A 1.5-V 14-Bit 100-MS/s Self-Calibrated DAC

Yonghua Cong, *Student Member, IEEE,* and Randall L. Geiger, *Fellow, IEEE*

***Abstract*—Large-area current source arrays are widely used in current-steering digital-to-analog converters (DACs) to statistically maintain a required level of matching accuracy between the current sources. This not only results in large die size but also in significant degradation of dynamic range for high-frequency signals. To overcome technology barriers, relax requirements on the layout, and reduce DAC sensitivities to process, temperature, and aging, calibration is emerging as a viable solution for the next-generation high-performance DACs. In this paper, a new foreground calibration technique suitable for very-low-voltage environments is presented which effectively compensates for current source mismatch, and achieves high linearity with small die size and low power consumption. Settling and dynamic performance are also improved due to a dramatic reduction of parasitic effects. To demonstrate this technique, a 14-bit DAC prototype was implemented in a 0.13-μm digital CMOS process. This is the first CMOS DAC reported that operates with a single 1.5-V power supply and achieves 14-bit linearity with less than 0.1 mm^2 of active area. At 100 MS/s, the spurious free dynamic range is 82 dB (62 dB) for signals of 0.9 MHz (42 MHz) and the power consumption is only 16.7 mW.**

***Index Terms*—Calibration, CMOS, digital-to-analog converter (DAC), high linearity, low voltage, self-calibration.**

I. Introduction

In MANY signal processing and telecommunication applications, the digital-to-analog converter (DAC) is a critical building block limiting the accuracy and speed of the overall system [1]–[4]. When applications require high speed and high resolution, the current-steering DAC architecture is almost exclusively used. It is widely known that inaccurate settled values and nonlinear switching transient both contribute to spectral harmonics in the DAC output [5]. These harmonics are the major factors limiting the spurious free dynamic range (SFDR) of the DAC. The inaccuracy of the settled values is mainly due to current source mismatch and their finite output impedance while the nonlinearity of the switching transients is primarily due to the parasitic effects in the current source cells [6]–[8]. Parasitic effects contribute to slow settling, glitches, clock feedthrough, etc. in the switching transient. Nonlinear on-chip components, clock jitter, and timing skew also contribute to dynamic nonlinearity. With an increase of sampling rate and/or input signal frequency, the SFDR degrades because longer portions of the sampling cycle are occupied by the highly nonlinear switching transients [8], [9].

Manuscript received April 14, 2003; revised July 14, 2003. This work was supported in part by Motorola, the Semiconductor Research Corporation, and the National Science Foundation.

Y. Cong was with the Department of Electrical and Computer Engineering, Iowa State University, Ames, IA 50011 USA. She is now with Broadcom Corporation, Irvine, CA 92619 USA (e-mail: congyh@yahoo.com).

R. L. Geiger is with the Department of Electrical and Computer Engineering, Iowa State University, Ames, IA 50011 USA (e-mail: rlgeiger@iastate.edu).

Digital Object Identifier 10.1109/JSSC.2003.819163

A conventional current source cell is usually implemented with a current source and a pair of switches driven by the digital signals. A DAC contains multiple copies of these current source cells and they are typically judiciously laid out in geometric arrays. Without any trimming or calibration, large-area current sources must be used to overcome the detrimental effects of random mismatch on yield [6], [7], [10]. For a high-resolution DAC, the current source array usually comprises several square millimeters of area. This large area introduces substantial parasitic effects as well as significant gradient effects. To compensate for gradient errors, complex biasing and routing schemes are generally used [6], [7], [11], [12], but these schemes increase the parasitic interconnect capacitances. Large parasitic capacitances severely reduce the settling rate thus causing the SFDR to drop rapidly at high frequencies.

The first goal of this work is to show that with a small-area self-calibration circuit, the area of the current source array can be dramatically reduced compared to what is required for achieving the same yield target without calibration. For a small-area current source array, the gradient effects become less significant, thus relaxing requirements on layout. Calibration also reduces sensitivities to process, temperature and aging. Small current source arrays exhibit reduced parasitic capacitances resulting in an improvement in settling and dynamic performance.

Another goal of this work is to develop a self-calibrated DAC architecture that is suitable for very-low-voltage processes. Paralleling the shrinking of process feature sizes to enhance performance of digital circuits is a comparable decrease in the supply voltage. The reducing supply voltages introduce additional challenges for designers using these state-of-the-art processes to build high-performance DACs. For example, reducing the effective gate-source voltages results in a severe deterioration of the matching and noise immunity of current sources [12]. Cascoding is widely used to enhance the output impedance of current sources and improve DAC dynamic linearity [6], [12], [13], but low supply voltages make it impractical to use cascoded current sources. In these cases, calibration becomes even more attractive.

Since most conventional calibration methods are not suitable for very-low-voltage operation [9], [14], a new foreground digital calibration scheme is used in this design that effectively calibrates the current source mismatches and achieves high linearity and high speed with a very small die size and low power dissipation. The new scheme can be readily implemented in fully digital low-voltage CMOS process leading to considerable cost reductions.

To demonstrate this technique, a 14-bit current-steering DAC prototype was implemented in a 0.13-μm digital CMOS process. This is the first CMOS DAC reported that operates

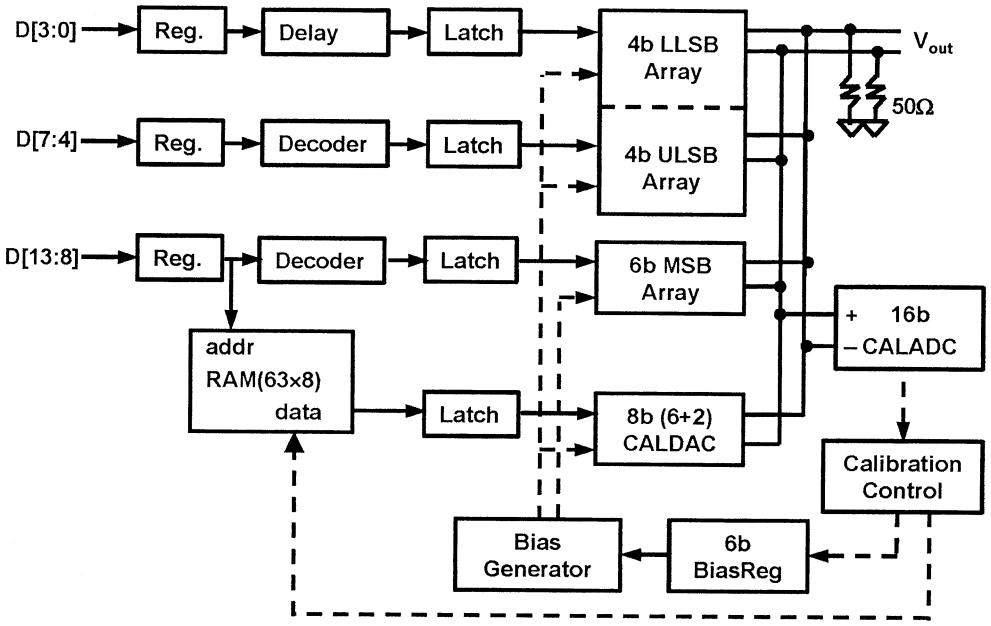

Fig. 1. DAC architecture.

with a single 1.5-V power supply and occupies an active area as small as 0.1 mm^2. It requires only 16.7-mW power at a 100-MHz sampling rate but still maintains state-of-the-art linearity and conversion rates.

