A 300-MS/s 14-bit Digital-to-Analog Converter in Logic CMOS

John Hyde, *Member, IEEE*, Todd Humes, *Member, IEEE*, Chris Diorio, *Member, IEEE*, Mike Thomas, and Miguel Figueroa

***Abstract*—We describe a floating-gate trimmed 14-bit 300-MS/s current-steered digital-to-analog converter (DAC) fabricated in 0.25- and 0.18-μm CMOS logic processes. We trim the static integral nonlinearity to $\pm$0.3 least significant bits using analog charge stored on floating-gate pFETs. The DAC occupies 0.44 mm^2 of die area, consumes 53 mW at 250 MHz, allows on-chip electrical trimming, and achieves better than 72-dB spur-free dynamic range at 250 MS/s.**

***Index Terms*—Digital-to-analog conversion, floating gate.**

I. Introduction

EMERGING standards for communications systems require digital-to-analog converters (DACs) with sample rates in the hundreds of megasamples per second and resolutions of 10–14 bits [1]. Designers typically use current-steering DACs for these applications because they are fast and can drive output loads without buffering. However, the static linearity of a current-steering DAC is sensitive to current-source mismatch. Designers often use large devices, randomized layouts, laser trimming, or continuous on-line electrical trimming [2]–[4] to reduce this mismatch. These techniques improve linearity, but at the expense of die area, power dissipation, or dynamic performance.

Analog-valued floating-gate MOSFETs are near-ideal current sources for a DAC, because they allow post-fabrication electrical trimming of their output current, and because they store a trim value almost indefinitely. They also allow small current-source transistors, because trimming removes matching constraints from the design equation. We have previously described floating-gate pFETs, fabricated in standard CMOS logic processes, that store analog charge on a floating gate with 16-bit resolution [5]. We have also described how to use these devices to trim a DAC current-source array [6], although the DAC described in [6] was capable of only static outputs (i.e., no dynamic performance). In this paper, we describe an entirely new 14-bit DAC with $\pm$0.3 least significant bit (LSB) integral nonlinearity (INL), which is an order of magnitude improved over the DAC in [6], and with dynamic performance that benefits from using small transistors. We also demonstrate how floating-gate trimming allowed us to port the DAC from a 0.25-μm process to a 0.18-μm process without redesign (other than scaling contacts and vias) and with improved performance. We show results primarily from the 0.25-μm DAC; we describe the 0.18-μm DAC.

Manuscript received August 9, 2002.; revised December 20, 2002. This work was supported by Impinj, Inc.

J. Hyde, T. Humes, C. Diorio, and M. Thomas are with Impinj, Inc., Seattle, WA 98103 USA (e-mail: todd.humes@impinj.com).

M. Figueroa is with the Department of Computer Science and Engineering, University of Washington, Seattle, WA 98195 USA.

Digital Object Identifier 10.1109/JSSC.2003.810049

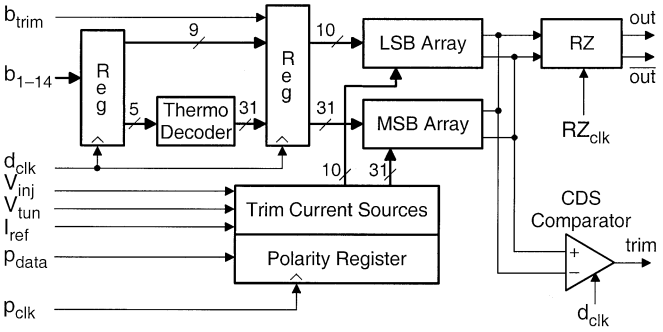

Fig. 1. DAC block diagram. The DAC comprises a 5-bit thermometer-decoded MSB section and a 9-bit binary LSB section. The return-to-zero switch reduces output glitch energy. The trim current sources adjust each bit over a $\pm$5 LSB range. The extra LSB $b_{\rm trim}$ (also trimmable) and the correlated-double-sampling comparator are used only during the trim process.

