Self-Calibration Techniques for a Second-Order Multibit Sigma–Delta Modulator

John W. Fattaruso, *Member, IEEE*, Sami Kiriaki, *Member, IEEE*, Michiel de Wit, *Senior Member, IEEE*, and Greg Warwar, *Member, IEEE*

***Abstract*—Design techniques for self-calibration of the digital-to-analog converter DAC in a multibit sigma–delta modulator are described. When used in conjunction with dynamic element matching, self-calibration provides linearity performance suitable for digital audio applications. The dynamic element matching circuitry provides the mechanism of determining device mismatch for self-calibration. Practical circuit details and an effective calibration method are discussed. Test results from a 1-μm CMOS test chip are presented. In this test system, a second-order loop with a 3-b quantizer achieves an 89-dB dynamic range and −91-dB harmonic distortion after calibration. In addition, a new method of detecting the presence of tones is described, using the entropy of the spectrum of the decimation filter output.**

I. Introduction

MANY approaches to designing oversampled sigma–delta modulators for high-precision analog-to-digital conversion have been recently reported in the literature [1]–[5]. One class of high-resolution modulators uses single-bit quantizers and feedback loops with orders greater than two to achieve the required quantization noise and harmonic distortion specifications. Another option is using a multibit quantizer, which allows a second-order loop to exhibit low quantization noise. This configuration has significant advantages over the single-bit system. Because of unconditional stability of a second-order loop, recovery from overload is assured. The portion of the decimation filtering that is performed by a sinc filter can use a simple, low-order filter. The general behavior of the modulator will approximate a linear feedback loop more closely than will a single-bit loop. Consequently the tone energy will be low, and the modulator will be able to follow a large-signal input more smoothly as the input level approaches full scale and avoid the increased noise that is present with single-bit quantizers. It is well known, however, that the overall linearity of the modulator will be limited to that of the multibit digital-to-analog converter (DAC), and some mechanism for correcting the linearity errors of a monolithic DAC is needed. While other authors have chosen digital error correction [6], the design reported here uses analog self-calibration [10]–[12] and dynamic element matching [7]–[9] to enhance DAC linearity.

The test chip for the self-calibration approach is a fully differential CMOS implementation for a single 5-V supply. Charge-redistribution DAC's are used in the modulator, and such DAC's built in untrimmed CMOS technology are typically limited to about 10 b of linearity due to random mismatches between nominally identical unit capacitors. These linearity errors due to capacitor mismatch are primarily removed by self-calibration of the unit capacitor values. This will improve the DAC linearity from about 10 b to about 14 b. The remaining capacitor mismatch errors may then be randomized with dynamic element matching. Dynamic element matching is a technique that will statistically eliminate the mismatches between the discrete devices in the DAC by assigning each unit DAC capacitor in a random order to the codes being converted. Although this results in the effective elimination of harmonics, white noise that is proportional to the degree of physical capacitor mismatch is introduced into the modulator. The test chip demonstrates that self-calibration of the eight DAC capacitors can reduce the inherent device mismatch to a level where high linearity and low noise are simultaneously possible when random dynamic element matching is used. The randomizing circuitry is also an essential tool for measuring the physical device mismatch, and it is used in the process of self-calibration.

Manuscript received May 17, 1993; revised July 30, 1993.

J. W. Fattaruso, S. Kiriaki, and M. de Wit are with the Integrated Systems Lab, Texas Instruments Inc., Dallas, TX, 75265.

G. Warwar is with Vitesse Semiconductor, Camarillo, CA.

IEEE Log Number 9212328.

II. Modulator Description

Refer to Fig. 1 for a block diagram of the modulator. The second-order loop is built around two switched-capacitor integrators. These are followed by a 3-b flash analog-to-digital converter (ADC), using an array of capacitive voltage dividers and seven latch comparators. This ADC is designed with seven transition levels between a full scale of ±3-V differential. Because of its location in the loop, any offset or matching errors in the ADC do not limit linearity of the modulator. Simulations show that the transition voltage levels of the ADC may have random errors on the order of 100 mV without degrading the precision of the modulator. The 7-b thermometer code output of the ADC comparators is converted to an output Gray code, and it is also fed back to drive the unit capacitors in the DAC's. Gray code is used to minimize the switching noise contributed by the three digital pad drivers. The DAC is designed with eight output voltages between a full scale of ±3.5-V differential. The difference between the ADC and DAC full-scale levels accounts for a factor of 4/3 of the second-stage gain of two shown in Fig. 1, with the additional factor of 3/2 coming from physical voltage gain in the second-stage integrator. This choice allows maximized statistical use of the output voltage swing of both integrators. Between the

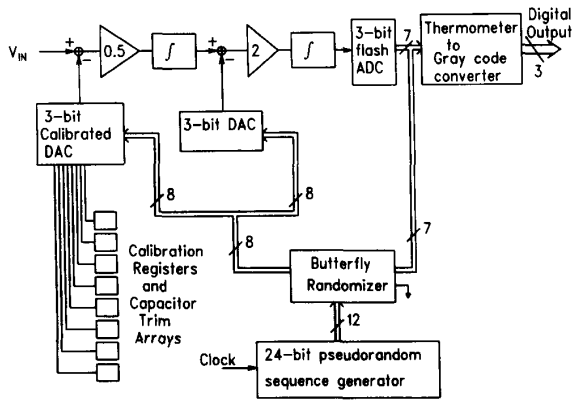

Fig. 1. Modulator block diagram.

