low-Order (<2) Designs

These designs are generally used for voice-quality communications and in some audio applications if the oversampling ratio is high enough (>128:1). First-order designs tend to have very poor performance for audio applications due to the strong tendency of these loops to produce correlated idling patterns, which result in audible tones in the noise floor. Second-order designs are vastly superior to first-order designs both in terms of the required oversampling ratio to achieve a particular signal-to-noise ratio as well as in the improved randomness of the idling patterns. Most of these designs are based on the pioneering work of Candy [1] at Bell Labs. Recently a number of researchers have shown that the second-order loop is not free of correlated idling patterns [2], a fact that is confirmed by the practical experience of the author, who has spent more time than he cares to think about listening to the quantization noise of various loops in the presence of small dc inputs.

0.2 Single-Loop High-Order (>2) Designs with Single-Bit Loop Quantizer

These designs offer improved signal-to-noise ratio for a given oversampling ratio as well as improved freedom from idling patterns and tones in the noise floor. The traditional objection to these loops is that they are impossible to stabilize. This objection was understandable a decade ago, but the fact that it still exists today is remarkable in view of the fact that commercial products using high-order loops have been available now for several years [3]. Engineers at Philips Laboratories were designing third-order loops in the late 1970s (although none of this work was published until later), and nonlinear stabilization of high-order loops was patented by Gould (now Martin Marietta) in the early 1980s [4], [5]. A large part of this paper is devoted to describing a methodology for designing stable single-bit high-order loops, as these loops show

* Presented at the 90th Convention of the Audio Engineering Society, Paris, France, 1991 February 19–22.

the most promise for achieving high performance using simple VLSI processes.

1 MODELING THE LOOP QUANTIZER

The traditional approach to designing a sigma–delta converter is to replace the comparator by an additive white-noise source with unity gain, as shown in Fig. 1. The transfer function of the noise input to the output can then be used as a figure of merit to determine the optimum noise-shaping filter. The transfer functions are simple to derive:

$$\mathrm{NS}(z) = \frac{Y(z)}{Q(z)} = \frac{1}{1 + H(z)} \tag{1}$$

$$\frac{Y(z)}{X(z)} = \frac{H(z)}{1 + H(z)} \tag{2}$$

where

$\mathrm{NS}(z)$ = noise-shaping transfer function
$Y(z)$ = z transform of output signal
$X(z)$ = z transform of input signal
$Q(z)$ = z transform of error signal added by quantizer.

This crude model is widely used in the design of low-order sigma–delta converters, but in the case of high-order loops it can cause erroneous results. Some of the deficiencies of this simple model when applied to high-order loops are listed below.

1) Assuming that the comparator has a particular "gain" is very dangerous. As an example, see Fig. 1(b), where we have put a gain of 10 in front of the comparator of the real circuit as well as in front of the linear model of the circuit. Clearly, the operation of the real circuit with the comparator will be unchanged, as the sign of a number is not changed by multiplication.

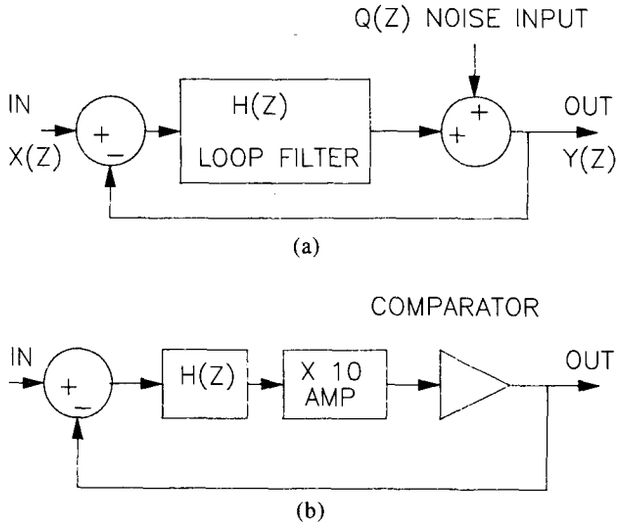

Fig. 1. (a) Linear model of loop showing comparator replaced by additive noise source. (b) Real sigma–delta with extra gain before comparator. Operation is unaffected, showing one of the drawbacks of the linear model.

This gain block could in fact be placed before the loop filter $H(z)$ with the same results. On the other hand, a gain of 10 inserted in the linear model of the loop will alter dramatically the transfer function to the output from both the noise and the signal inputs.

2) The comparator has the property that its output power is constant (assuming ± 1 output levels). Parseval's power theorem relates the power in the time-domain representation of a signal to the area under the power spectral density of a signal. Therefore we conclude that the area under the square of the noise-shaping spectrum must be constant. The linear model does not predict this.

Fortunately there is still a way that we can use the linear model and get results that are reasonably accurate. The input to a comparator operating in a real sigma–delta loop consists of a random idling component superimposed on a slowly varying (relative to the sampling rate) signal component due to the input signal. If we place an average-responding voltmeter on both the input and the output of the comparator [Fig. 2(a)], we can sweep the input over the specified input range and observe the voltages V_1 and V_2. Fig. 2(b) shows a typical plot of V_2 versus V_1, and we can define the gain of the comparator at a particular input to be the slope of this curve. The measured gain can be explained by convolving the probability density function (pdf) of the idling noise with the ideal voltage transfer function of a comparator, as shown in Fig. 2(c). This approach is

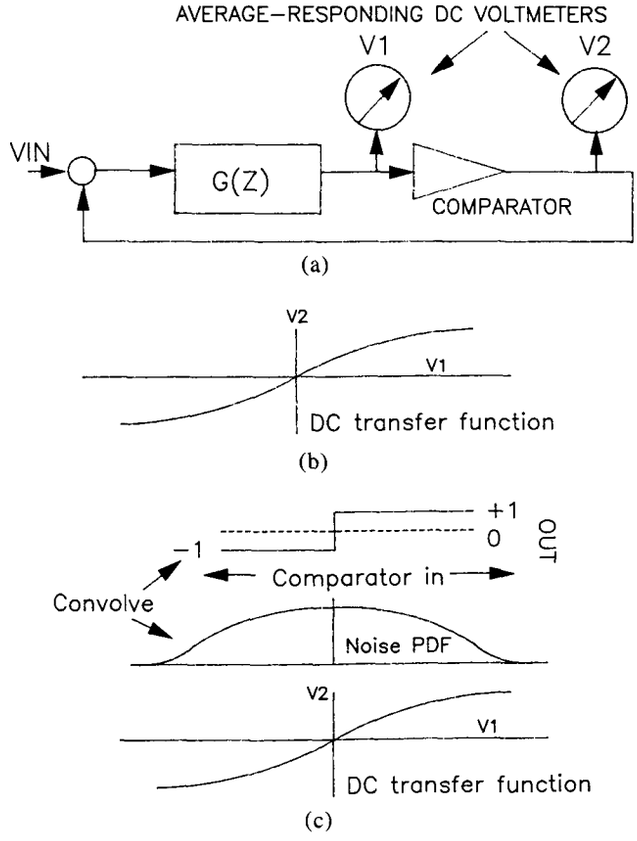

Fig. 2. (a) In-circuit measurement of comparator gain. (b) Typical measured gain curve V_2/V_1. (c) Average gain derived by convolving input pdf with ideal transfer function.

not as accurate as the direct measurement approach of Fig. 2(a), as the pdf of the idling noise is assumed to be invariant with the input level, which is not correct. The pdf of the idling noise becomes asymmetrical for larger inputs.

This procedure does not tell us how to design a sigma–delta converter, as the gain of the comparator must be measured in a working circuit. But it does provide us with an analysis tool that can be used in an iterative design process. For example, if we design a loop using the linear model with an assumed gain of 1, we can then simulate this loop and measure the average dc input and output of the comparator for a range of inputs as described. If this gain does not agree with the gain assumed in the linear model, we can change the loop filter in a direction to increase (or decrease) the actual gain to make it agree with the assumed linear-model gain. Since the gain varies inversely with the magnitude of the idling pdf at the comparator input, a loop filter that increases the magnitude of the noise-shaping transfer function at high frequencies will tend to increase the width of the pdf at the comparator input and hence decrease the comparator gain. Usually only a few such iterations are necessary before agreement is reached between the assumed gain used in the linear model and the actual gain as measured in the circuit.

One potential problem with this approach is that if the starting point of the iterative process is grossly wrong, then the loop will oscillate and no in-circuit gain measurements can be made. There are two solutions to this problem. One involves using rules of thumb that are obtained by the painful experience of other designers. The rule of thumb for single-bit systems is that the noise-shaping response of the linear model should on average be about 3 dB over most of the out-of-band portion of the frequency spectrum, assuming that the response is reasonably flat and not peaky at these frequencies. An example of a typical noise-shaping response is shown in Fig. 3.