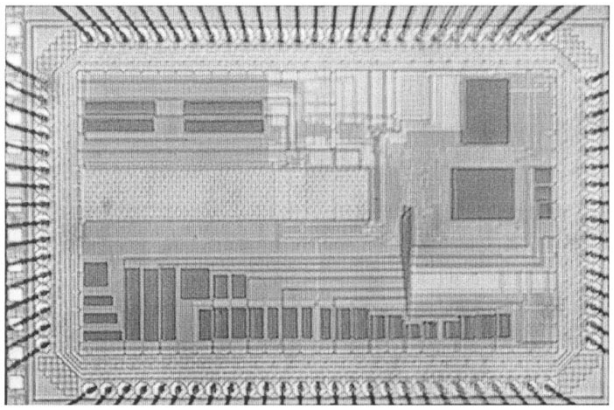

Fig. 2. DAC micrograph. The total DAC core area is 0.44 mm^2 for both the 0.25- and 0.18-μm implementations; the trim circuits account for 14% of this area. We used 0.25-μm five-metal and 0.18-μm six-metal single-poly CMOS logic processes.

II. DAC Architecture

Fig. 1 shows the DAC architecture, and Fig. 2 a die plot. The DAC is segmented as five thermometer-decoded most significant bits (MSBs), nine binary-decoded LSBs, and an additional LSB for trimming. The digital circuitry comprises a 14-bit input data register, a 5-to-31 thermometer decoder to set the MSB current switches, and a 41-bit register to drive the differential-pair switches for the MSB, LSB, and trim-LSB sections. The 41-bit

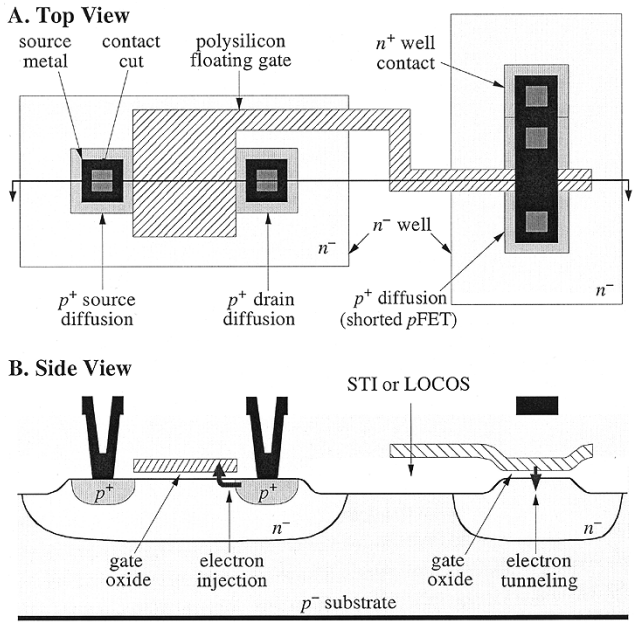

Fig. 3. Floating-gate pFET and its associated tunneling junction, showing the electron tunneling and injection locations. Both pFETs share a common floating gate.

register uses an internally regulated low-voltage supply to minimize its voltage swings (and thereby glitch energy) during differential-pair switching.

The thermometer and binary current sources are arranged in a single current-source array. Each of the 41 current sources comprises a static (untrimmable) source and an associated floating-gate trimmable source. Each trimmable source can trim the output current over a ± 5 LSB range. A current mirror in each trim cell allows us to either add the trim current to, or subtract the trim current from, the associated static source, providing a bidirectional trim capability. A digital polarity bit holds the trim state (add or subtract) for each source. The return-to-zero (RZ) switch at the array output shorts the differential output wires together during codeword switching, further reducing output glitch energy. We use the correlated-double-sampling (CDS) comparator for trimming.

III. Electrically Trimmable Current Source

The heart of our current-source trim is a floating-gate pFET. Fig. 3 shows the structure, comprising two p-channel MOSFETs with a shared floating gate. We use the first transistor for hot-electron injection and a second, with shorted drain, source, and well, for electron tunneling. There is no direct electrical connection to the floating gate nor is there a second polysilicon coupling capacitor; consequently, the structure in Fig. 3 is compatible with CMOS logic processes.