The self-calibration scheme is discussed in Section II. The circuit implementation of critical constituent blocks and layout strategies are described in Section III. Measurement results are presented in Section IV.

II. Calibration Schemes

The segmented differential output DAC architecture used in this work is shown in Fig. 1. The 14-bit DAC is segmented into a 6-bit thermometer-decoded MSB array, a 4-bit thermometer decoded upper LSB (ULSB) array and a 4-bit binary-weighted lower LSB (LLSB) array. The MSB array will be self-calibrated. The 8-bit LSB array comprised of the ULSB and LLSB arrays is not calibrated since it needs only to maintain 9-bit accuracy, which is practically achievable with acceptable yield without calibration. The calibration circuitry includes a slow-speed 16-bit calibration analog-to-digital converter (CALADC), a small 8-bit binary-weighted current-steering calibration DAC (CALDAC), a 63-word 8-bit-per-word SRAM, a bias generator, and some calibration control logic. The ADC is used to measure the differential output of the DAC so that the calibration can be implemented in the digital domain. The 63-word SRAM stores the measured errors associated with each of the 63 MSB input codes. During the normal conversion (i.e., after calibration), the MSB digital inputs address the SRAM and the error corresponding to this MSB code is read out and used to drive the CALDAC. The output current from the CALDAC is summed to the output current of the main DAC to provide the overall DAC output current. Since the digital calibration does not need to be done frequently, the CALADC and the control logic can be put into a sleep mode during normal operation, thereby making the power dissipation associated with the calibration process negligible and minimizing effects of any noisy transients associated with the calibration circuitry during normal operation.

A. MSB Array Calibration

The basic idea behind the calibration scheme is illustrated in Fig. 2. Ideally, the 63 current sources in the MSB array are identical and each of them provides an output current that is equal to that provided by the overall LSB array. Here the LSB array is comprised of 255 unit current sources plus one additional dummy unit current source. If the total current of the LSB array is denoted as I_{LSB}, when the LSB array is turned off and the MSB inputs increase from 1 to 63, the DAC output should increase in equal steps according to the sequence $I_{\mathrm{LSB}}, 2I_{\mathrm{LSB}}, \ldots, 63I_{\mathrm{LSB}}$. The difference between these ideal values and the actual DAC outputs, denoted as $e(j)$ $(1 \leq j \leq 63)$, are stored in the SRAM during calibration. During normal operation, these errors are corrected with the CALDAC. The full-scale output of the CALDAC is determined by the maximum possible correctable error of the 63 MSB codes. The nominal inputs of the CALDAC are set to b10 000 000 so that the adjustment can be done in both directions. With the nominal inputs, the differential output of the CALDAC (including a dummy unit current source) is approximately equal to zero, while for single-ended applications, the CALDAC will nominally introduce a dc offset at the DAC output.

A block diagram of the MSB array calibration is shown in Fig. 3. Initially, all inputs of the DAC are set to 0, the inputs to the CALDAC are set to b10 000 000 and the offset of the DAC is measured by the CALADC and stored as D_{off}. In the following, D_{off} will be removed from all digital outputs of the CALADC unless specifically mentioned otherwise. Following measurement of D_{off}, a counter counting from 0 to 63 with unit increments controls the MSB inputs. When the counter is at 0,

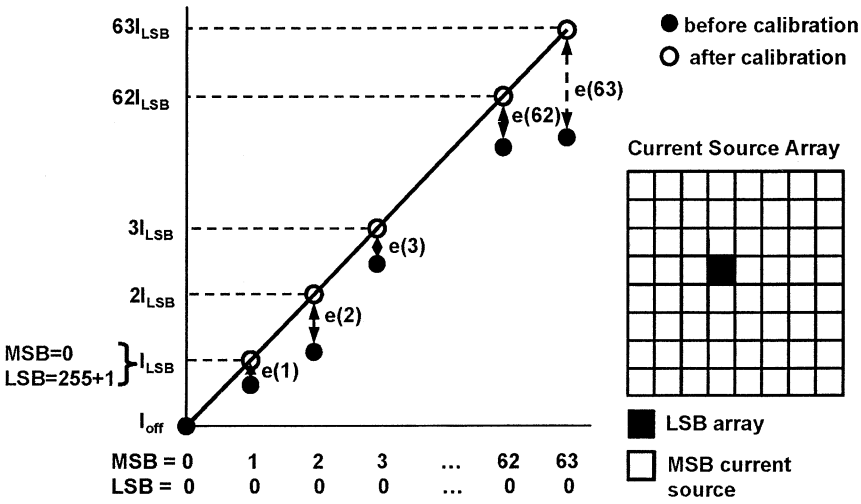

Fig. 2. Conceptual illustration of MSB array calibration.

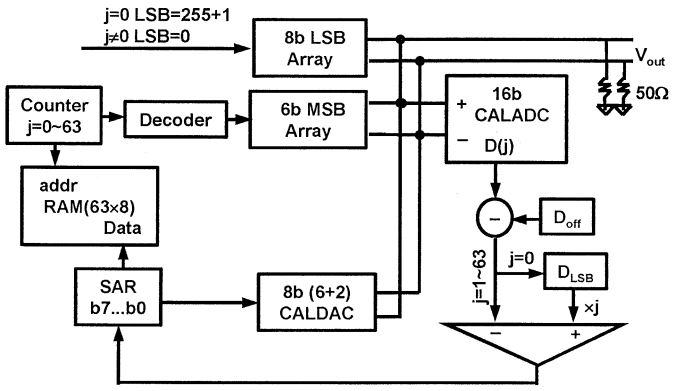

Fig. 3. Block diagram of MSB array calibration.