The other approach is to simulate the circuit initially with an infinite quantizer in the loop and start with a linear noise-shaping response that has enough gain at high frequencies that more than one quantization level will be active during idling. The noise-shaping response can then be adjusted until only one level is used for the idling patterns (with occasional excursions into the adjacent levels). At this point the rest of the quantizer levels can be discarded, leaving only a single comparator. This approach is accurate enough to provide a starting point for the remainder of the iterative process outlined.

The fact that the gain of the comparator falls with the input level [Fig. 2(b)] explains why higher order single-bit modulators become unstable for inputs that approach overload. Such systems are called conditionally stable systems, as they are only stable for some range of gains. There are many circuit solutions to this problem, some of which are discussed later. Most higher order single-bit systems limit the input range to somewhere between 50 and 70% of the clipping limit (the point at which the input exceeds the comparator feedback levels).

One criticism of this method for deriving the comparator gain is that it is valid only for frequencies much lower than the sampling rate. For frequencies close to the sampling rate, the implicit time averaging does not apply. Fortunately, the frequencies where most of the open-loop zeroes of $H(z)$ occur are well below the sampling rate (typically by a factor of 30:1), so the view of the quantizer as a linear element is still reasonably accurate.

2 REVIEW OF CONDITIONAL STABILITY

It is a common misconception that any negative feedback loop with more than unity gain at the point where the phase shift of the loop filter is 180° is unstable. Since we are assuming negative feedback, this point corresponds to 0° of phase shift around the complete loop.

Whether or not such a loop is unstable depends on the Nyquist plot of the open-loop system. The Nyquist plot of a conditionally stable system is shown in Fig. 4(a). The criterion for stability is whether or not there is a net encirclement of the (−1, j0) point, which in this case there is not. Note that the points labeled A and B represent points where the loop phase shift is 0° and the loop gain is greater than unity. If we now reduce

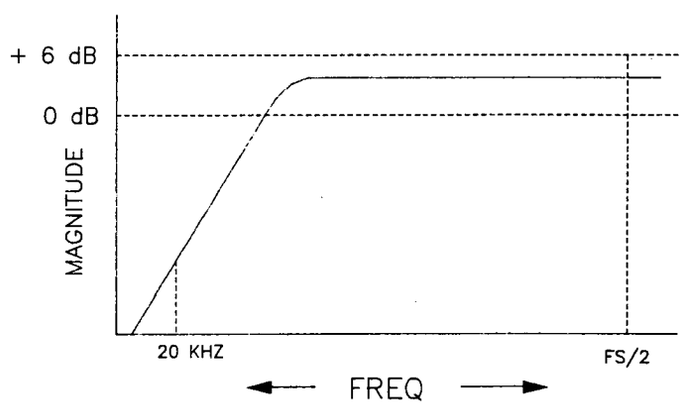

Fig. 3. Typical noise-shaping response for stable operation with 1-bit quantizer.

the open-loop gain of the system, we can see that there will be a net encirclement of the $(-1, j0)$ point, causing instability [Fig. 4(b)]. It is this mechanism that is responsible for oscillations in higher order sigma–delta converters. If the gain of the comparator as defined in Sec. 0 is reduced to some critical point by an input signal that is too close to the input clipping level, then the average gain of the comparator is reduced, causing oscillations.

Another view of the same mechanism is shown in Fig. 5. Here we have the traditional Bode plot of the open-loop response. If there are N low-frequency poles and $N - 1$ high-frequency zeros that occur before the unity-gain crossover point, we can say that the system is stable if the slope of the response is <12 dB per octave at the point where the open-loop gain is unity. Fig. 5 shows what happens as the gain is reduced, and it is clear that the unity-gain crossover now occurs where the slope is greater than 12 dB per octave. This is not as rigorous an approach as the Nyquist analysis, but it does illustrate the problem in an intuitively appealing manner.

In general, conditionally stable systems are only stable over a range of loop gains, and real-world conditions such as comparator overload or turn-on transients can cause oscillations to occur that do not go away even after the input that caused them has disappeared.

3 LOOP FILTER DESIGN FOR HIGH-ORDER SINGLE-BIT LOOPS

In this section we introduce a number of different noise-shaping transfer functions and examine their performance when used in the design of single-bit loops. We make heavy use of the "comparator gain" analysis presented in Sec. 2. In many cases we start by assuming an infinite quantizer and show how the noise-shaping response must be modified for use in a single-bit system. Our approach is to start with the noise-shaping function NS(z), as it is this function that will determine our signal-to-noise ratio and hence the number of bits in the final output. The loop filter $H(z)$ can always be derived from NS(z) by simple loop algebra. All transfer functions will be z-domain functions, implying that a discrete-time loop filter will be used (most commonly employing switched-capacitor circuitry). The results we obtain are valid both for A/D loops where the loop filter is an analog circuit as well as for D/A loops where the loop filter is digital. Brick-wall decimation filters are assumed throughout.

There is a fundamental restriction that arises from causality constraints, which limits our choice of noise-shaping transfer functions. This restriction states that the loop must not contain any delay-free loops. This implies that the loop filter $H(z)$ must have a z^{-1} delay, and therefore the order of the numerator of H(z) must be one less than the order of the denominator of H(z),

$$H(z) = \frac{a_1 \cdot z^{-1} + a_2 \cdot z^{-2} + \cdots}{b_0 + b_1 \cdot z^{-1} + b_2 \cdot z^{-2} + \cdots} . \quad (3)$$

The most important conclusion we can draw from this condition is that *the first value of the impulse response of NS(z) must be 1.*

This conclusion is intuitively obvious if we consider what happens if we apply an impulse to the noise input of the linear model shown in Fig. 1. Since the loop filter $H(z)$ has a delay of at least 1, the loop cannot respond to this input until the next cycle, and therefore the impulse instantly appears at the output with unity weighting.

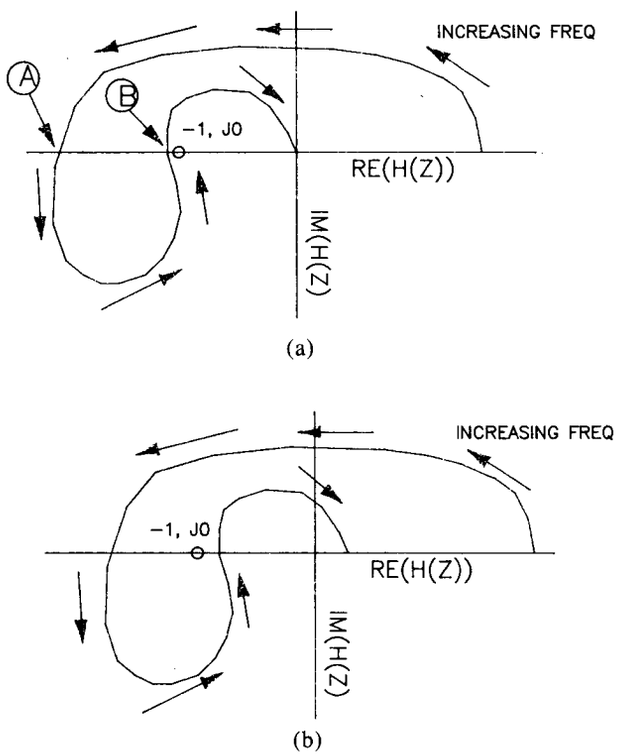

Fig. 4. (a) Conditionally stable Nyquist plot. (b) Unstable Nyquist plot.

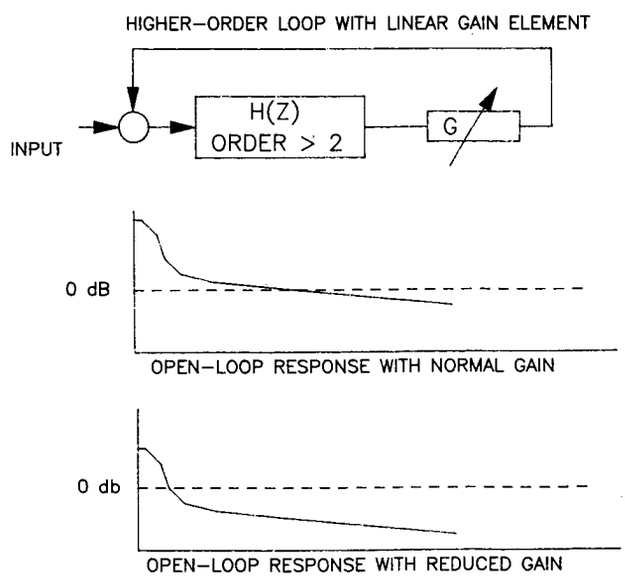

Fig. 5. Bode plot of instability mechanism in conditionally stable system.