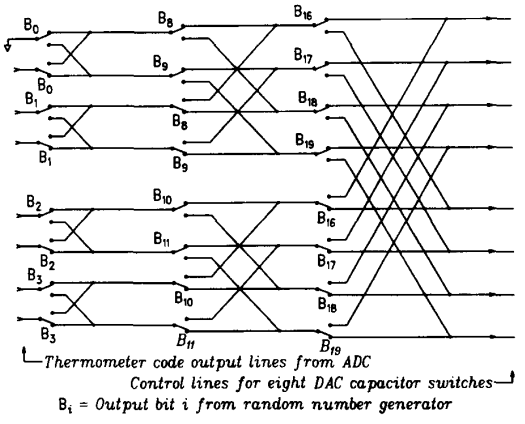

Fig. 2. Butterfly randomizer network.

ADC comparators and the DAC capacitor switches there is an intervening randomizer block that assigns unit capacitors in the DAC's to corresponding thermometer code lines in one of 4 096 combinations. The randomizer is a "butterfly" array of CMOS switches [9], shown in Fig. 2. The particular ordering used at any modulator clock period is determined by the state of 12 b taken from a 24-b pseudorandom number generator. When the modulator is used to actively convert samples, the clock of the random number generator is run at the same rate as the modulator, giving a new pseudorandom reordering for each input sample.

System simulations were again used to determine that the linearity of the DAC feeding back into the *second* integrator does not significantly affect the overall modulator linearity. Introducing random capacitor mismatch at the 10-b level for this DAC in the simulations did not significantly affect noise or distortion performance of the modulator, and therefore this DAC need not be calibrated. It was also clear from simulations that the capacitor voltage coefficient typically observed for the metal-to-polysilicide capacitors in the process to be used would not be responsible for any significant generation of harmonics.

The eight unit capacitors in the DAC feeding the *first* integrator may be trimmed by loading digital correction values

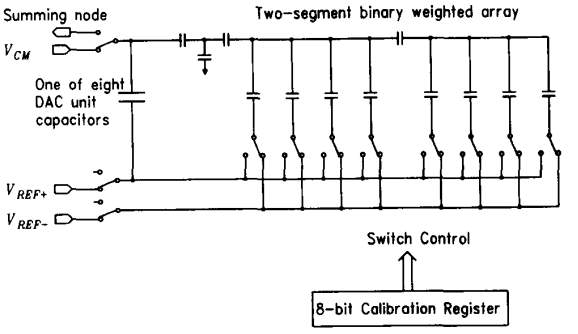

Fig. 3. Capacitor calibration networks.

into eight calibration registers. Fig. 3 is a simplified schematic of the capacitor trimming circuitry for each of the eight DAC unit capacitors. Although fully differential switched-capacitor circuits are used in the modulator, a simplified half-circuit is shown in the figure. The effective capacitive value of each of the unit capacitors may be trimmed by a parallel capacitive network. Binary bits loaded into the corresponding 8-b calibration register select switch positions for a series-coupled binary weighted capacitor array. The parallel capacitance of this array can alter the unit capacitor value over a total range of one part in 256, so the trimming resolution is one part in 2^{16}. A total of eight of these correction arrays and calibration registers are needed in the modulator.

Simulations of the modulator system must also be used to find an appropriate design of the pseudorandom number generator. The sequence governing the rotation of the DAC capacitors will repeat after some time interval, since it is generated by a practical number generator of finite length. The mechanism of dynamic element matching is a complex combination of modulation of the input signal with a low-frequency rotation sequence in the presence of the nonlinearity of the DAC characteristic. When simulations are performed on the modulator and decimation system, assuming a Gaussian distribution of capacitor mismatch errors, the baseband output spectrum exhibits many frequency bins containing concentrations of harmonic energy. As the length of the pseudorandom sequence is increased, this harmonic energy is distributed among an increasing number of bins, reducing the peak energy in any single bin. Simulations showed that selecting a shift register length of 20 b reduces the harmonic energy far below the expected level of device thermal noise. Since adding a few extra bits to the shift register consumes minimal chip area, a conservative length of 24 b was chosen. Of these 24 output bits, 12 were arbitrarily selected to drive the switches of the butterfly randomizer.

The method of random number generation must also be chosen carefully. Hardware random number generators often use linear feedback shift registers, where a logical exclusive OR of high-order bits are fed back to recirculate into the low-order shift input. Although this method may be satisfactory for generating bit sequences interpreted as binary numbers, it is not appropriate for this case where rotating patterns of bits cause repeatable patterns of capacitor application in the DAC's, visible as baseband harmonics.

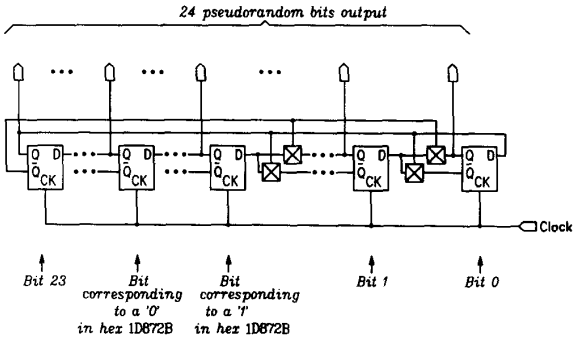

Fig. 4. Psuedorandom number generator.

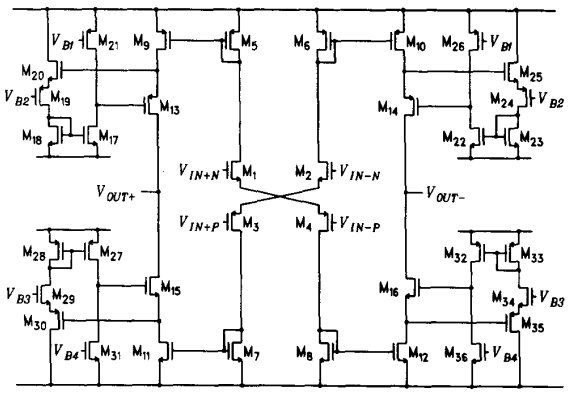

Fig. 5. Class AB op amp for first integrator.