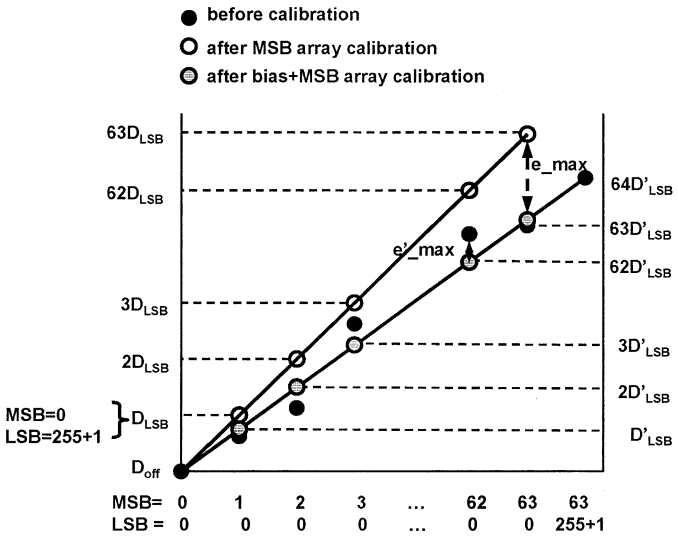

Fig. 4. Conceptual illustration of gain error accumulation and bias calibration.

the MSB inputs are b000 000, and the inputs of the LSB array are set to b11 111 111. The differential output of the DAC, measured by the CALADC and denoted as $D_{\rm LSB}$, is stored in a register. For each subsequent count of the counter, the inputs of the LSB array are all set to 0, and $D_{\rm LSB}$ will be used as the reference to find the error associated with each MSB code. For example, when the counter output is j, $1 \leq j \leq 63$, the MSB input is equal to j, and the desired output of the DAC, when digitized by the CALADC, should be equal to $j \cdot D_{\rm LSB}$. If not, the difference between $j \cdot D_{\rm LSB}$ and the CALADC output, $D(j)$, is defined to be the error associated with the MSB input code j and is stored in word j of the SRAM.

However, in reality, the error corresponding to MSB code j is not simply equal to the digital value $D(j) - D_{\rm off} - j \cdot D_{\rm LSB}$, instead, it is determined through an 8-step successive approximation process utilizing the CALDAC and a successive approximation register (SAR). This is because, to avoid overflow, the input swing of the CALADC should be slightly larger than the output swing of the DAC and the resulting offset and gain mismatch between the ADC and the DAC have to be compensated. The SAR output serves as the input to the CALDAC. In the first step of the successive approximation, the MSB of the SAR is set to 1 while the other bits are set to 0. The corresponding output of the CALDAC together with the output of the main DAC as measured by the CALADC are compared to the desired value $j \cdot D_{\rm LSB}$. If it is smaller, the MSB of the SAR remains at 1, otherwise it is reset to 0. The same process is repeated for each of the lower bits of the SAR and the DAC output becomes closer to the ideal value with the process. After the approximation process is completed, the error code that then resides in the SAR is stored in the SRAM at address location j.

B. Bias Calibration

The above calibration algorithm is subject to gain error accumulation since the gain depends strongly on the measured value of $D_{\rm LSB}$. This can be seen from Fig. 4. Ideally, the total current of the LSB array $D_{\rm LSB}$ is equal to 1/64 of the DAC full scale. However, if $D_{\rm LSB}$ deviates even a small amount from its ideal value, the effect of this deviation will be accumulated during the calibration process as the MSB codes increase, and this can result in a rather large gain error. The gain error will not hurt the linearity of the DAC, but it increases the tuning range requirements for the CALDAC. This problem can be solved by tuning

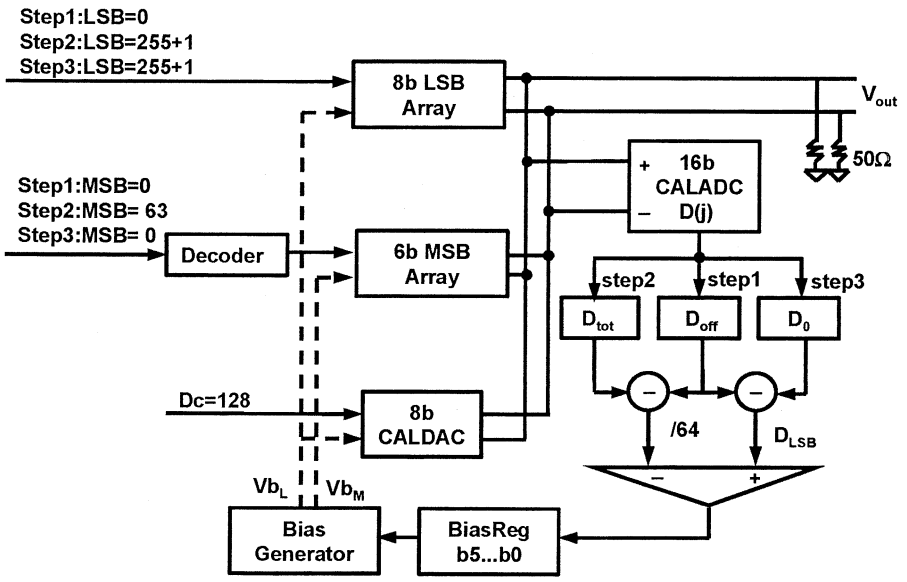

Fig. 5. Block diagram of bias calibration.

the bias voltage of the LSB array so that the current of the LSB array is equal to 1/64 of the DAC full-scale output. Once $D_{\rm LSB}$ has been defined, this value of $D_{\rm LSB}$ is used to calibrate the MSB array using the procedure described in the previous section. Independent of the MSB array calibration, the bias voltage of the LSB array is also calibrated using a successive approximation process utilizing the CALADC and a bias calibration DAC.

Fig. 5 shows the block diagram of the bias calibration. The bias generator provides two bias voltages, $Vb_{\rm M}$ and $Vb_{\rm L}$. $Vb_{\rm M}$ is a fixed bias driving the MSB array, while $Vb_{\rm L}$ drives the LSB array and the CALDAC. The successive approximation calibration process for $Vb_{\rm L}$ uses a 6-bit DAC embedded in the bias generator and a bias register to store the calibration code. The approximation process is similar to that used in the MSB array calibration except that here the target is to make $D_{\rm LSB}$ equal to 1/64 of the DAC full scale output. During this bias calibration, the input of the CALDAC is set to its nominal value.

In summary, the overall calibration procedure is comprised of a LSB array bias calibration and a MSB array linearity calibration. The bias calibration is a six-step successive approximation process. Each step determines one bit of the bias DAC. The MSB array calibration is comprised of 63 cycles, each being used to correct the nonlinearity of one of the 63 MSB codes. Each MSB calibration cycle is comprised of an eight-step successive approximation process and each approximation step determines one bit of the CALDAC.

In the calibration algorithm described, it has been assumed that the CALADC has a voltage input and that the DAC output currents are observed as voltages on the 50-Ω load resistors depicted in Fig. 1. Actually, the algorithm calibrates the static errors of the DAC current outputs. It is independent of the value of the resistive load as long as the load resistance is constant at dc and the ADC reference can be adjusted so that its input swing is slightly larger than the DAC output voltage swing. The algorithm can be applied to either differential output or single-ended output DACs. In applications where the users do not provide dc resistive loads, for example when a transformer is utilized to perform differential to single-ended conversion, or when the output reconstruction filter is of band-pass nature, the calibration algorithm described is not directly applicable. In such cases, the DAC output currents can be switched to on-chip dummy resistive loads for calibration and subsequently directed to the desired load after the calibration is complete.