Charge stored on the floating gate determines the gate potential, and, consequently, the channel current of a current-source pFET that shares this gate. We use impact-ionized hot-electron injection (IHEI) [7] to add electrons to the floating gate, and Fowler–Nordheim (FN) tunneling [8] to remove them. We apply ~ -4.8 V drain-to-source across the injection transistor to cause IHEI at the drain, and ~ 10 V to the shorted transistor's drain, source, and well to tunnel electrons off the floating gate. The

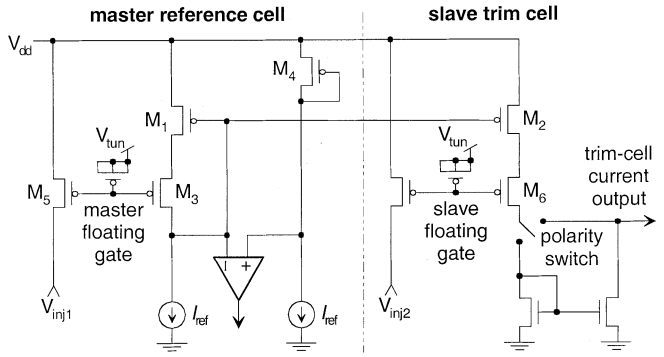

Fig. 4. Trim cell. A single master reference controls 41 slave cells. The master cell adjusts M_2's gate voltage to ensure that floating-gate transistor M_6 provides a constant output current despite temperature variations in transconductance and mobility. We trim the floating-gate voltages using tunneling and injection (see Fig. 3).

tunneling transistor's n^- well acts as a high-voltage implant, allowing us to apply 10 V (causing FN tunneling through the shorted transistor's gate oxide) without incurring p-n junction breakdown. IHEI and tunneling enable bidirectional updates to the floating-gate charge, and, consequently, bidirectional updates to the current-source output. IHEI is a very precise and controllable process, allowing us to write accurate charge values to the floating gate.

To ensure adequate charge retention, all of the floating-gate pFETs are 3.3-V devices with 70-Å gate oxides. Accelerated leakage experiments on the floating-gate trim cell shows that leakage-induced changes in the output current will not cause the DAC's INL to exceed one LSB after ten years.

IV. Trimming the DAC

To trim the DAC, we must adjust the current in the 41 trim current sources. The trim circuitry comprises a single master and 41 slave cells. We begin by describing the trim cells themselves, then the trimming algorithm.

Although a floating-gate pFET's gate charge is nonvolatile, its channel current still varies with temperature, so we must compensate the trim current sources for temperature-induced variations in carrier mobility and threshold voltage V_t. Fig. 4 shows the master cell and a single slave. We begin with the master cell. We trim the gate charge on pFET M_3 in the master reference cell until the comparator toggles, indicating that M_1 is biased in its triode regime. If subsequent temperature changes increase M_3's transconductance, M_3's drain voltage will rise, pushing back on M_1's gate and increasing M_1's triode resistance to ensure constant channel current.

We adjust the charge on slave-cell pFET M_6's gate to obtain the desired trim-cell output current. M_1 and M_2 share a gate, so changing M_1's gate voltage to compensate M_3's drain current for temperature also compensates M_6's drain current for temperature. Ideally, when the trim current is equal to the reference current, we achieve perfect temperature compensation. A polarity switch determines whether the trim cell adds or subtracts its output from the corresponding bit's static output current.