The method used for pseudorandom number generation in this system is one usually employed in software routines [18], [19]. It is also based on a shift register, where bits are shifted to more significant locations, with a "zero" bit inserted into the least significant bit, on each clock. However, if a "one" bit is shifted out of the most significant bit, then the entire contents of the register is exclusive ORed with a special binary constant. For a 24-bit register the appropriate constant is hexadecimal 1D872B. This algorithm can be implemented very efficiently in CMOS technology with the circuit shown in Fig. 4.

Fig. 5 is the fully differential op amp used in the first integrator. It is designed to give a gain of about 80 dB with a single 5-V supply. Simulations show that a high-gain amplifier is needed with the nonlinear large-signal transfer characteristic typical of practical op amps. Not shown in the figure are the conventional switched-capacitor common-mode feedback devices [14]. The opamp uses a gain-enhancement cascode technique [15], [16] with a new topology for the auxiliary cascode feedback amplifiers (for example, M_{17}–M_{21}). This feedback amplifier design has the advantage of allowing the maximum output swing possible, within a few hundred millivolts of the supply rails, but only requiring a few devices per amplifier.

The op amp input devices M_1–M_4 are in a cross-coupled class AB configuration. Usually this arrangement is driven by a pair of level-shift devices that establish the quiescent current through the amplifier by setting the bias voltage between n-

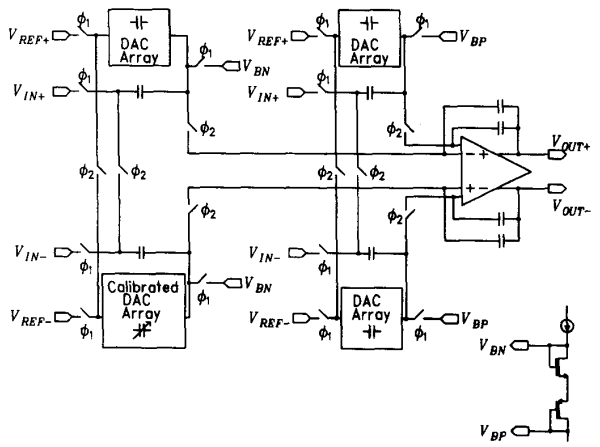

Fig. 6. First integrator schematic.

channel and p-channel gates [17]. However, these level-shift devices can contribute a large amount of input-referred noise because there is no isolating voltage gain between them and the input terminals. The first integrator in this modulator uses two identical switched-capacitor feedback networks around the op amp, as shown in Fig. 6. One network drives the n-channel differential inputs V_{IN+N} and V_{IN-N}, and the other the p-channel inputs V_{IN+P} and V_{IN-P}. The quiescent bias for the op amp is established by precharging these capacitor networks to two different common-mode bias levels, V_{BN} and V_{BP}. These voltages are derived across diode-connected devices that match the n- and p-channel op amp input transistors. The kT/C noise contributed by the integrator switches is minimized by increasing the size of the capacitors over what would be required in a conventionally biased integrator. Circuit optimization of the op amp device widths showed that it was possible to provide greater transient current to slew this increased capacitance and consume less overall power than in the conventional case of many noise-contributing devices operating at high bias current levels. Only one of the four DAC capacitor arrays actually needs to be trimmed to make the differential charge inserted by each unit capacitor pair matched with the others. Slight differences in common-mode charge will be attenuated by the common-mode rejection of the system.

The second integrator is of a conventional design, with a two-input class AB op amp using level shift devices. Minimizing input referred noise at this integrator is not important for low modulator noise.

III. Calibration

Because of the dynamic element matching, output energy concentrated at harmonics of the input signal due to capacitor mismatch will be eliminated, but as a tradeoff white noise will be introduced into the modulator. The power of this white noise is proportional to the variance among the DAC capacitor values. If the DAC capacitors were ideally trimmed to exactly match each other in value, then any random reordering of how these capacitors are used will not introduce any modu-

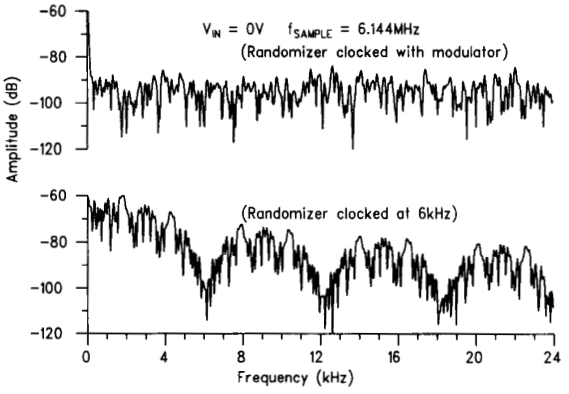

Fig. 7. Measured noise during calibration.

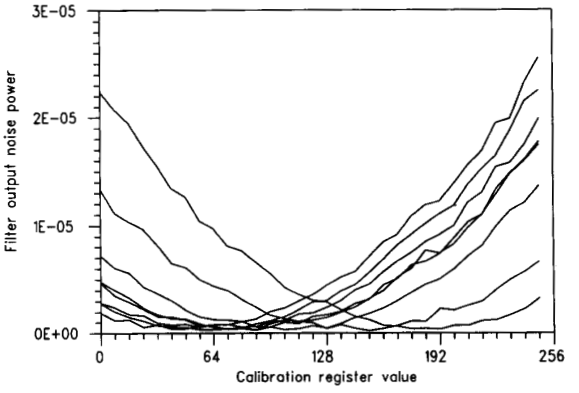

Fig. 8. Calibration noise versus calibration register value.

lator noise. Therefore any noise power that is introduced by dynamic element matching is a direct measure of the degree of capacitor mismatch. This noise power may be measured by simply calculating the variance of the digital samples at the output of the decimation filter with the analog input voltage set to zero. The modulator may then be calibrated by adjusting the eight 8-b calibration register words with a multidimensional numerical optimization algorithm to minimize this power. When the optimum values for the calibration registers have been found, the effective values of all the DAC capacitors will be close to identical.