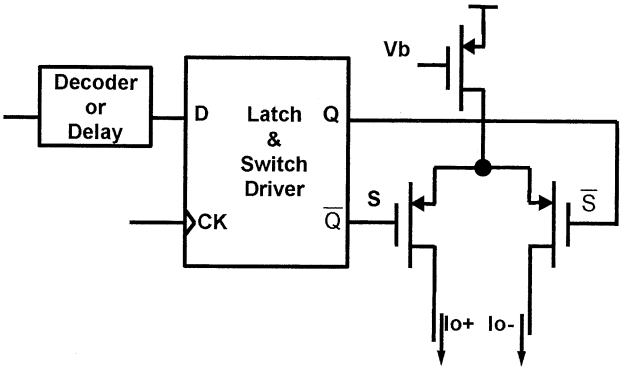

Fig. 6. Current cell.

C. Design Optimization

In any implementation of the proposed architecture, the overall performance is strongly dependent upon the matching accuracy of the current source arrays, the resolution and accuracy of the CALADC, CALDAC, and the bias generator. These issues were addressed with a careful statistical analysis and Monte Carlo simulations. Based on a model given in [15], to achieve integral nonlinearity (INL) and differential nonlinearity (DNL) of less than 0.5 LSB for a 14-bit DAC at the 99% yield level without any trimming or calibration requires a relative standard deviation of the unit current sources, $\sigma_{\rm u}$, of less than 0.22% to overcome the effects of random process variations in addition to a layout that is insensitive to gradient effects. With this $\sigma_{\rm u}$, the minimum gate area for the current source

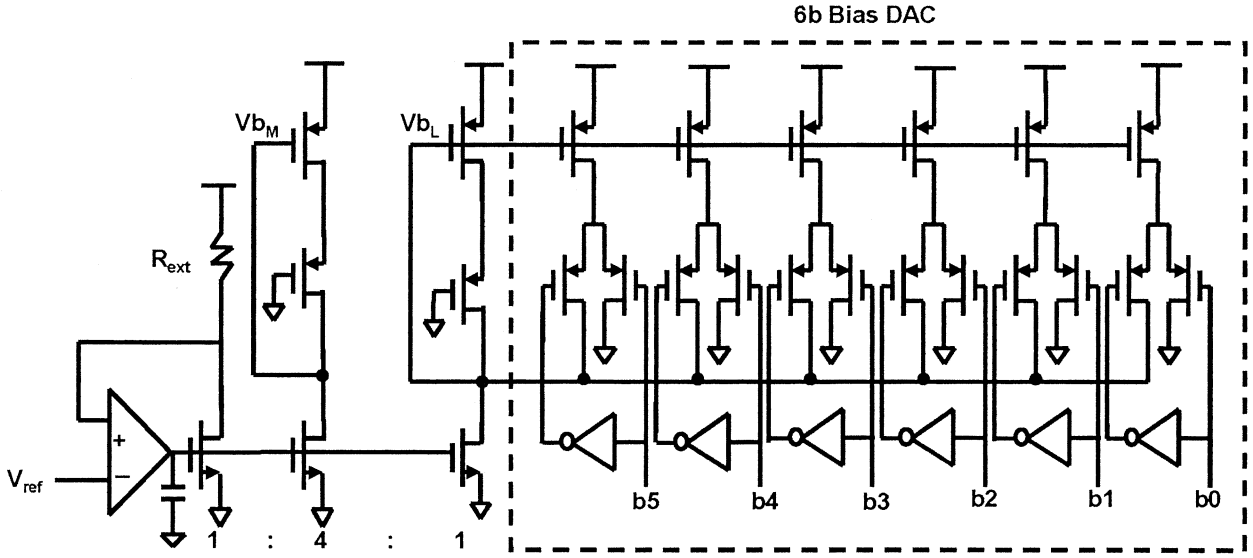

Fig. 7. Bias generator.

transistors can be determined [6], [7], [10]. In the specific 0.13-μm CMOS process used in our implementation, the total gate area $(W \cdot L)$, without calibration, would need to be about 1.5 mm^2 to overcome the effects of random current source mismatch.

After calibration, the overall errors of the DAC include both the residual errors of the calibrated MSB array and the intrinsic errors of the uncalibrated LSB array. The errors of the two arrays are uncorrelated, and in the implementation we allocated half (0.25 LSB) of the overall error budget (0.5 LSB) to each of these arrays. To keep the nonlinearity error of the 8-bit LSB array less than 0.25 LSB, the relative standard deviation of each unit current source, σ_{u}, can be about four times larger than that required for a 14-bit DAC without calibration [15]. The implication of this less stringent standard deviation requirement is a reduction of the gate area of the current sources by about a factor of 16.

Statistical analysis for a 0.25-LSB error budget for the calibrated MSB array suggests that the gate area required for the MSB array can be made as small as that of the LSB array. Thus, compared to the DAC without calibration, the total gate area of the current source array is reduced by over a factor of 500.

One concern associated with dramatically reducing the gate area of current sources in the calibrated structure is that smaller current sources might suffer more from $1/f$ noise. However, although the size of the current sources is dramatically reduced with the calibration, they are still relatively large long-channel devices. Meanwhile, the output noise is not only determined by the gate area and the transconductance of the current sources, but also strongly dependent on the bandwidth of the bias circuits and the ratio between the bias current and the full-scale current. Simulations show that in this design, to steer 10-mA full-scale current, the $1/f$ noise spectral density of the output current is around $160\,\mathrm{nA}/\sqrt{\mathrm{Hz}}$ at 1 Hz and the thermal noise spectral density is $160\,\mathrm{pA}/\sqrt{\mathrm{Hz}}$. Therefore, for a full-scale single-tone sinusoid signal, the noise floor is -147 dB FS/Hz.

After calibration, the residual errors of the MSB array are mainly determined by the errors of the CALADC and the CALDAC. The implementation in the next section was designed so that the errors of the CALADC and the CALDAC were both limited to 0.25 LSB. Therefore, the CALADC performance was set to require at least 16 bits of resolution and accuracy. The CALDAC was also designed to have at least 0.25-LSB resolution but the INL requirement of the CALDAC is less important. As long as it is monotonic, the successive approximation processes can proceed in the right direction. As mentioned before, the tuning range of the CALDAC is determined by the maximum possible error of the 63 MSB codes. Based on a statistical analysis, the CALDAC needs to have a tuning range of $\pm 2^5$ LSB, and hence, 8 bits of resolution.