This fundamental rule helps narrow down the choice of NS(z) to those functions whose first impulse value of the time response is 1. We usually choose NS(z) to be some cascade of first- and second-order IIR sections,

$$\mathrm{NS}(z) = \frac{g_1 \cdot (1 - a_1 \cdot z^{-1}) \cdot (1 - b_1 \cdot z^{-1} + b_2 \cdot z^{-2}) \cdots}{g_2 \cdot (1 - c_1 \cdot z^{-1}) \cdot (1 - d_1 \cdot z^{-1} + d_2 \cdot z^{-2}) \cdots} . \quad (4)$$

The restriction that the first value of the impulse response equal 1 can then be translated into the requirement that $g_1/g_2 = 1$, as g_1/g_2 is the z^0 term of the ratio of polynomials, Eq. (4).

We now examine three different noise-shaping polynomials for their suitability as noise-shaping functions in a single-bit system.

3.1 Nth-Order Pure Differentiation

Let us start with the simplest noise-shaping function, pure Nth-order differentiation of the form

$$\mathrm{NS}(z) = (1 - z^{-1})^N . \quad (5)$$

This function meets our requirements that the first value of the impulse response be equal to 1. The magnitude response of NS(z) for various values of N is shown in Fig. 6. Note that at high values of N we obtain more effective suppression of quantization noise at low frequencies, and more gain at high frequencies.

Fig. 7 shows a histogram of quantization levels used for idling only for various loop orders, assuming an infinite quantizer [6]. These plots are obtained by running a loop for thousands of cycles and counting the number of occurrences of each quantization level in the loop quantizer. Note that the number of levels used increases at a rate proportional to 2^N, which makes sense since the high-frequency noise-shaping gain also increases at this rate. From our previous discussion of comparator gain versus assumed linear-model gain, it appears that this noise-shaping function is not a good candidate for loops with 1-bit quantizers. As a rule of thumb, any high-order loop that uses many levels for idling with an infinite quantizer will oscillate when the extra levels it needs are taken away and a 1-bit quantizer is used. The reason for this can be explained using the pdf-derived comparator gain argument outlined earlier.

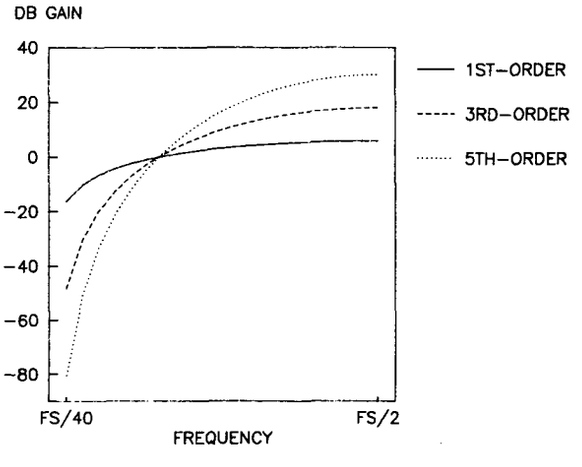

Fig. 6. Pure differentiated noise-shaping response. NS(z) = $(1 - z^{-1})^N$, N = 1, 3, 5.

The fact that a loop uses many levels for idling implies that the random idling-waveform input to the quantizer is large. If we use only a 1-bit quantizer, then the width of the pdf at the comparator input will be large, causing its low-frequency gain as described in Sec. 1 to be very small. This small gain causes oscillations to occur, as described in Sec. 2.

3.2 Butterworth High-Pass Response

The problem encountered in the previous noise-shaping function was the large high-frequency noise-shaping gain for large N, which causes the idling waveform at the comparator input to be very large, resulting in low comparator gain and hence instability. We can modify the pure differentiating response by introducing poles into NS(z),

$$\mathrm{NS}(z) = \frac{(1 - z^{-1})^N}{D(z)} \quad (6)$$

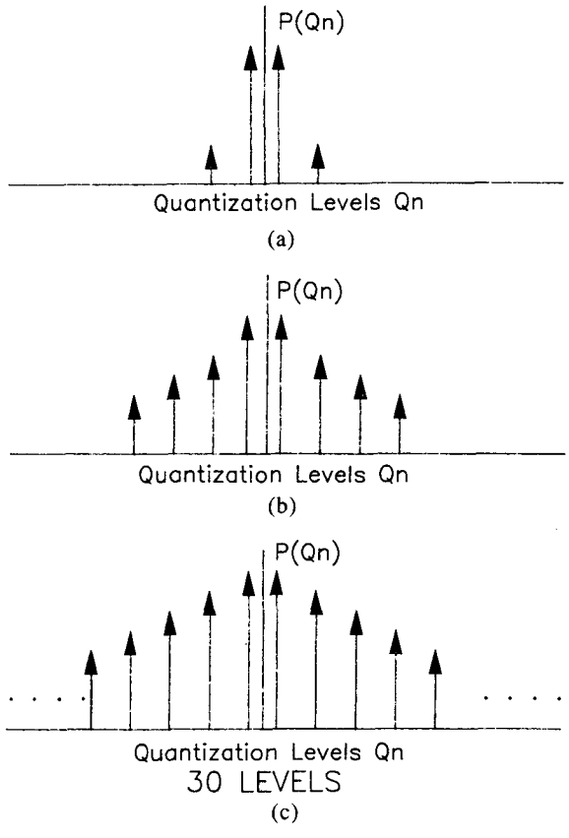

Fig. 7. Number of quantization levels used for idling for noise shaper using Nth-order differentiation as noise-shaping transfer function. (a) First order. (b) Third order. (c) Fifth order.

where $D(z)$ is an Nth-order z polynomial.

The purpose of adding $D(z)$ is to flatten the high-frequency portion of NS(z), as shown in Fig. 8(a). NS(z) is now simply an Nth-order high-pass function. With the correct choice of $D(z)$, a Butterworth alignment of the poles can be obtained, resulting in a maximally flat high-frequency region of NS(z).

The high-frequency gain of NS(z) is now a design parameter that is independent of order N and in fact can be adjusted by varying the high-pass corner frequency of NS(z). We must be careful, however, to ensure that the first coefficient of NS(z) is 1, as no delay-free loops are permitted. In practice, what we do is to make an initial guess for the high-pass corner frequency. Using a commercial digital IIR filter design package, we plug in the high-pass corner frequency and specify a Butterworth high-pass response. We take the output from this program, find the value of the highest power of z (the first value of the impulse response), and multiply by whatever constant is required to make this value equal 1. Using this technique for different corner frequencies, we can derive a family of high-pass functions that all meet the requirement that the first coefficient equals 1, as shown in Fig. 8(b).

While it is possible to implement all of these curves using the linear model, only one of them will survive the transformation from linear model to sigma–delta loop. This is because the differing amounts of high-frequency noise-shaping gain produce differing amounts of idle noise at the comparator input and hence different average comparator gains, where the comparator gain is defined by the pdf convolution argument given in Sec. 1. The design procedure then becomes an iterative approach, as follows.

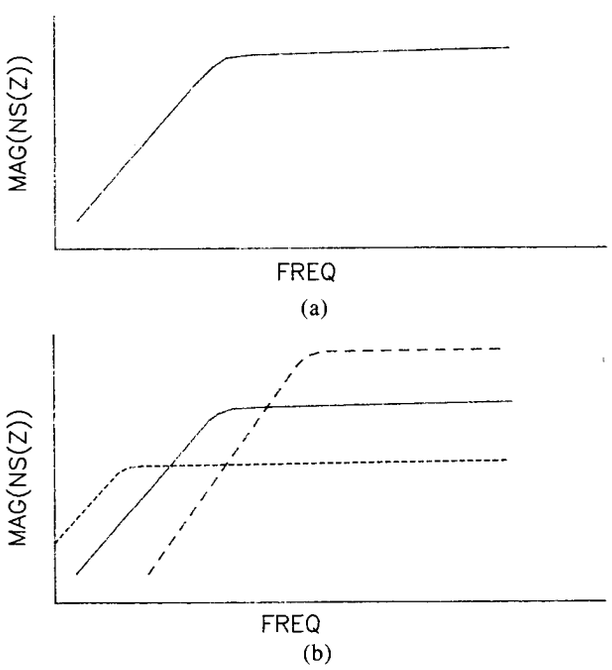

Fig. 8. (a) Noise-shaping function using Butterworth alignment of poles. (b) Family of Butterworth curves with different corner frequencies and scaled for first value of impulse response = 1.

1) Design an Nth-order Butterworth high-pass filter using any commercially available program. Start with a cutoff frequency between 0.1 and 0.02 of the sampling frequency.