For this calibration process to be effective, the noise from the randomized DAC capacitors must be the dominant noise contribution at the output. This process is made statistically valid in the presence of modulator quantization and device noise by driving the random number generator clock *during calibration* at a frequency equal to the cutoff frequency of the decimation filter that follows this modulator. For example, with an oversampling ratio of 128 to 1, the randomizer would be clocked at a rate of 1/256 times the modulator clock during the calibration phase. This injects *colored* noise into the modulator loop containing most of its energy at frequencies that will survive the decimation filter, and with a power proportional to the extent of capacitor mismatch. This low-frequency noise energy will pass through the modulator and filter and will appear much larger than the intrinsic modulator noise energy. Therefore the numerical routine to minimize this noise power will converge where all capacitors have essentially equal values. After calibration the randomizer clock can be set back to the modulator clock frequency to allow active conversions to be run.

The calibration process can be carried out effectively at test time with an external processor, followed by the optimum calibration register values set into ROM on chip. This technique has been found to be effective for self-calibration of a charge-redistribution ADC, as the required calibration correction for capacitor ratios has been found to be stable over temperature and lifetime variations [13]. The processor would read the digital stream from an appropriate decimation filter, compute the statistical variance and trim the calibration values appropriately. Note that *no precision analog source* is needed in the calibration process.

IV. Experimental Results

The test chip was built in a 1-μm CMOS process with metal-to-polysilicide capacitors. It includes the modulator, the pseudorandom number generator, and the butterfly randomizer. The calibration registers on chip are a static shift register, and the randomizer clock is fed externally during calibration. In a practical version of this modulator an additional digital frequency divider would be needed to generate the randomizer clock during calibration. For all measurements the modulator was clocked at 6.144 MHz. The decimation filtering, with an oversampling ratio of 128 to 1, was performed externally by a test processor.

The measured results in Fig. 7 show the effectiveness of slowing down the randomizer clock during calibration to within the passband of the decimation filter. The plots are the spectra found at the output of the decimation filter with the input voltage set to zero. For this test the DAC capacitor calibration registers are all set to nominal values, and there is capacitor mismatch at about the 10-b level due to processing variations. There is a significant level of noise introduced by running the randomizer at the modulator rate, as seen in the top spectrum, but any progress toward calibration would quickly lower this noise into the modulator noise floor. The lower spectrum was measured with the random number generator running at an example rate of 6 KHz. Now the noise introduced by randomizing the mismatched capacitors is colored, and it is concentrated at frequencies that survive the decimation filter. The increase in power level at the filter output is evident, providing a statistically dominant mismatch measurement for the entire calibration process.

One-dimensional variations in the calibration cost function are shown graphically in Fig. 8. This plot is the measured noise power at the filter output as each of the eight calibration registers are swept over their full range. Only one register is varied for each curve, with the other seven held at their optimum values. When the noise level reaches a minimum, capacitor matching is optimum. Note that since all capacitors may be trimmed, there are actually many possible combina-

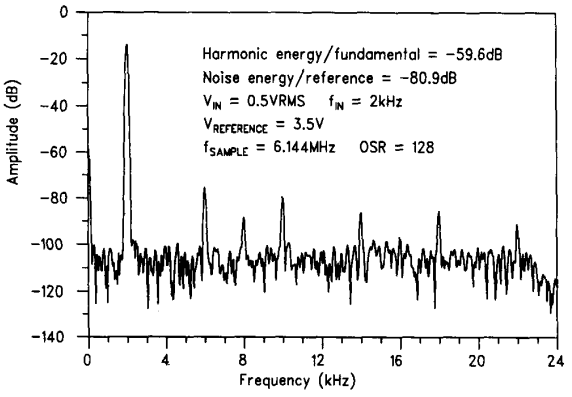

Fig. 9. Uncalibrated DAC and randomizer off.

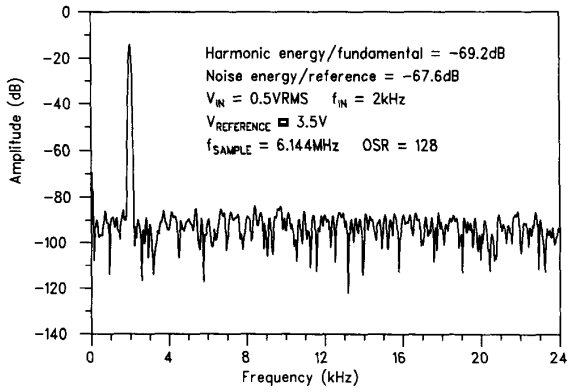

Fig. 10 Uncalibrated DAC and randomizer on.

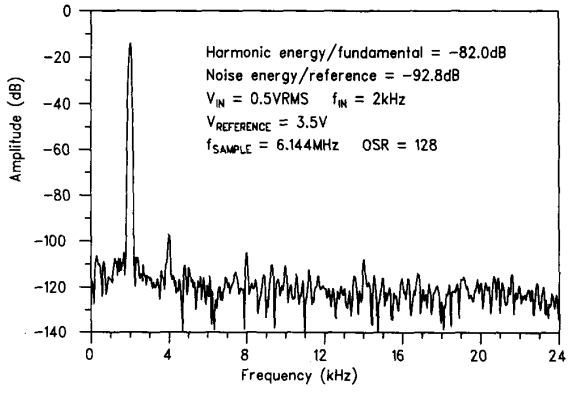

Fig. 11. Calibrated DAC and randomizer off.