III. Prototype Implementation

The 14-bit differential output DAC was implemented in a 0.13-μm digital CMOS process with a single 1.5-V power supply. Considering the low power supply voltage, the current source cells were formed with a single-transistor current source and a pair of switches as shown in Fig. 6. The switches are controlled by a compact latch and switch driver that minimizes the fluctuation at the common source node of the switches. The switches operate in the saturation region when they are on thus creating a cascode configuration to enhance the output impedance. When there is enough voltage headroom, small-size cascode devices can be added on either side of the switch pair to further improve the output impedance and bandwidth of the current source [13]. However, in this design we used a single 1.5-V supply to realize a 0.5-V single-ended output voltage swing when driving a 50-Ω load on both the inverting and noninverting outputs. Only 1-V voltage headroom was left for the switch pair and the current source. To keep all the transistors in the saturation region, the effective gate–source voltage $V_{\mathrm{d,sat}}$ of the current source was limited within 400 mV. With

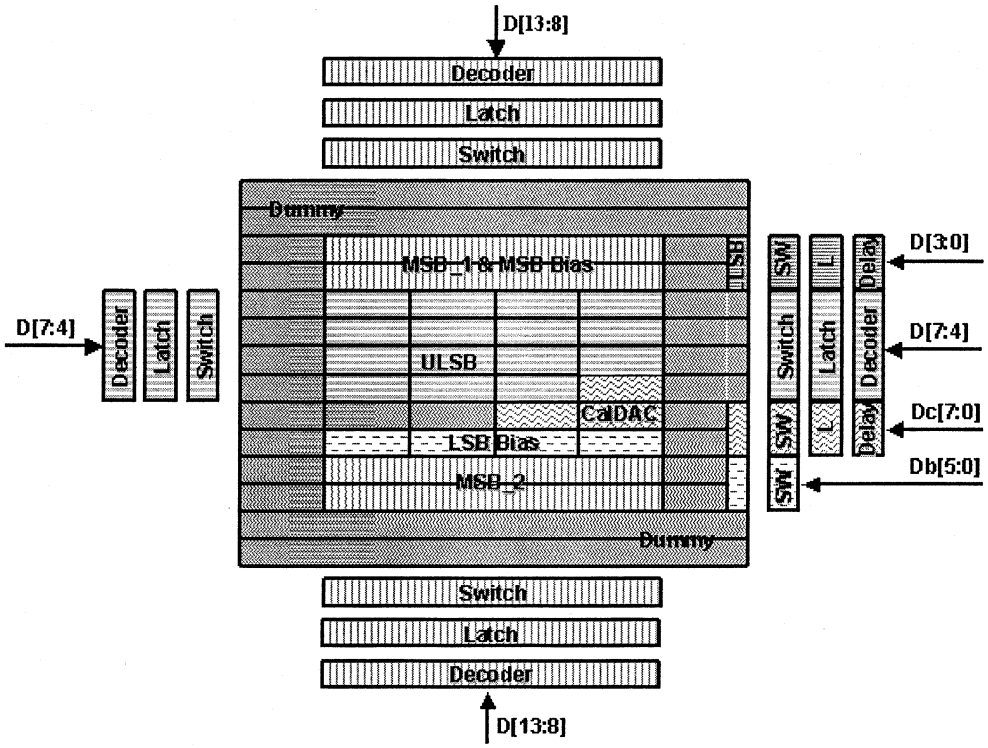

Fig. 8. Floor plan of current cell arrays.

additional cascode transistors, the $V_{\mathrm{d,sat}}$ of the current source has to be further reduced. If the $V_{\mathrm{d,sat}}$ of the current source were to be reduced by a factor of 2, the width of the current source would be increased by a factor of 4 to maintain the same current output and matching. This increases in width would not only increase the active area but also the parasitic diffusion and interconnection capacitances of the current source array. On the other hand, since the size needed for the current sources is dramatically reduced by calibration, the percent area overhead for including additional cascode transistors would be much larger than what would be experienced in a DAC without any trimming or calibration. As a result, while the pole at the drain of the current source was drawn lower due to the increased parasitic capacitance, the extra pole introduced by the cascode stage would also become nontrivial, limiting the improvement of the current source bandwidth. Another drawback of reducing the $V_{\mathrm{d,sat}}$ of the current sources is that it would increase the current source transconductance and the noise floor at the DAC outputs. For all of these reasons, additional cascode devices were not used in this implementation.

Fig. 7 shows the structure of the bias generator. The bias current of the DAC is set by a reference voltage and an external resistor. This bias current is mirrored through two nMOS transistors to drive the diode-connected pMOS transistors that generate the two bias voltages Vb_{M} and Vb_{L}. Vb_{M} is the bias of the MSB array. The pMOS transistors used to generate this bias match the current sources in the MSB array. Vb_{L} is the bias of the LSB array. The pMOS transistors used to generate this bias match the current sources in the LSB array. The six rightmost pMOS transistors are binary weighted and form the 6-bit bias calibration DAC. These transistors can be switched either to ground or to the nMOS current source to adjust the bias voltage

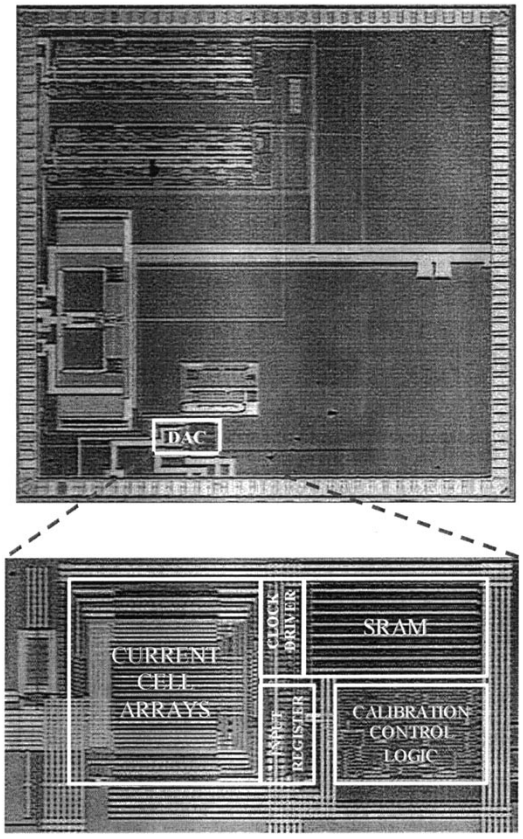

Fig. 9. Chip micrograph.

Vb_{L}. The tuning range of the bias DAC is ± 4 LSB and the minimum resolution is 1/8 LSB. To realize a bidirectional trim of