We show the trim algorithm in Fig. 5. The algorithm comprises two main steps. First, we trim all thermometer sources to a median value. Trimming to the median minimizes the trim

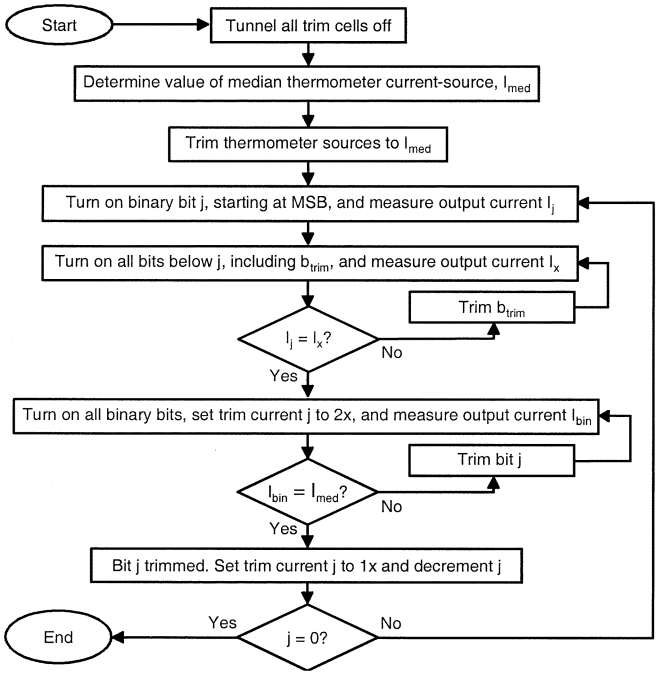

Fig. 5. Trim procedure. We first turn off all the current sources, then successively turn on and trim individual sources, proceeding from MSB to LSB. The algorithm needs only an on-chip CDS comparator (i.e., no off-chip instruments).

range for all the sources. Second, we trim each bit in the binary section using a thermometer source as a reference. The trimming is top-down, meaning that we start with the MSB and end with the LSB. Top-down trim is superior to bottom-up trim, both because it uses the largest current source as a reference, and because it minimizes error propagation by halving the trim range at each successive bit. To trim bit j, we add the currents from sources $(j-1)$ to the LSB plus an extra LSB $b_{\rm trim}$, and follow the procedure in Fig. 5 to match this sum to the current in source j. Because we trim bit j's current to half of bit $(j+1)$'s current, we need to double bit j's current during the trim. We include a gain-of-two current mirror ($2\times$ mode) in the trim cell, although we have omitted it from Fig. 4 for clarity. The algorithm requires a single CDS comparator to trim all the current sources. Although the present DAC uses an off-chip state machine to control the trim, we will integrate this state machine on-chip in future designs.

Tunneling cells individually would require high-voltage switches. Because our n-well-based high-voltage switches are large, we avoid them in the present design and simply tunnel all 41 trim cells simultaneously, by applying 10 V to the 41 well-tunneling junctions. Our present design uses an off-chip high-voltage source; future designs will incorporate a 10-V charge pump on-chip.

After tunneling, we trim the cells using injection. Electron injection requires that we apply roughly -1.5 V to the injection pFET drains. We generate -1.5 V using individual single-stage charge pumps for each of the 41 trim cells. Injection efficiency depends on the magnitude of the drain-to-channel electric field; higher fields generate more hot electrons by impact ionization, resulting in larger gate currents [9]. The drain-to-channel electric field is maximal for subthreshold operation and decreases

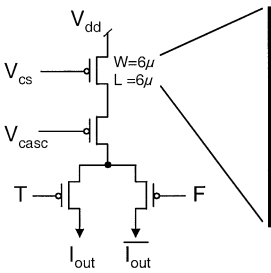

Bit	Current (μA)		PMOS Channels		
	Total	Each	Series	Parallel	Total
1-5 (31)	312.50	9.77	1	32	32
6	156.25	9.77	1	16	16
7	78.13	9.77	1	8	8
8	39.06	9.77	1	4	4
9	19.53	4.88	2	4	8
10	9.77	2.44	4	4	16
11	4.88	1.22	8	4	32
12	2.44	1.22	8	2	16
13	1.22	0.61	16	2	32
14	0.61	0.61	16	1	16
X	0.61	0.61	16	1	16

Fig. 6. Current-source array. The current sources comprise combinations of series and parallel-connected 6 μm$\times$ 6 μm pFETs with cascode transistors. The table shows the transistor count and connectivity for each of the current sources in the array. For example, bit 10 comprises four parallel combinations of four series-connected units (6 μm$\times$ 6 μm) pFETs for a total of 16 pFETs.