2) Examine the first value of the impulse response (coefficient of the highest power of z). Call this value C_1. The high-frequency gain of NS(z) after normalizing the first coefficient to 1(by multiplying the function by $1/C_1$) will be $1/C_1$. We have empirically determined that a good value for the high-frequency gain of NS(z) is about +3 dB, so a good value for C_1 is 0.707. This value will usually give a comparator gain that is reasonably close to the assumed linear gain. We are assuming comparator output levels of +1 in this example; other values would lead to different empirically optimum values of C_1.

3) If the value of C_1 is not within 10% of 0.707, then go to step 1 and change the high-pass corner frequency. Increasing the corner frequency will decrease C_1, which increases the high-frequency gain of NS(z) after normalization, resulting in increased idling activity at the comparator input and hence lower comparator gain in the actual sigma–delta circuit. Repeat this loop until the correct corner frequency is found.

4) Find the loop filter $H(z)$ from NS(z) by the equation

$$H(z) = \frac{1 - \mathrm{NS}(z)}{\mathrm{NS}(z)} . \tag{7}$$

5) Plug $H(z)$ into a discrete-time simulator, or write your own difference equations for the loop directly from $H(z)$. This is quite easy to do.

6) Add low-pass filters to the comparator inputs and outputs so that the dc value of the input and output signals can be observed.

7) Sweep the inputs over a range of dc values and plot the comparator dc transfer function. Find the gain around zero of the comparator and compare to the gain assumed in the linear model (usually unity gain was assumed for the linear model).

8) If the gain is too low, go back to step 1 and decrease the high-pass corner frequency. If the gain is too high, increase the corner frequency. Repeat until the gain is close to unity.

Once this iterative process is complete, the correct NS(z) and $H(z)$ have been found and the loop will be stable over some range of dc inputs. Further simulation can be performed to ascertain the stable input range and the low-frequency signal-to-noise ratio.

3.3 Inverse Chebyshev with Complex Zeros on the Unit Circle

The Butterworth high-pass response of Sec. 3.2 can be modified to move the real stopband zeros at the (1 + j0) point out on the unit circle to produce nulls in the noise-shaping transfer function at frequencies other than dc. An example of such a response for a third-order system is shown in Fig. 9, where there is real zero at dc and one complex pair at the edge of the audio band. Compared with the Butterworth alignment,

improved signal-to-noise ratio can be obtained as the area under the audio portion of the noise-shaping response is reduced.

The equation for a general Nth-order noise shaper of this type is given by

$$\mathrm{NS}(z) = \frac{\prod \text{first-order terms} \cdot \prod \text{second-order terms}}{D(z)} . \tag{8}$$

For a given order N, the designer must decide how many complex zero pairs will be moved to nonzero frequencies on the unit circle. For example, a fifth-order system could have three real zeros and one complex pair, or it could have one real zero and two complex pairs. Generally, using the largest number of complex pairs possible for a given order results in the best noise-shaping characteristic, although there may be circuit-related reasons for choosing otherwise.

The same iterative design methodology outlined for the Butterworth case can be used here as well, the only difference being that the commercial IIR design program should be instructed to design an inverse Chebyshev filter instead of a Butterworth filter, and that the stopband edge should be entered as the audio-band edge. Note that the design program will result in a filter that is equiripple in the stopband. This is not necessarily the optimum solution in terms of the total integrated noise over the audio band, and some adjustment of the complex zero locations may yield several decibels of improvement in signal-to-noise ratio. It is also advisable to do this optimization in conjunction with the known frequency response of the decimation filter, as the noise from the noise shaper rises at a very fast rate at frequencies just beyond the last complex pair frequency, which is typically quite close to the decimator cutoff frequency.

4 COMPUTATION AND IMPLEMENTATION OF G(z)

The NS(z) design approach given in Sec. 3 must now be converted to a loop filter $H(z)$ using the simple equation

$$H(z) = \frac{1 - \mathrm{NS}(z)}{\mathrm{NS}(z)} . \tag{9}$$

Note that by insisting that the first impulse value of NS(z) is a 1, we have guaranteed that $H(z)$ has a unit delay, as the term $1 - \mathrm{NS}(z)$ will cancel the z^0 term in the numerator.

There are numerous topologies that can be used to implement $H(z)$, and in fact $H(z)$ can be split into numerator and denominator sections, and these can be implemented separately in different locations within the loop. In [7] a state-variable topology is described that can synthesize arbitrary noise-shaping functions, including both the Butterworth and the Chebyshev designs discussed here.

Fig. 10 shows two topologies that can be used to implement the Butterworth design (Sec. 3.2). We show a fifth-order topology for convenience, but any order can be implemented using the same topology. These topologies are equivalent in that they produce the same NS(z) for a given set of coefficients, but they are not equivalent in terms of the transfer function $Y(z)/X(z)$.

The topology of Fig. 10(a) is guaranteed to produce a flat frequency response over the audio band, as $H(z)$ must have high gain in the audio band in order to provide effective noise shaping. This is not the case for the topology of Fig. 10(b), where part of $H(z)$ is distributed into the input nodes of the integrators. For the Butterworth case described earlier, it turns out that the Y/X response becomes a Butterworth low-pass response, and therefore the audio-band response can be made quite flat with little difficulty.

To implement the inverse Chebyshev polynomial NS(z), we can modify these topologies as shown in Fig. 11. Here we show feedback paths around groups of two integrators, which can be shown to shift their poles away from dc along the unit circle in conjugate pairs. Again the flow-graph inversion rule may be applied to obtain two topologies that are identical from a noise-shaping point of view.

To determine the correct coefficients for any of these four topologies, we must derive a z-transform expression for the open-loop transfer function of each topology and equate it with the desired $H(z)$ [which was computed from NS(z) using the iterative methodology given in Sec. 3). The algebra can become a bit messy, and the authors have had good success using commercially available nonlinear equation solvers to help out, that is, we can easily write a nonsimplified algebraic expression from the topology itself and equate this to the desired $H(z)$ with known coefficients. The unknown coefficients are then easily found using the equation solver.

After the coefficients are determined, a difference equation simulation is usually performed, either by writing difference equations directly from the topology

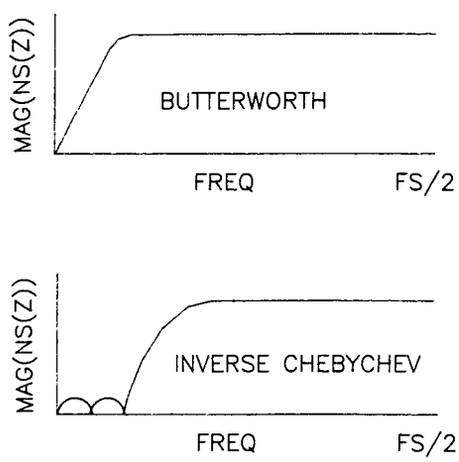

Fig. 9. Modification of Butterworth curve to give inverse Chebyshev noise-shaping response.

in some high-level language, or by using a general-purpose difference-equation simulation tool. After the circuit has been simulated to verify that it meets signal-to-noise and stability requirements, it is usually necessary to scale the integrator gains and coefficients so that the peak levels at the integrator outputs are within certain bounds dictated by the voltage limits of the actual circuit. It is then a simple step to convert the topology into a switched-capacitor circuit with the appropriate capacitor values.

5 NONLINEAR STABILIZATION

High-order loops with single-bit quantizers always have the potential for oscillation when the input exceeds the stable input range or when the circuit is first turned on. The reason for this instability is the dependence of comparator gain on input level as described.

There are two possible fixes for this problem. The most obvious fix is to sense instability either by looking for large integrator voltages or by sensing large strings of consecutive 1's or 0's in the bit stream, and then resetting the circuit. By clever design of both the modulator and the decimator, it is possible for the decimated output signal to be well-behaved while multiple reset events are occurring.

The other method of stabilization relies on nonlinear stabilization techniques [4], [5]. Nonlinear stabilization is implemented by placing clippers on the output of each integrator signal (or any other set of valid state variables in the system). The thresholds of these clippers are selected such that they are slightly higher than the worst-case voltages encountered during normal (stable) operation. This effectively puts a bound on the available state space of the modulator, and limits it to the stable region of operation. Another view of this circuit is that it dynamically alters the gain/phase characteristics of the loop when the loop begins to oscillate such that stability is maintained. For example, it is obvious that an integrator that is heavily clipped at its output for a sine-wave input will exhibit less phase shift (where phase shift is defined on a zero-crossing basis) than the theoretical 90° of an unclipped integrator.

Either of these schemes can be used to ensure global stability, and the choice should be made according to which is more convenient to implement from a circuit design point of view.