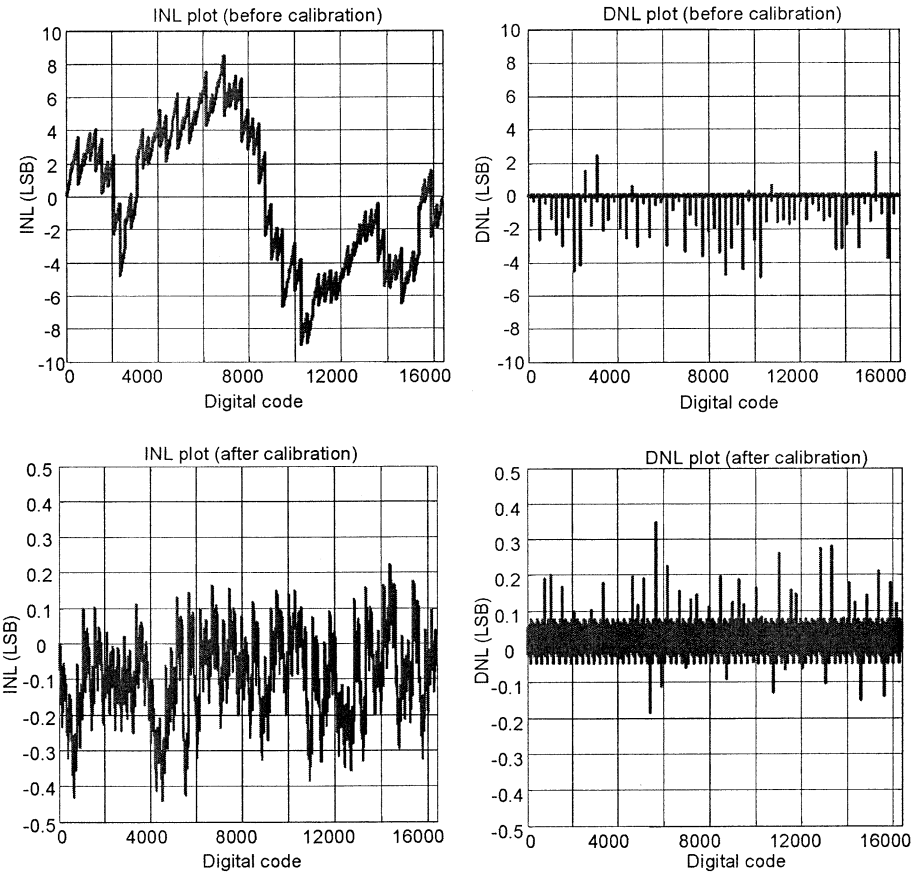

Fig. 10. Static linearity.

the bias voltage, the control bits of the bias calibration DAC are nominally set to b100000.

The floor plan of the current cell arrays including the bias generator is shown in Fig. 8. All current sources are placed at the center being surrounded by the switches, latches and decoders. This arrangement minimizes the interconnections between the current sources and their switches, and hence, the parasitic capacitance at the common source node of the switches. This is essential for improving the settling performance and the dynamic linearity. Since the ULSB array is not calibrated, it is placed in the middle of the current source array to minimize the gradient effects. The MSB array to be calibrated is partitioned into two rows being placed on the top and the bottom of the LSB array, respectively.

Fig. 9 shows the micrograph of the test chip along with an expanded view of the DAC itself. The overall active area of the DAC, exclusive of the CALADC, is only 0.1 mm^2. The remaining area on the chip is occupied by test circuits unrelated to this prototype. The current cell arrays appear on the left side of the DAC die and the SRAM and calibration control logic appear on the right side of the die. In the middle are the clock driver and the input registers.

As the main goal of this prototype is to demonstrate the calibration concept, the CALADC was implemented off-chip using a commercial 16-bit Σ-Δ ADC. Since the ADC converts only dc signals, it can be readily implemented on-chip using a low-order 16-bit Σ-Δ modulator with a high oversampling ratio. We anticipate that in the 0.13-μm process, the area for the ADC would not be larger than that required for the DAC array.

IV. Measurement Results

About 100 chips have been tested. The following results are all from one chip but they are representative of the performance we have seen on all chips. The measured INL and DNL plots are shown in Fig. 10. Before calibration, the INL is 9 LSB and the DNL is 5 LSB. The major errors are due to mismatch of the 63 current sources in the MSB array. After calibration, both the INL and the DNL are belorw 0.5 LSB.

Fig. 11 shows the single-tone spectrum of a full-scale differential output with a 100-MHz sampling rate. After calibration, for a sinusoid signal at 0.9 MHz, the SFDR is 82 dB. When the signal is close to the Nyquist rate, the SFDR is about 62 dB. Fig. 12 shows a two-tone spectrum. For two sinusoid inputs at 23.5 and 24.5 MHz, the SFDR is 66 dB. Fig. 13 shows the SFDR as a function of the normalized input signal frequency at sampling rates of 50 and 100 MHz. The x axis is the ratio between the signal frequency and the sampling frequency. At low signal frequencies, the SFDR is mainly determined by the static linearity. With calibration, it is boosted from 62 to 82 dB. At high signal frequencies, the dynamic nonlinearity is the dominant contributor to the SFDR and the improvement of SFDR with calibration becomes less significant. Even so, for signals near the Nyquist rate, the DAC can still maintain over 60-dB SFDR.

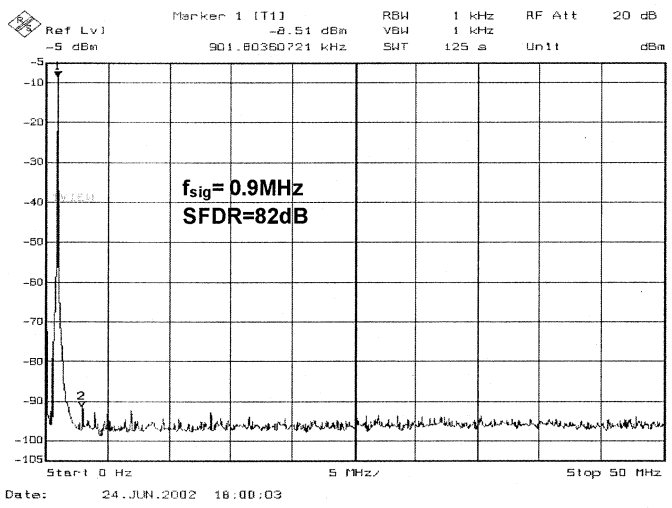

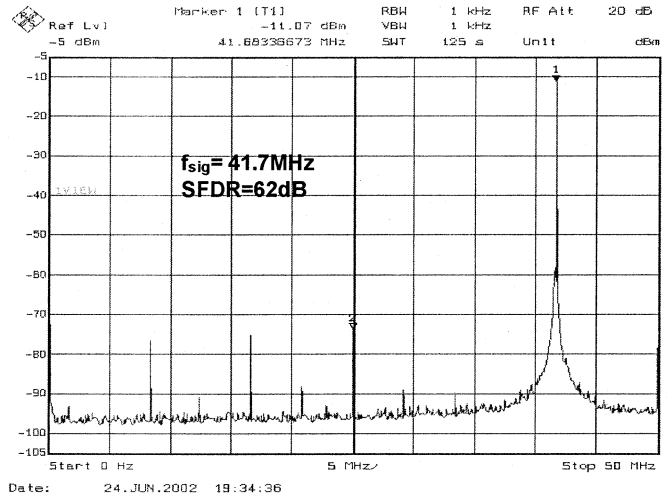

Fig. 11. Single-tone spectrum at 100 MS/s.