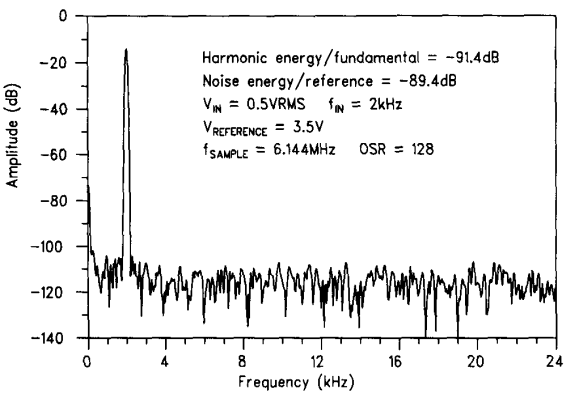

Fig. 12. Calibrated DAC and randomizer on.

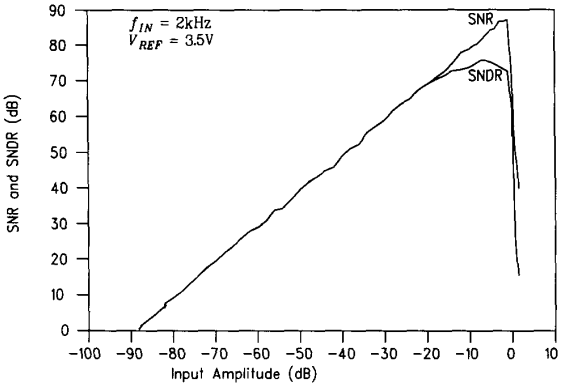

Fig. 13. SNR and SNDR versus input amplitude.

tions of calibration register values that minimize the variance among their values. Any numerical minimization routine can iteratively find an optimum combination of register values. For the experimental testing, a simple gradient descent algorithm was used.

Figs. 9 through 12 are measured test results with a 0.5-VRMS low-distortion sine wave at 2 kHz applied at the input. The figures are four audio passband output spectra of the decimation filter. The first plot in Fig. 9 shows the spectrum at the filter output with all of the calibration registers set to a nominal value and the randomizer turned off. (This is the case of a simple second-order modulator with an uncalibrated DAC.) The harmonics of the input signal are those due to the nonlinearity of the first-stage DAC, originating from capacitor mismatch at about the 10-b level. The next plot in Fig. 10 is the result of turning the randomizer on. The energy in the harmonics has been spread out into white noise by the dynamic element matching, raising the noise floor by about 13 dB. Next, in Fig. 11, the randomizer is turned back off and the calibration registers are loaded with values derived from the optimization procedure. Note that the linearity of the DAC is greatly improved, resulting in harmonic energy about 22 dB lower than that in the uncalibrated case. Finally, in Fig. 12, the randomizer is turned on with these optimum calibration words in place. The remaining nonlinearity is spread out to raise the noise floor by only about 3 dB, leaving both harmonics and noise at about the −90-dB level.

Fig. 13 is the characteristic of the modulator as the input level is swept. Note that a high signal-to-noise ratio is maintained very close to full scale by the multibit quantizer. The signal-to-noise plus distortion curve tails off due to nonlinearity in the op amps and the input sampling network.

Although some theoretical work seems to predict that all quantization noise in a second-order loop with a 3-b quantizer

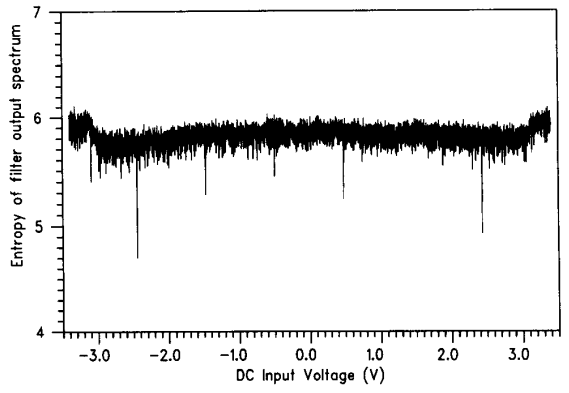

Fig. 14. Spectral entropy versus dc input voltage.

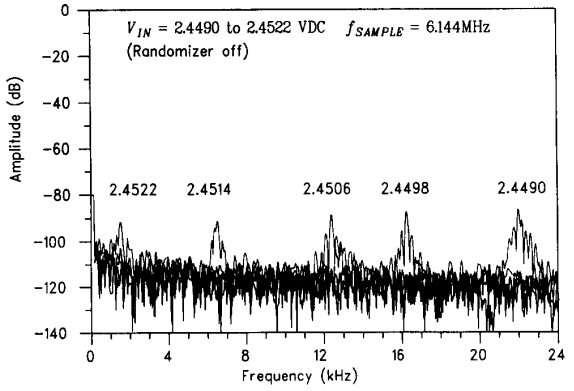

Fig. 15. Tone energy at various dc input voltages.