			R	T31	T30	T29	B6	T27	T26	T25	T24		T22	T21	T20	T19	BX	T17	T23	T18	T28
			T5	T15	T10	T16	B6	T14	T13	T12	T11	B1	T9	T8	T7	T6	BX	T4	T3	T2	T1
		V	T1	T2	T3	T4	B6	T6	T7	T8	T9		T11	T12	T13	T14	BX	T16	T10	T15	T5
			T28	T18	T23	T17	B6	T19	T20	T21	T22	B2	T24	T25	T26	T27	BX	T29	T30	T31	R
		C	T28	T18	T23	T17	B6	T19	T20	T21	T22		T24	T25	T26	T27	BX	T29	T30	T31	R
			T1	T2	T3	T4	B6	T6	T7	T8	T9		T11	T12	T13	T14	BX	T16	T10	T15	T5
		V	T5	T15	T10	T16	B6	T14	T13	T12	T11	B1	T9	T8	T7	T6	BX	T4	T3	T2	T1
			R	T31	T30	T29	B6	T27	T26	T25	T24	B4	T22	T21	T20	T19	BX	T17	T23	T18	T28
		P,T	T28	T18	T23	T17	B5	T19	T20	T21	T22	B4	T24	T25	T26	T27	B8	T29	T30	T31	R
			T1	T2	T3	T4	B5	T6	T7	T8	T9	B3	T11	T12	T13	T14	B8	T16	T10	T15	T5
		C	T5	T15	T10	T16	B5	T14	T13	T12	T11	B1	T9	T8	T7	T6	B8	T4	T3	T2	T1
			R	T31	T30	T29	B5	T27	T26	T25	T24		T22	T21	T20	T19	B8	T17	T23	T18	T28
		C	R	T31	T30	T29	B7	T27	T26	T25	T24	B2	T22	T21	T20	T19	B8	T17	T23	T18	T28
			T5	T15	T10	T16	B7	T14	T13	T12	T11	B1	T9	T8	T7	T6	B8	T4	T3	T2	T1
		V	T1	T2	T3	T4	B7	T6	T7	T8	T9		T11	T12	T13	T14	B8	T16	T10	T15	T5
			T28	T18	T23	T17	B7	T19	T20	T21	T22		T24	T25	T26	T27	B8	T29	T30	T31	R

Fig. 7. Current-source array layout. We optimized the current source array to compensate for linear and quadratic gradient errors. The figure depicts the location of the individual unit transistors (6 μm$\times$ 6 μm size) and their location within the array. We show only one of the four quadrants in the array for clarity. The symmetry axes are the bold lines on the right and bottom of the array.

as the MOSFET transitions above threshold. Consequently, we operate our floating-gate pFETs with channel currents near threshold, thereby allowing us to use small charge-pump capacitors. Therefore, 41 individual charge pumps are smaller than one large charge pump and a decoder.

V. Current-Source Array

Fig. 6 shows the configuration of the current-source array and output differential switch. We use standard cascoded current sources with differential output switches. We did not optimize the differential switches for harmonic performance because the RZ switch at the DAC output ensures low output spurs. The RZ switch shorts the outputs together during the first half of the sampling clock cycle, and applies the differential currents to the output during the second half. We implement the current sources as series and parallel combinations of 6 μm$\times$ 6 μm unit transistors (see the table in Fig. 6) similar to Van der Plas *et al.* [4]. For example, we designed the 312.5-μA thermometer current sources as 32 parallel unit transistors, and the LSB as of 16 series-connected unit transistors.