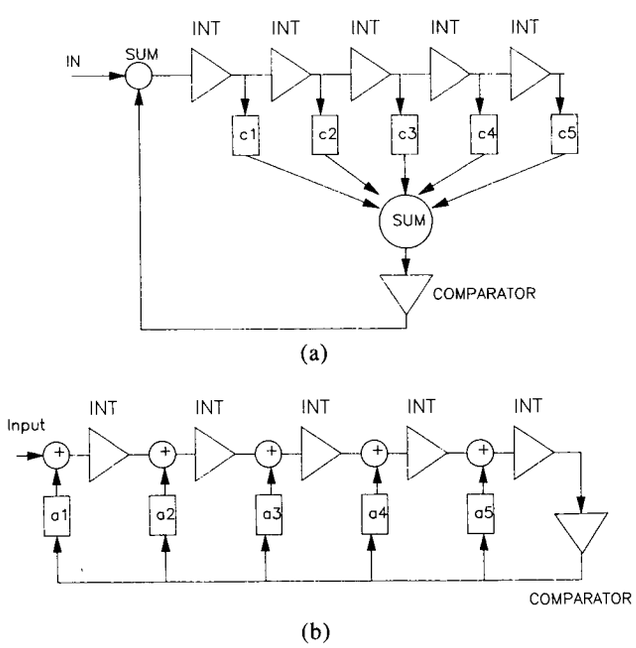

Fig. 10. (a) Fifth-order loop with all poles at dc. (b) Inverted loop with five poles at dc.

6 IMPLEMENTATION OF A STEREO SIGMA–DELTA 18-BIT A/D CONVERTER INTEGRATED CIRCUIT

In this section we describe a dual 20-kHz bandwidth, 18-bit sigma–delta A/D converter with an oversampling ratio of 64 and a noise floor that is 105 dB below full scale. The converter is implemented on two chips: a 10-V 3-μm CMOS modulator chip and a 5-V 1.5-μm CMOS decimator chip.

6.1 Noise-Shaping Modulator IC

The techniques described in this paper were used to design a fifth-order noise-shaping filter with an inverse Chebyshev characteristic for the noise transfer function [8]. The loop filter uses two resonator sections and one integrator, resulting in the noise transfer function shown in Fig. 9.

A block diagram of the circuit is shown in Fig. 12. The modulator chip contains a dual-output 3.0-V reference, clocking and interface circuitry, and two fifth-order, 1-bit, fully differential switched-capacitor sig-

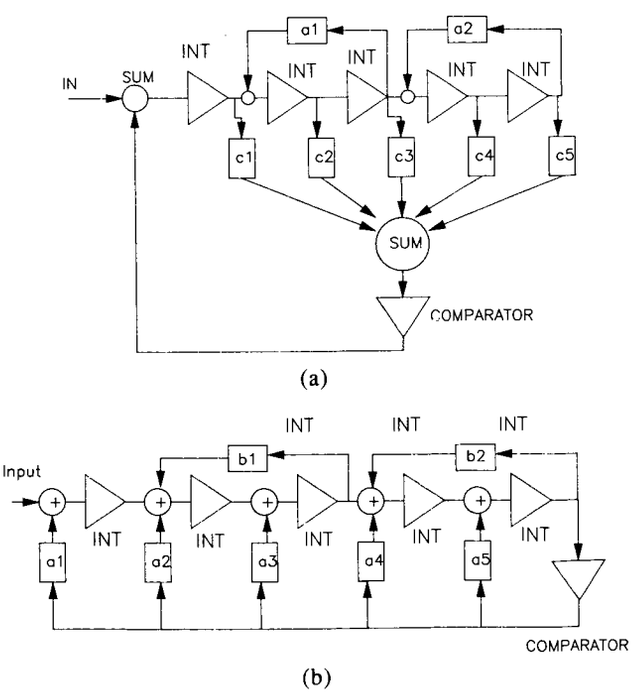

Fig. 11. (a) Fifth-order loop with two complex pairs and one real pole. (b) Inverted fifth-order loop with two complex pairs and one real pole.

ma–delta modulators. Each modulator contains five amplifiers, a comparator, switches, capacitors, and associated nonlinear stabilization circuitry.

6.2 Switched-Capacitor Loop Filter Design

Fig. 13 shows a single-channel single-ended representation of the switched-capacitor architecture used. Note that this architecture corresponds to the inverted architecture shown in Fig. 11(b). This design was chosen to avoid having to use an extra active summing amplifier at the comparator input [a requirement of the noninverted topologies of Figs. 10(a) and 11(a)]. Extra signal paths were added from the input to each integrator input to provide more flexibility in the choice of signal frequency responses. This only affects the signal frequency response and has no effect on the noise-shaping transfer function.

A fully differential architecture was used to increase signal swing and power supply rejection while reducing the effects of substrate coupling and digital feedthrough. Since switched-capacitor circuits are true sampling devices, even small amounts of high-frequency noise coupling through the supplies or any other path will fold down and cause noise and/or distortion in the audio band. The switches are sized to meet the requirement of 18-bit settling in one-half the clock period, and the *RC* bandwidths of the switch–capacitor combination may be 30 MHz or more. There is therefore little intrinsic filtering of high-frequency noise that may cause foldovers.

The reference switching scheme is shown in Fig. 14. The signals X and $\overline{X}$ are the true and complement outputs of the comparator. Instead of using both a positive and a negative reference, the feedback sense of the reference is inverted by swapping phase 1 and phase 2 clocks on the reference input sampling switches. In one case, the capacitor is short-circuited in phase 2 and in phase 1 it is connected between V_{ref} and the integrator input, delivering a positive packet of charge to the integrator. In the other case, the capacitor is connected between V_{ref} and ground on phase 2, and between ground and the integrator input on phase 1, delivering a negative charge packet to the integrator. Depending on the decision of the comparator, these charge packets are steered on phase 1 into either one side or the other side of the differential integrator. This scheme has the important advantage that the current drawn from the reference does not depend on the signal. Were this not true, some nonlinear version of the signal would appear on the reference and cause distortion, as the reference is a multiplicative input to the system.

The differential architecture is fully pipelined, with each integrator having an effective delay of one sample.

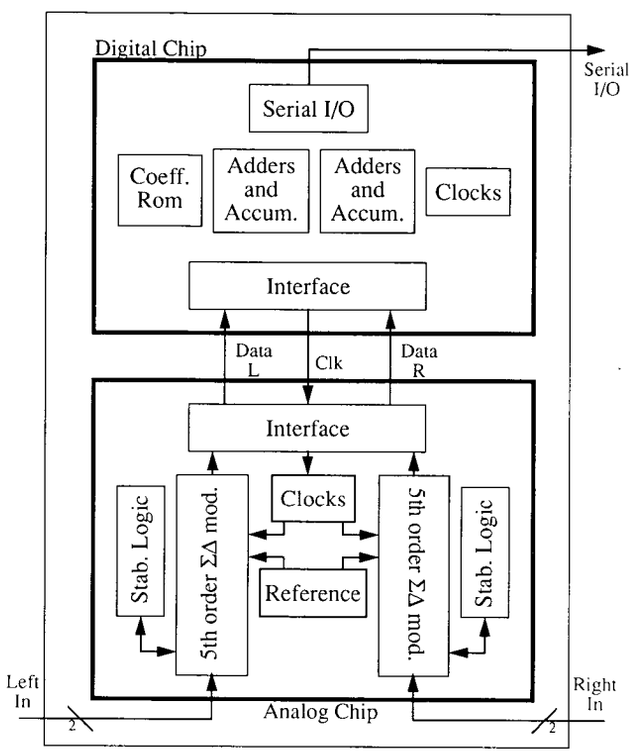

Fig. 12. Block diagram of converter.

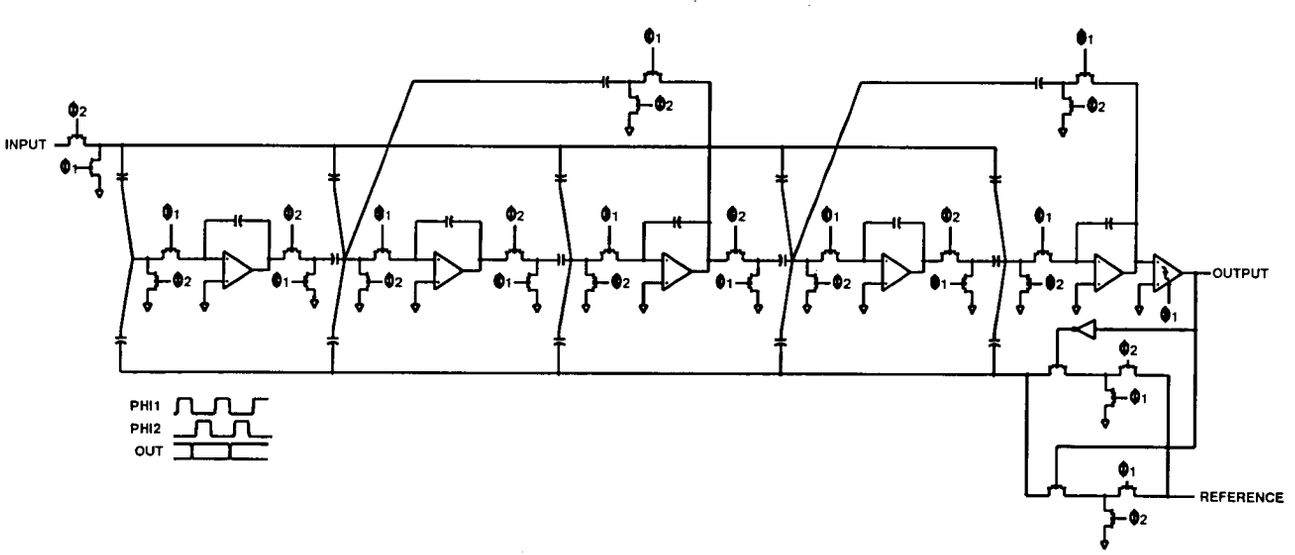

Fig. 13. Simplified switched-capacitor topology.