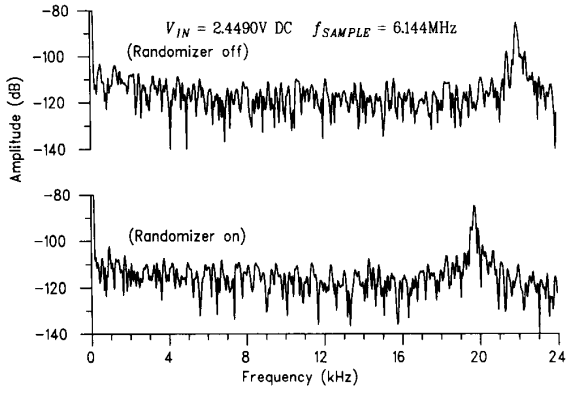

Fig. 16. Tone energy with and without randomization.

will be white [20], there are a few low-level tones that are measurable from the test chip. The tones exhibit a high-pass shaped characteristic in the frequency domain as the modulator dc input level is varied. This usually is observed with tones inherent to the modulator system. However, since all attempts to simulate these tones have failed so far, a mechanism of parasitic feedback within the test system cannot be ruled out as a cause.

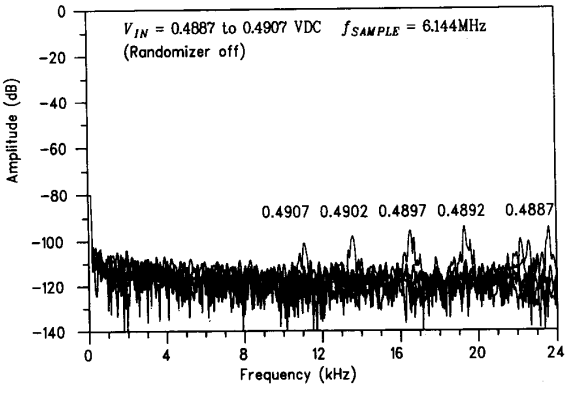

Fig. 17. Tone energy at various dc input voltages.

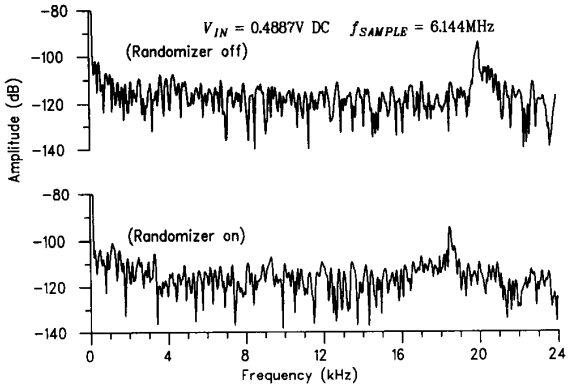

Fig. 18. Tone energy with and without randomization.

Fig. 14 shows the result of a systematic search for tones. The dc input level is swept with an increment of 1 mV, and the *entropy* of the spectrum found at the decimation filter output is measured. Entropy is usually a measure of structure or information content in a statistical system. This same measure may be applied to a frequency spectrum as a measure of structure. Spectral entropy will be high for white noise, and it will be lower for any spectrum where some discrete frequencies are more energetic than the rest. This measure is inherently independent of the actual noise level of the spectrum. Therefore, the entropy is a natural quantification to be used in a search for tones, where the presence of one or more predominant frequencies over a noise floor of any level is to be reported.

Assume that P_i is the power found in frequency bin number i in the discrete Fourier transform of the decimation filter output. Then define

$$\rho_i = P_i / \sum_i P_i$$

as the fractional bin power. The entropy is then defined in the usual way as

$$S = -\sum_i \rho_i \cdot \ln(\rho_i).$$

For a spectrum of white noise, $\rho_i = 1/n$, where n is the number of frequency bins. In this case $S = -\ln(1/n)$, the

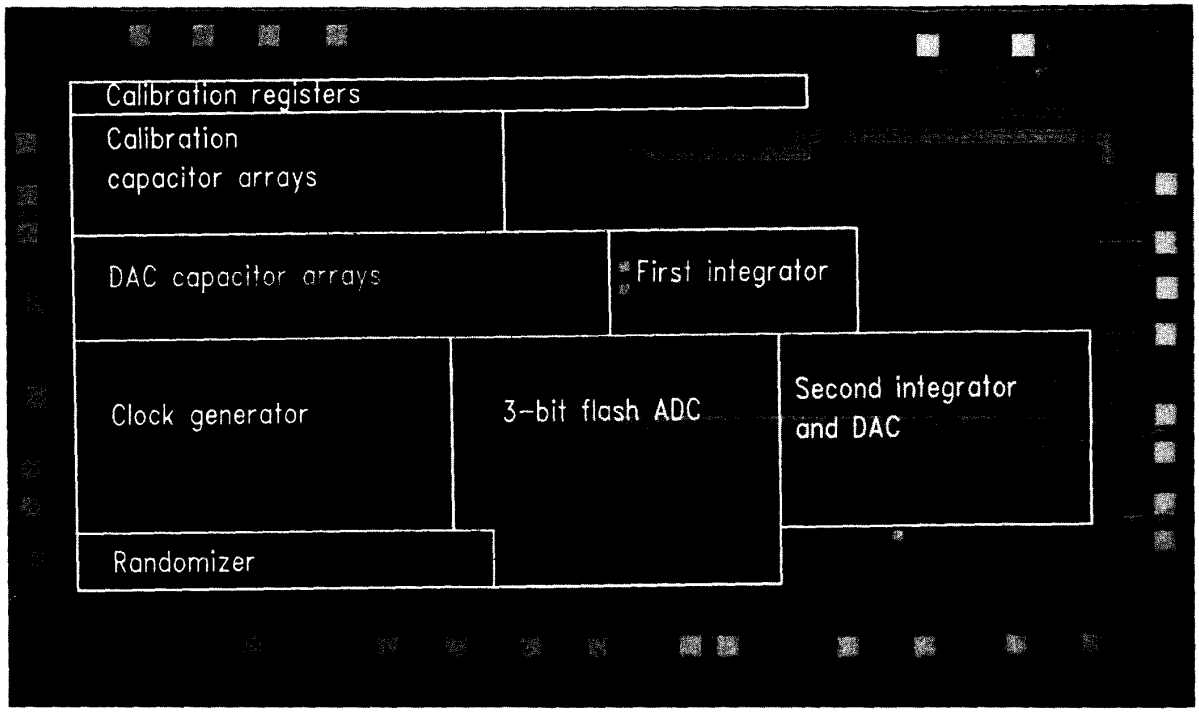

Fig. 19. Chip photomicrograph.