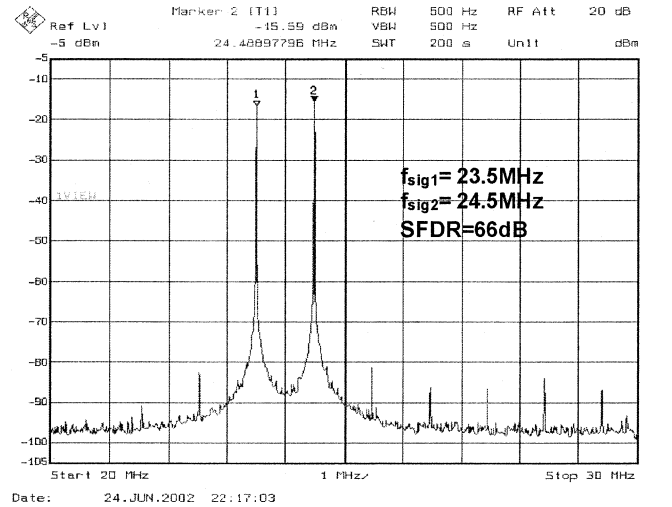

Fig. 12. Two-tone spectrum at 100 MS/s.

It is much higher than what can be attained in a DAC designed under the same conditions but without any trimming or calibration. The maximum sampling rate is 180 MHz. In this prototype, the maximum sampling rate is limited by the access time

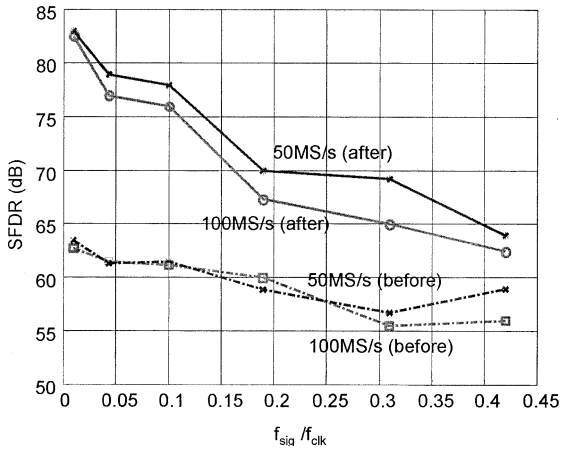

Fig. 13. Dynamic linearity.

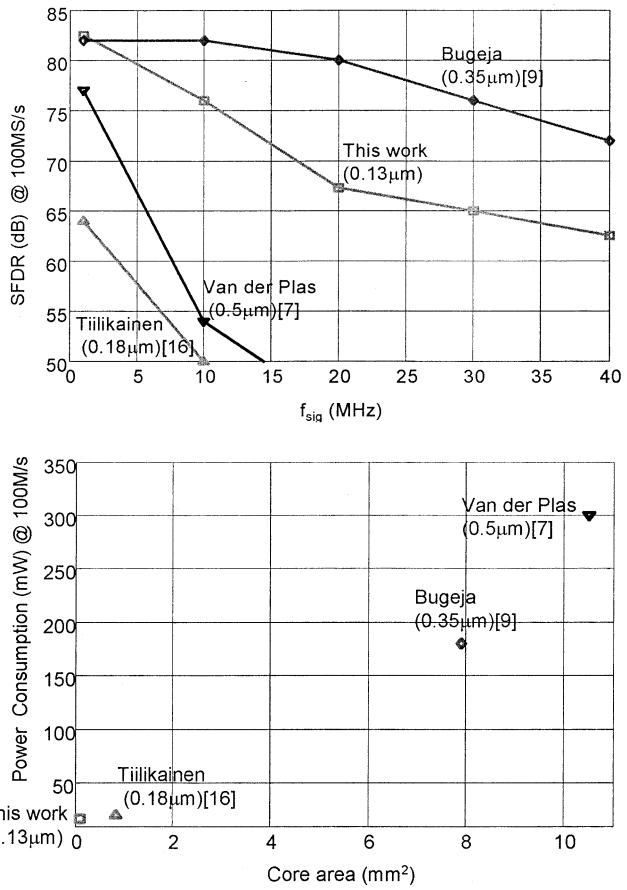

Fig. 14. Comparison with the prior art.

of the SRAM. As shown in Fig. 1, the binary-to-thermometer decoding and the read operation of the SRAM are supposed to be finished in a half clock cycle assuming the clock has 50% duty cycle.

In Fig. 14, the SFDR, area and power consumption of this work are compared with the prior art. The DAC reported in [7] achieves 14-bit static linearity without any trimming or calibration. However, due to the large current source array and complex routing used, the SFDR degrades rapidly at high frequency. With the calibration and the dramatic parasitic reduction, the SFDR of this design drops much slower than the prior art ex-

TABLE I
PERFORMANCE SUMMARY

Resolution	14 bits
Maximum Sampling Rate	180MHz
Full-Scale Current	10mA
Output Swing	0.5V (single-ended) and 1V (differential)
Supply Voltage	1.5V (nominal) 1.25V (minimum)
INL	9LSB (before) 0.43LSB (after)
DNL	5LSB (before) 0.34LSB(after)
Single-tone SFDR @ 50MS/s	64dB(before) 83dB(after) @ f_{sig}= 0.45MHz 59dB(before) 64dB(after) @ f_{sig}=21 MHz
Single-Tone SFDR @ 100MS/s	63dB(before) 82dB (after) @ f_{sig}= 0.9MHz 56dB(before) 62dB(after) @ f_{sig}=42 MHz
Two-Tone SFDR @ 100MS/s	66dB(after) @ f_{out1}=23.5 MHz, f_{out2}=24.5 MHz,
SFDR Temperature Variation	-3dB (-40°C~50°C)
Power Dissipation in 1.5V Supply	15mW(analog)/1.7mW(digital) @ f_{clk}=100MHz, f_{sig}=42MHz
Active Area	0.1mm^2 in 0.13μm CMOS process

cept when compared with the 14-bit DAC reported in [9] where a return-to-zero technique was used at the cost of halving the signal power. Return-to-zero can also be used in this design to further improve the SFDR.

The calibration and optimization used in the proposed structure have provided a dramatic reduction in both power and die area when compared to what is achievable at the same performance level without calibration. The active area of the DAC is only 0.1 mm^2. At 100 MS/s, for a 41.5-MHz output signal the power dissipation is only 16.7 mW for a 1.5-V supply. Of this power, 15 mW is due to the 10-mA full-scale current of the analog portion and only 1.7 mW is consumed by the digital circuitry owing to the small feature size of the process.

The above data were all measured at room temperature. However, it is found that the DACs in this design have low sensitivity to temperature variations. Without redoing the calibration, the SFDR changes by at most 3 dB for temperature over the $-40\,^{\circ}$C–$50\,^{\circ}$C range. The experimental results also show that the DACs can tolerate power supply voltages as low as 1.25 V with less than 3-dB degradation of SFDR. More features are given in Table I.