This avoids the double-settling problem, where two operational amplifiers are connected in series during the settling phase of the circuit. This would require much faster settling times from each operational amplifier.

The notches in the noise transfer function are implemented by feeding back around pairs of integrators, as shown in Fig. 13. As described, this feedback moves the open-loop poles of the feedback filter along the unit circle to frequencies in the audio band. It is interesting to note that in the open-loop configuration, these sections would in theory oscillate with constant amplitude at their resonant frequency. Fortunately the application of negative feedback to the system turns these open-loop poles into zeros on the unit circle, where they serve to reduce in-band quantization noise.

6.3 Circuit Noise Considerations

A fundamental noise source exists in all switched-capacitor circuits that arises from the finite resistance R of the MOS switch in conjunction with the $R-C$ bandwidth of the switch-capacitor combination. Every time a switch opens and leaves charge on a capacitor, there is a random error voltage due to the thermal noise of the switch resistance R. The power of this noise may be found from:

$$\text{Noise power} = KT/C$$

where

K = Boltzmans constant
T = temperature in degrees kelvin
C = capacitance in farads.

A variety of other noise sources besides KT/C noise also exist and careful attention was given to minimizing these sources. The theoretical quantization noise of the loop (as simulated using a difference-equation discrete-time simulator) is at the -112-dB level. One noise source that is frequently overlooked is the digital truncation noise that occurs at the decimator output when the accumulator is truncated to 18 bit. This noise is by definition at the 18-bit level (110 dB). This truncation is properly dithered by the analog circuit noise so as to remove any distortion effects that might otherwise occur.

The capacitors were sized from KT/C noise considerations, with the first-stage capacitors being the dominant noise source and hence the largest capacitors. These capacitors must be quite large to give 18-bit noise performance, and this results in large operational amplifiers with high bias currents to meet the 18-bit settling requirement. The total noise of the circuit is the power sum of the theoretical loop quantization noise, the KT/C noise, the digital truncation noise, and the operational amplifier thermal noise.

Fig. 15 shows the simulated and measured signal and quantization error transfer functions of the modulator where the comparator is replaced by a linear buffer. This turns the modulator into a switched-capacitor filter instead of a sigma–delta modulator, making it easy to measure these responses. The clock was slowed down in order to accommodate the external circuitry needed to make the measurement. Excellent agreement was found between simulated and measured results.

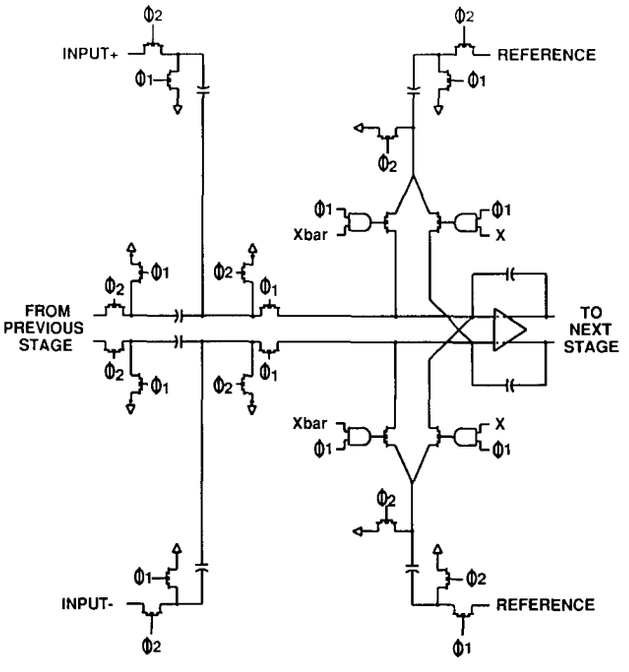

Fig. 14. Reference switching scheme with signal-independent reference load.

6.4 Stabilization Using Integrator Reset

As mentioned, higher order modulators are stable over a limited range of 1's density at the output. Instability in this modulator is sensed digitally, by counting the number of consecutive 1's or 0's in the modulator bit stream. A sufficiently long string of either 1's or 0's indicates modulator instability and triggers circuitry which resets the state in the integrators to put the modulator into a stable operating condition. When a signal continuously overloads the modulator for an extended period of time (dc, for example), the modulator goes into a mode where it detects overload, resets, comes out of reset, and after some time goes into overload again. As long as the average value of the 1-bit stream remains high enough, the digital clipper in the decimator will prevent any of this noisy behavior from appearing at the output.

6.5 Reference Design

A single 3.0-V bandgap reference with separate left and right output stages is included in the part, with a reference voltage being fed back to each channel independently using a switched-capacitor single-ended to differential converter and a selector controlled by the comparator, as represented in Fig. 3. The bandgap reference takes advantage of parasitic bipolar devices available in a standard CMOS process. These devices are stacked to reduce the effects of MOS

threshold mismatches on the output voltage and temperature drift performance.

6.6 Operational-Amplifier Design

The performance of the input amplifier in higher order loops determines overall converter performance. Therefore, to achieve the stated performance, the following features were designed into our front-end amplifier:

1) Fully differential, including a common-mode loop
2) Settling to 18-bit accuracy (23 μV) in one-half clock cycle (120 ns)
3) Gain linearity of 0.001%, or 10-μV peak nonlinearity over the entire 3-V peak output range
4) Chopper stabilization (between the clock phases) to eliminate offset and $1/f$ noise contributions
5) Input-referred thermal noise contribution below the modulator noise floor of −105 dB full scale.

A simplified circuit diagram of the amplifier is shown in Fig. 16. A single-stage folded-cascode architecture was chosen for both the differential and the common-mode paths. This topology has an inherently high bandwidth, which enables the amplifier to meet the settling time specifications. Compensation for both the common-mode and the differential-mode paths is provided using the switched-capacitor circuit capacitances surrounding the amplifier.

The common-mode output voltage is sensed using PMOS source followers on each output driving a resistive and capacitive summing network at the output. The summed voltage is compared to ground through an NMOS differential pair, and the difference current fed back to the amplifier to control the common-mode current.

The amplifier linearity is limited by impact ionization current in the NMOS output devices. This impact ionization current causes an effective nonlinear resistive load on the amplifier, and this will cause distortion when reflected through the finite gain of the amplifier to the input terminals. This problem is corrected by placing an NMOS transistor between the output and the NMOS current source and driving its gate with half the difference between the output and the minus analog supply. This has the effect of maintaining a low