highest possible value. For a 1024-point FFT, n = 512 and S = 6.238. The entropy of most of the data points plotted in Fig. 14 approaches this theoretical limit. At any dc input level where there is structure in the spectrum, the entropy will be lower than its maximum limit. Tones in the modulator are indicated by the downward spikes in the figure. In practice, entropy should be calculated for a spectrum accumulated as the time average of several measured spectra. For Fig. 14, four measured spectra were averaged to give a statistically meaningful computation of entropy. As the dc input level approaches full scale, the modulator approaches overload and the noise level increases dramatically. The slightly increased entropy seen at the extremes of Fig. 14 occurs when the noise level has increased to unusable levels.

As can be seen from Fig. 14, tones in this modulator are visible centered around the few dc input voltages that correspond to the DAC output levels. Fig. 15 shows the migration and shaping of the worst-case tone through the baseband as the dc input is slowly varied. This tone is around the DAC output level at 2.5 V. As can be seen in Fig. 16, turning the randomizer on has little effect on the tone energy, other than a slight effective dc level shift. Figs. 17 and 18 show similar results from a tone of much lower energy around 0.5 V. The frequency shaping of this tone can be more clearly seen.

The DAC's in this modulator have output levels at ±0.5, ±1.5, ±2.5 and ±3.5 V. Since there is no DAC output level at zero volts, no tone was found that would be present in the case of a low-level input signal or an idle channel. The few tones that are present at these large dc offsets should be swept through with large signals. In this sense a multibit quantizer with levels spaced to step over 0 V gives an effective tone behavior that is superior to single-bit quantizers with tones in the neighborhood of 0 V.

Fig. 19 is a die photo of the test chip. The circuit occupies 5800 μm × 3200 μm of silicon area in the 1-μm process (including bond pads), and consumes 100 mW of power from the 5-V supply with the randomizer running.

V. Summary and Discussion

ADC's in the past have been calibrated using algorithmic methods to calculate the necessary correction values for capacitor mismatch. This paper demonstrates the technique of using noise power measurements in a sigma–delta modulator for self-calibrating a multibit DAC. This method uses randomization circuitry that lowers distortion further with dynamic element matching.

The noise floor of −89 dB observed in testing of this modulator was much higher than expected. During design the contributions of the expected noise sources of modulator quantization, amplifier device thermal noise and switch kT/C noise had all been determined through simulations to be less than 108 dB below the full-scale signal level. The multibit DAC easily allowed this low level of theoretical quantization noise, and the extra complexity of the self-calibration system was viewed as a worthwhile tradeoff for superior stability and tone behavior. During testing, however, the observed noise level did not follow the dependence on frequency, oversampling ratio or chip temperature that would lead to its explanation as any of the expected noise sources. Since the observed noise

floor did vary significantly with the power-supply voltage fed to the on-chip multiphase clock generator, it appears the limitation is due to digital switching noise coupled through the substrate in this particular implementation. If the modulator system detailed here is fabricated with techniques for reduced substrate noise now being discussed in the literature [21], we expect that its performance can compare favorably to higher order modulators with single-bit quantizers, without their attendant problems of potential instability and intrusive tones at low signal levels.

ACKNOWLEDGMENT

The authors would like to gratefully acknowledge the contribution of technical discussions with B. Brandt, R. Hester, Y. LeDuc, K. Tan, J. Hellums, J. Hochschild, and D. Gata, the layout expertise of D. Franck, the test support of R. Hardin, and the helpful comments of the reviewers.