V. CONCLUSION

A new calibration scheme has been proposed for the design of very-low-voltage high-performance DACs. It has been shown that this approach can be used to achieve high linearity along with a dramatic reduction in die area and power dissipation when compared with conventional designs that maintain comparable yield. The settling rate and dynamic linearity also show significant improvements due to the reduction of parasitic capacitances associated with a much smaller die area. Experimental results show that in a 0.13-μm digital CMOS process and with a single 1.5-V supply, 14-bit resolution and accuracy can be achieved with only 0.1 mm^2 of active area (not including a 16-bit low-order Σ-Δ modulator) and 16.7 mW of power consumption at 100 MS/s. The SFDR at 100 MS/s is 82 dB in the lower signal frequencies and over 60 dB for signal frequencies up to the Nyquist rate.

ACKNOWLEDGMENT

The authors would like to thank D. Garrity and B. Newman, of the Semiconductor Product Sector of Motorola, for their constructive input on this project.

REFERENCES

[1] J. Vankka, J. Ketola, O. Vaananen, J. Sommarek, M. Kosunen, and K. Halonen, "A GSM/EDGE/WCDMA modulator with on-chip D/A converter for base station," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, 2002, pp. 236–237.

[2] I. Mehr, M. Prabir, and D. Paterson, "A 12-bit integrated analog front end for broadband wireline networks," *IEEE J. Solid-State Circuits*, vol. 37, pp. 302–309, May 2001.

[3] W. Claes, W. Sansen, and A. Puers, "A 40-μ A/channel compensated 18-channel strain gauge measurement system for stress monitoring in dental implants," *IEEE J. Solid-State Circuits*, vol. 37, pp. 293–301, Mar. 2002.

[4] C. Yang, V. Stojanovic, S. Modjtahedi, M. Horowitz, and W. Ellersick, "A serial-link transceiver based on 8-G samples/s A/D and D/A converters in 0.25-μm CMOS," *IEEE J. Solid-State Circuits*, vol. 36, pp. 293–301, Nov. 2001.
[5] J. Wikner and N. Tan, "Modeling of CMOS digital-to-analog converters for telecommunication," *IEEE Trans. Circuits Syst. II*, vol. 46, pp. 489–499, May 1999.
[6] J. Bastos, A. Marques, M. Steyaert, and W. Sansen, "A 12-bit intrinsic accuracy high-speed CMOS DAC," *IEEE J. Solid-State Circuits*, vol. 33, pp. 1959–1969, Dec. 1998.
[7] G. Van der Plas, J. Vandenbussche, W. Sansen, M. Steyaert, and G. Gielen, "A 14-bit intrinsic accuracy Q^2 random walk CMOS DAC," *IEEE J. Solid-State Circuits*, vol. 34, pp. 1708–11 718, Dec. 1999.
[8] A. Rugeja and B. Song, "A 14-b, 100-MS/s CMOS DAC designed for spectral performance," *IEEE J. Solid-State Circuits*, vol. 34, pp. 1719–1731, Dec. 1999.
[9] ——, "A self-trimming 14-b 100-MS/s CMOS DAC," *IEEE J. Solid-State Circuits*, vol. 35, pp. 1841–1852, Dec. 2000.
[10] J. Bastos, M. Steyaert, and W. Sansen, "A high-yield 12-bit 250-MS/s CMOS D/A converter," in *Proc. IEEE Custom Integrated Circuits Conf.*, May 1996, pp. 431–434.
[11] Y. Cong and R. Geiger, "Switching sequence optimization for gradient error compensation in thermometer-decoded DAC arrays," *IEEE Trans. Circuits Syst. II*, vol. 47, pp. 585–595, July 2000.
[12] T. Miki, Y. Nakamura, M. Nakaya, S. Asai, Y. Akasaka, and Y. Horiba, "An 80-MHz 8-bit CMOS D/A converter," *IEEE, J. Solid-State Circuits*, vol. SC-21, pp. 983–988, Dec. 1986.
[13] A. Van den Bosch, M. Steyaert, and W. Sansen, "SFDR-bandwidth limitations for high-speed high-resolution current-steering CMOS D/A converters," in *Proc. IEEE Int. Conf. Electronics, Circuits and Systems*, 1999, pp. 1193–1196.
[14] D. Groeneveld, H. Schouwenaars, H. Termeer, and C. Bastiaansen, "A self-calibration technique for monolithic high-resolution D/A converters," *IEEE J. Solid-State Circuits*, vol. 24, pp. 1517–1522, Dec. 1989.
[15] Y. Cong and R. Geiger, "Formulation of INL and DNL yield estimation in current-steering D/A converters," in *Proc. IEEE Int. Symp. Circuits and Systems*, 2002, pp. 149–152.
[16] M. Tiilikainen, "A 14-bit 1.8-V 20-mW 1-mm^2 CMOS DAC," *IEEE J. Solid-State Circuits*, vol. 36, pp. 1144–1147, July 2001.

Yonghua Cong (S'00) received the B.S. degree in computer communications and the M.S. degree in communication and electronic systems from the University of Electronic Science and Technology of China in 1993 and 1996, respectively, and the Ph.D. degree in computer engineering from Iowa State University, Ames, in 2002.

From 2001 to 2002, she was with the Semiconductor Product Sector, Motorola, as an intern working on high-performance DACs. Since 2002, she has been a Staff Scientist with the Analog and RF Microelectronics Group, Broadcom Corporation, Irvine, CA. Her current research interests include high-speed and high-resolution data converters and low-voltage mixed-signal IC design.

Dr. Cong received the Research Excellence Award from Iowa State University in 2002.

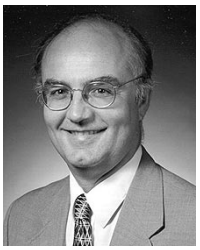

Randall L. Geiger (S'75–M'77–SM'82–F'90) received the B.S. degree in electrical engineering and the M.S. degree in mathematics from the University of Nebraska, Lincoln, in 1972 and 1973, respectively, and the Ph.D. degree in electrical engineering from Colorado State University, Fort Collins, in 1977.

He served as a Faculty Member with the Electrical Engineering Department, Texas A&M University, College Station, from 1977 to 1990. Since 1991, he has been a Member of the Faculty with the Electrical and Computer Engineering Department, Iowa State University, Ames, where he is currently the Willard and Leitha Richardson Professor. His teaching and research interests are in the fields of analog and mixed-signal VLSI design, specifically in the areas of amplifier design, test and built-in self test of mixed-signal circuits, data converter design, device modeling, and design for yield.

Dr. Geiger is a past member of the Board of Governors, a past Vice President of Publications, and a past President of the IEEE Circuits and Systems Society. He has served as an Associate Editor of the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS PART II and as the Circuits and Systems Society editor for the *IEEE Circuits and Devices Magazine*. He was a member of the IEEE Publications Board and the IEEE Periodicals Council and is a past chair of the Transactions Committee of the Periodicals Council. He has served in various capacities on the technical program committees and on the organizing committees for the IEEE International Symposium on Circuits and Systems and the IEEE Midwest Symposium on Circuits and Systems. He received an IEEE Fellow Award in 1990, the Meritorious Service Award of the IEEE Circuits and Systems Society in 1996, the Golden Jubilee Medal of the IEEE Circuits and Systems Society in 2000, and the IEEE Millennium Medal in 2000.