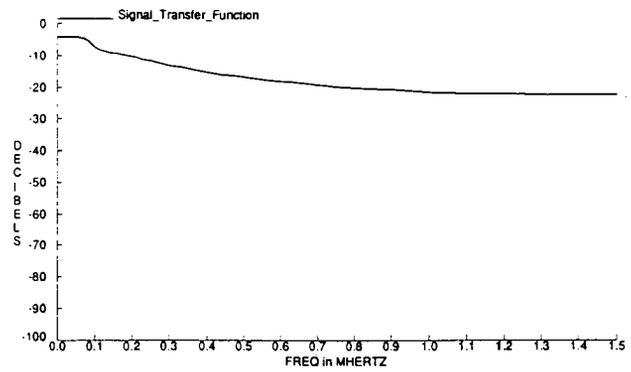

DC Signal Transfer Function 1.5MHz

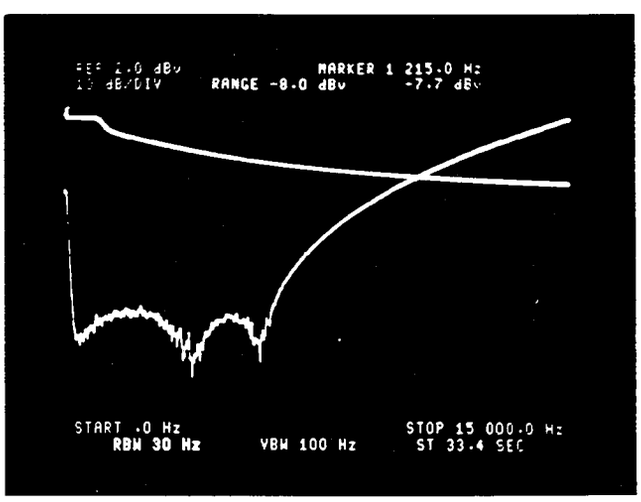

DC Error Transfer Function 50KHz

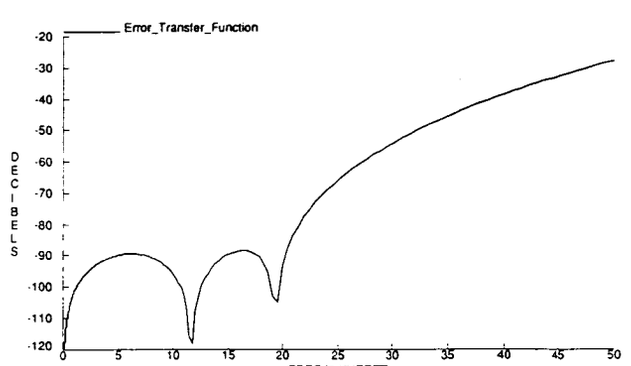

Fig. 15. Measured and predicted transfer functions.

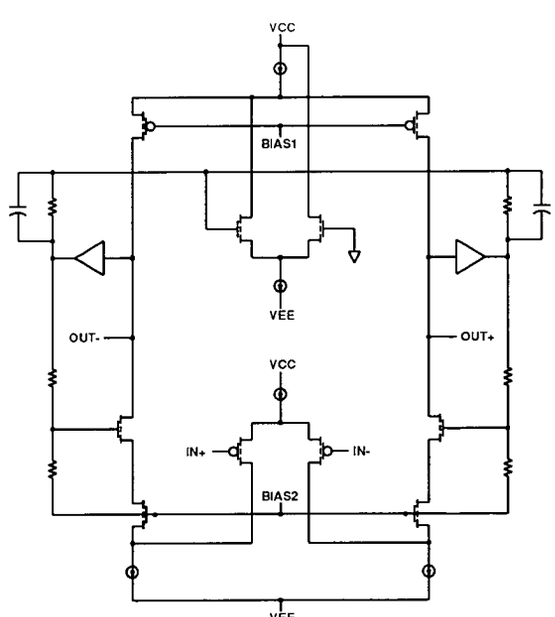

Fig. 16. Simplified schematic of Differential operational amplifier.

drain–source voltage across all NMOS signal transistors. Since impact ionization is a strong function of the drain-to-source voltage, this circuit dramatically reduces impact ionization. Without the correction circuitry, the third harmonic distortion products were only 88 dB below the fundamental signal. The correction circuitry improves the distortion performance by almost 20 dB.

7 DECIMATOR IC

The decimator chip has been described previously [9]. It is designed to have flat response to 20 kHz when operating at an output sample rate of 44.1 kHz. In-band images are suppressed by the stopband attenuation of the filter, which is more than 115 dB. The decimation filter frequency response is shown in Fig. 17.

The decimation is accomplished in one step rather than the more usual multiple steps. Parallel processing is used to implement a 4096-tap FIR filter with a 64-times oversampled 1-bit input. The 1-bit nature of the input signal can be exploited to design a multiplier-free filter structure. A multiply–accumulate operation with 1-bit data can be done with only an accumulator and a ROM.

One advantage of the one-step decimation approach is that the full stopband attenuation figure (−115 dB) is maintained over the entire region from 24 kHz up to 2.8 MHz (64 × 44.1 kHz–24 kHz). Multirate decimation structures usually exhibit periodic regions of poor attenuation due to the interaction between the image responses of the various stages of filtering.

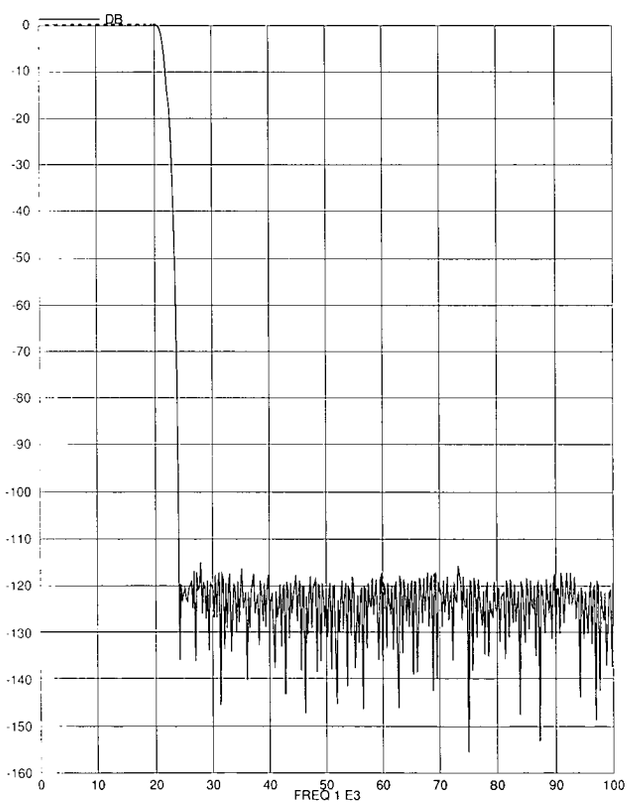

Fig. 17. Decimator frequency response.

Before data are output to the serial port, they are passed through a digital clipping circuit. This circuit prevents accumulator wraparound from appearing at the output.

8 SERIAL OUTPUT

The serial output format from the decimator is very flexible and has several different modes. The serial out pins are as follows:

1) Serial data out
2) Bit clock
3) L/R clock
4) Word clock
5) Bit clock field size select (16 or 32)
6) Mode control pins (2 pins decoded for 4 states).

The mode select pins control whether the serial port is in the master mode (where data and clocks are all outputs) or in the slave mode (where data and clocks are all inputs). In the master mode the data are justified toward the end of a 32-bit serial field and shift out MSB first.

There are several slave modes where bit clock, L/R clock, and word clock are controlled by the user. The word clock may be used in slave mode as an input to determine whether data are left-justified, right-justified, or triggered by a word-clock transition. The field size may be selected between 16 and 32 bits by the size pin. Obviously, only 16 bits may be shifted out when the field size is selected to be 16 bits.

In slave mode, all signals supplied externally are resynchronized to the master clock supplied by the user. This is necessary to ensure that there are no digital outputs that change just before the input is sampled on the modulator chip. Interference effects between these signals can easily degrade the performance to a mere 16 bits (or less).

9 MEASURED PERFORMANCE

The measured performance of the converter is shown in Table 1. These results indicate performance obtained from the first silicon. Fig. 18 is a plot of signal/(total harmonic distortion + noise) versus signal for the converter. The noise performance of the converter may be derived from the low-signal portion of the graph, where the THD + noise is dominated by noise and not distortion (unlike successive-approximation converters where low-level THD + noise is dominated by distortion). The flattening of this curve at high signal levels is due to harmonic distortion effects. Fig. 19 shows both the modulator and the decimator chips.

10 CONCLUSION

A design procedure has been given that results in stable high-order modulators with 1-bit quantizers. This procedure relies on a definition of the comparator gain that is derived from the dc transfer function of the

comparator measured in the sigma–delta loop. This low-frequency gain view of the quantizer can be thought of as a convolution of the input pdf to the comparator with the actual input–output transfer function of the comparator.

The authors have used this technique to design an 18-bit stereo sigma–delta converter IC with 105-dB dynamic range. This IC uses a differential switched-capacitor loop filter to implement a fifth-order stable 1-bit noise-shaping modulator. This performance is achieved in a simple CMOS process using no bipolar devices (except for parasitic pnp's in the reference) or thin-film resistors.

Table 1. Measured converter performance.

Converter noise floor	−105 dB FS
Interchannel crosstalk	−110 dB at 1 kHz
Input range	6 V (differential)
Signal bandwidth	20 kHz
Decimator passband ripple	0.001 dB
Decimator stopband attenuation	115 dB
Modulator die size	6.8 × 4.7 mm
Decimator die size	6.8 × 5.5 mm
Total power dissipation	1.1 W
Package type	28-pin 600-mil plastic DIP

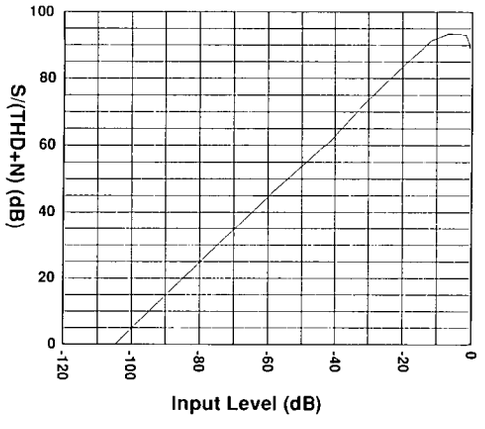

Fig. 18. Signal/(THD + N) versus input level.