REFERENCES

[1] J. C. Candy, "A use of limit cycle oscillations to obtain robust analog-to-digital converters," *IEEE Trans. Commun.*, vol. COM-22, pp. 298–305, Mar. 1974.
[2] J. C. Candy, "A use of double integration in sigma–delta modulation," *IEEE Trans. Commun.*, vol. COM-33, pp. 249–258, Mar. 1985.
[3] B. E. Boser and B. A. Wooley, "The design of sigma–delta modulation analog-to-digital converters," *IEEE J. Solid-State Circuits*, vol. 23, no. 6, pp. 1298–1308, Dec. 1988
[4] B. Brandt, D. E. Wingard, and B. A. Wooley, "Second-order sigma–delta modulation for digital audio signal acquisition," *IEEE J. Solid-State Circuits*, vol. 26, no. 4, pp. 618–627, April 1991.
[5] P. Ferguson, Jr. *et al.*, "An 18b 20 kHz Dual $\sum\Delta$ A/D Converter," in *1991 ISSCC Dig. Tech. Papers*, pp. 68–69.
[6] M. Sarhang-Nejad and G. C. Temes, "A high-resolution multi-bit $\sum\Delta$ ADC with digital correction and relaxed amplifier requirements," *IEEE J. Solid-State Circuits*, vol. 28, no. 6, pp. 648–660, June 1993.
[7] R. J. van de Plassche, "Dynamic element matching for high-accuracy monolithic D/A converters," *IEEE J. Solid-State Circuits*, vol. 11, no. 6, pp. 795–800, Dec. 1976.
[8] B. E. Boser, "Design and implementation of oversampled analog-to-digital converters," *Ph.D. dissertation*, Stanford Univ., pp. 100–101, 1988.
[9] L. R. Carley and J. Kennedy, "A 16-bit 4th order noise-shaping D/A converter," in *Proc. 1988 CICC*, pp. 21.7.1–21.7.4.
[10] K. S. Tan, "On board self-calibration of A/D and D/A converters," U.S. Patent 4 399 426, August 16, 1983.
[11] H. S. Lee, D. A. Hodges, and P. R. Gray, "A self-calibrating 15-b CMOS A/D converter," *IEEE J. Solid-State Circuits*, vol. SC-19, no. 6, pp. 813–819, Dec. 1984.
[12] K. S. Tan *et al.*, "Error correction techniques for high-performance differential A/D converters," *IEEE J. Solid-State Circuits*, vol. SC-25, no. 6, pp. 1318–1327, Dec. 1990.
[13] M. de Wit, K.-S. Tan, and R. Hester, "A low power 12-b analog-to-digital converter with on-chip precision trimming," in *1992 Symp. VLSI Circuits Dig. Tech. Papers*, pp. 100–101.
[14] D. Senderowicz *et al.*, "A family of differential NMOS analog circuits for a PCM codec filter chip," *IEEE J. Solid-State Circuits*, vol. SC-17, no. 6, p. 1014–1023, Dec. 1982.
[15] E. Säckinger and W. Guggenbühl, "A high-swing, high-impedance MOS cascode circuit," *IEEE J. Solid-State Circuits*, vol. 25, no. 1, pp. 289–298, Feb. 1990.
[16] K. Bult and G. J. G. M. Geelen, "A fast-settling CMOS opamp for SC circuits with 90-dB DC gain," *IEEE J. Solid-State Circuits*, vol. 25, no. 6, pp. 1379–1384, Dec. 1990.
[17] R. Castello and P. Gray, "A high-performance micropower switched-capacitor filter," *IEEE J. Solid-State Circuits*, vol. SC-20, no. 6, p. 1126, Dec. 1985.
[18] D. E. Knuth, *The Art of Computer Programming*, second ed. Reading, MA: Addison-Wesley, 1981, vol. 2, sec. 3.2.2.
[19] H. Chamberlin, "Computer bits: Random number generators," *Popular Electron.*, vol. 15, no. 3, pp. 89–92, Mar. 1979.
[20] R. M. Gray, "Quantization noise spectra," *IEEE Trans. Inform. Theory*, vol. 36, no. 6, pp. 1220–1244, Nov. 1990.
[21] D. Su *et al.*, "Experimental results and modeling techniques for substrate noise in mixed-signal integrated circuits," *IEEE J. Solid-State Circuits*, vol. 28, no. 4, pp. 420–430, April 1993.

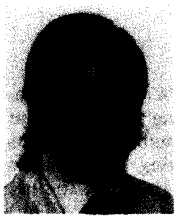

John W. Fattaruso (S'75–M'84–S'86–M'86) was born in Iowa in 1956 and grew up in Berkeley, CA. He received the B.S. (highest honors), M.S., and Ph.D. degrees in electrical engineering from the University of California, Berkeley, through 1986.

He has been a Hertz Foundation Fellow, Teaching Associate, Research Assistant, and Instructor at the University of California, Berkeley. In 1979 he worked in the Digital Signal Processing R&D group at Hewlett-Packard, Santa Clara, CA, and in 1985 he served as a consultant to Seeq Technology, San Jose, CA. He is currently a Member of the Technical Staff of Texas Instruments, Dallas, TX, working in the Semiconductor Process and Design Center on MOS analog VLSI technology. His research interests include analog circuit design, circuit simulation and optimization, neural networks, and numerical analysis.

Dr. Fattaruso is a member of Eta Kappa Nu, Tau Beta Pi, and Phi Beta Kappa. He currently holds five patents in circuit design.

Sami Kiriaki (S'81–M'84) received the B.S. and M.S. degrees in electrical engineering from the University of Rhode Island, Kingston, RI, in 1982 and 1984, respectively.

He has been a Research Assistant at the University of Rhode Island, Kingston, where his main research work concentrated on the development of switched capacitor circuits. Upon graduation he joined the VLSI Design Laboratory of Texas Instruments, Dallas, TX, as a member of the Technical Staff, where he started working on the design and test of high-resolution analog-to-digital data converters. He is currently engaged in the design of low-power low-voltage oversampled ADC's and DAC's. His research interest is the design of analog circuits, especially those for data converters.

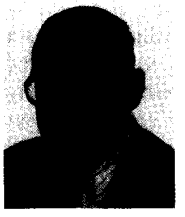

Michiel de Wit (M'85–SM'90) was born in Amsterdam, the Netherlands, in 1933. He received the B.S. degree in physics from Ohio University, Athens, OH, in 1954 and the Ph.D. degree in physics from Yale University in 1960.

In 1959 he joined the Central Research Laboratory of Texas Instruments, Inc., Dallas, TX, where he has been engaged in theoretical and experimental investigations of impurity structures in solids and has developed novel Q-switching techniques for laser target designators. In 1974 he joined the CCD Signal Processing branch, which has evolved to the Analog Design branch, where he is currently engaged in the development of monolithic CMOS A/D and D/A converters. He is currently a Senior Member of the Technical Staff in the Semiconductor Process and Design Center.

Dr. de Wit has coauthored 26 publications and 28 technical presentations and holds four patents.

Greg Warwar (S'87-M'89) received the B.S. and M.S. degrees in electrical engineering from Rice University, Houston, TX, in 1987 and 1989, respectively.

In 1989 he joined the VLSI Design Laboratory at Texas Instruments, Inc., Dallas, TX, as a Member of the Technical Staff. There he worked on the design of oversampling sigma–delta A/D converters and switched-capacitor filters. In 1992 he joined Vitesse Semiconductor in Camarillo, CA, where he is involved in the design of high —speed GaAs circuits for telecommunications applications.