11 ACKNOWLEDGMENT

The authors gratefully acknowledge the contributions of several persons to the success of this project: B. Scharf and NCR Semiconductor for process development; P. Sheppard, K. Hawley, and R. Berry for layout; C. Roberts and T. Goida for packaging; U. Moon and A. Wellinger for fixturing and measurements; and the members of the Analog Devices and Ensoniq design communities for their insight.

12 REFERENCES

[1] J. Candy, "A Use of Double Integration in Sigma–Delta Modulation," *IEEE Trans. Commun.*, vol. COM-33, pp. 249–258 (1985 Mar.).

[2] R. Gray, "Spectral Analysis of Quantization Noise in a Single-Loop Sigma–Delta Modulator with dc Input," *IEEE Trans. Commun.*, vol. 37, pp. 588–599 (1989 June).

[3] D. R. Welland, B. del Signore, and E. J. Swanson, "A Stereo 16-bit Delta Sigma A/D Converter for Digital Audio," presented at the 85th Convention of the Audio Engineering Society, *J. Audio Eng. Soc. (Abstracts)*, vol. 36, p. 1034 (1988 Dec.), preprint 2724.

[4] R. W. Harris, "Enhanced Delta Modulation Encoder," U.S. patent 4,509,03, filed 1982 Dec. 1, assigned to Gould Inc.

[5] R. W. Adams, "Companded Predictive Delta Modulation: A Low-Cost Conversion Technique for Digital Recording," *J. Audio Eng. Soc.*, vol. 32, pp.

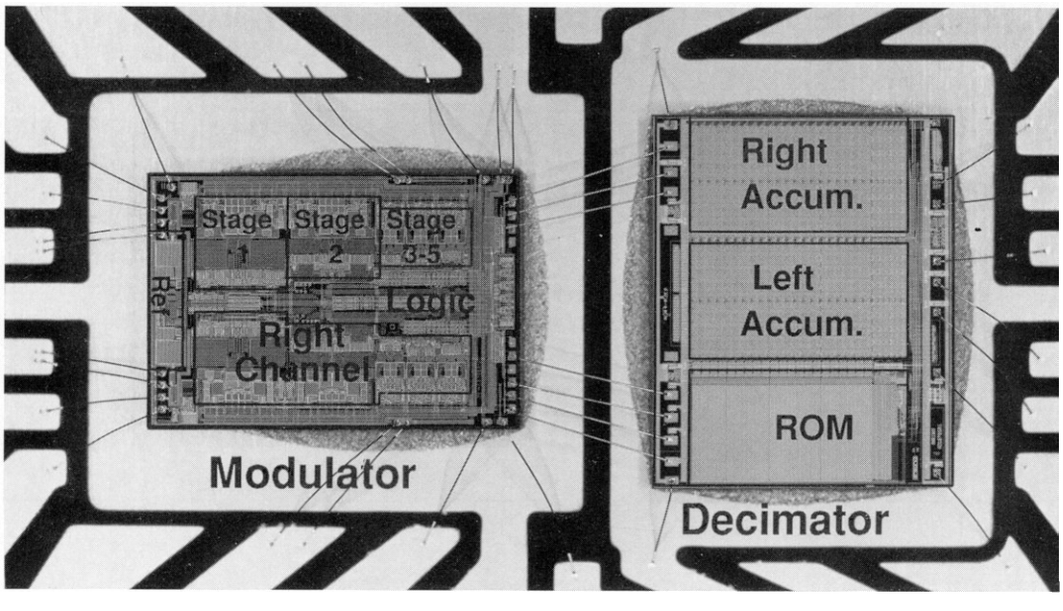

Fig. 19. Integrated circuit.

659–672 (1984 Mar.).

[6] 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.).

[7] W. L. Lee and G. C. Sodini, "A Topology for Higher-Order Interpolative Coders," in *Proc. 1987* ISCAS.

[8] P. F. Ferguson, A. Ganesan, and R. W. Adams, "One-Bit Higher Order Sigma–Delta A/D Converters," in *Proc. 1990 ISCAS*, vol. 2, pp. 890–893.

[9] J. Dattorro, A. Charpentier, and D. Andreas, "The Implementation of a One-Stage Multirate 64:1 FIR Decimator for Use in One-Bit Sigma–Delta A/D Applications," presented at the AES 7th International Conference: Digital Audio, Toronto, Ont., Canada, 1989 May 14–17.

THE AUTHORS

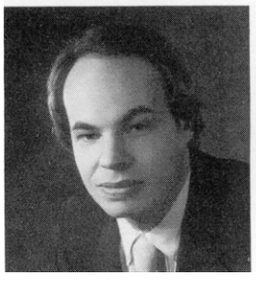

R. Adams

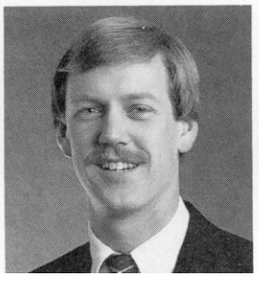

P. F. Ferguson, Jr.

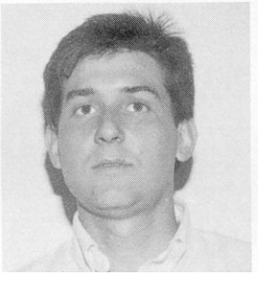

S. Vincelette

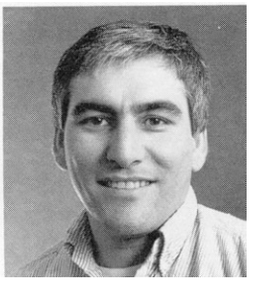

A. Volpe

Robert Adams was born in Syracuse, NY, in 1954. He studied electrical engineering at Tufts University in Medford, MA, where he graduated in 1976. In 1978 he joined the staff of dbx Inc., where he was responsible for the design of a broad range of audio products.

Since 1982 Mr. Adams has specialized in the areas of analog-to-digital conversion and digital signal processing. From 1983 to 1989, he was director of research at dbx. In April of 1989, he joined the staff of Analog Devices, where he is involved in converter and signal processing products aimed at the audio market. A member of the AES, Mr. Adams holds a number of patents in the area of audio processing.

•

Paul F. Ferguson, Jr., was born in Portland, ME, in 1962. He received the A.B. degree from Dartmouth College in 1984, and the M.S.E.E. degree from MIT in 1986. He joined Analog Devices in 1986, where he has been involved in converter design in the semiconductor division as a senior design engineer. His interests include $\Sigma\Delta$ converters, sample-hold amplifiers, and self-calibrating converters in CMOS, BiCMOS, and Bipolar technologies.

•

Scott Vincelette received the B.S.E.E. degree from University of Connecticut in 1986, and the M.S.E.E. degree from Lehigh University in 1989. In 1986 he joined Wang Laboratories, Lowell, MA, where he worked on microprocessor design. Since 1989 he has been with Analog Devices, Wilmington, MA, involved in oversampled converter design for audio applications.

•

Apparajan Ganesan[1] graduated from the Indian Institute of Technology, Madras, with a bachelor's degree in electrical engineering in 1975. In 1976 he received a master's degree in system theory from the State University of New York in Stony Brook, Long Island.

He has worked at General Instruments Corporation, Bell Labs, RCA, and several divisions of AT&T, and has been a consultant to the Hazeltine Corporation. He has also served as an adjunct faculty at several schools and has taught courses in the area of device physics, computer-aided design, modeling, and simulation and integrated-circuit design. Mr. Ganesan has authored and coauthored several papers in the area of linear signal processing and has been awarded five patents. He was the New Business Development Manager at Analog Devices. His current interests are in the areas of computer-aided design, integrated circuit design, and system definition and engineering.

•

Anthony Volpe received a Bachelor of Science degree in Electrical Engineering from Southeastern Massachusetts University in 1979. He joined Analog Devices semiconductor division where he has worked on temperature transducers, voltage references, and the design of sampling A/D converters.

•

Robert J. Libert[1] received a B.S.E.E. from Lowell Technological Institute (now the University of Lowell), Lowell, MA, in 1973. He joined Analog Devices Inc. in 1973 and worked as a converter product/test/trim engineer. In 1984, he became a design engineer and developed a series of monolithic zener references. Recently he has been working on digital audio converter products.

[1] Photograph was not available at press time.