We chose to limit the trim range to ± 5 LSBs to minimize trim sensitivity. We optimized the current-source array to ensure ± 5 LSB intrinsic matching using standard matching techniques [10]. Fig. 7 shows one quadrant of the current-source

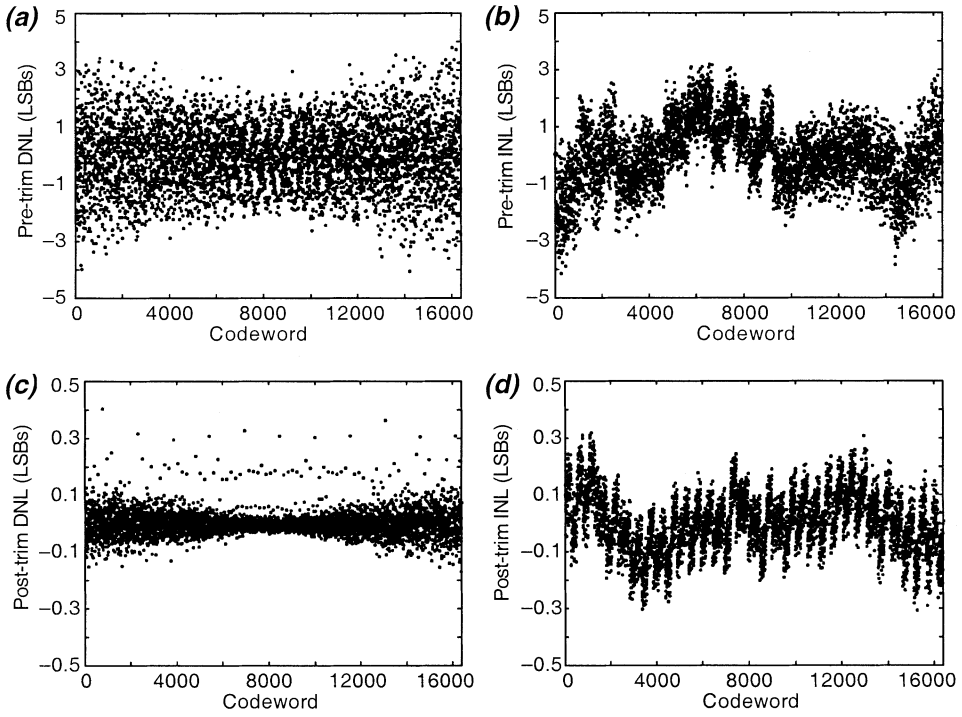

Fig. 8. Static performance. (a), (b) The DAC's pre-trim DNL and INL. (c), (d) The post-trim DNL and INL. Trimming improves both DNL and INL by roughly one order of magnitude.

array layout. We indicate the axes of symmetry by the two bold lines on the right side and bottom of the array. A T prefix denote the thermometer bits; a B prefix denotes binary segments. The blank cells denote dummy transistors whose sources and drains are together tied to the $V_{\rm dd}$ supply. The array includes additional current references that occupy areas normally reserved for dummy transistors. These include the C current sources used to bias the cascode transistors, R current sources for the reference diode, V current sources to bias the internal voltage regulator, and the *PT* current sources for biasing the polarity and trim circuitry, respectively.

We designed the array to compensate for both linear and quadratic gradient errors. We modeled and simulated the matching errors using BSIM3v3.2 transistor models provided by the foundry; these simulations predicted that the series-connected unit transistors would have systematic matching errors up to 9.9%. However, theory as well as EKV models predicted no such systematic matching errors. As can be seen by the pre-trim INL data in Fig. 8, we achieved an intrinsic linearity of ± 4 LSBs, validating the theoretical analysis and the EKV models. We compensated the cascode array for linear but not for quadratic gradient errors, primarily because the source and drain lines needed routing channels which limited our ability to compensate quadratic errors in the floorplan.

VI. Performance

Fig. 8 shows measured differential nonlinearity (DNL) and INL, both before and after trimming. Trimming improves the INL to less than ± 0.3 LSB. If we had designed for 0.3 LSB INL using intrinsic matching rather than trimming, the area of the current-source array would have increased by two orders of magnitude. Similarly, if we had used continuous (capacitor-based) electrical trimming rather than a floating-gate trim, we would have needed to trim continuously, increasing die area and degrading dynamic performance.

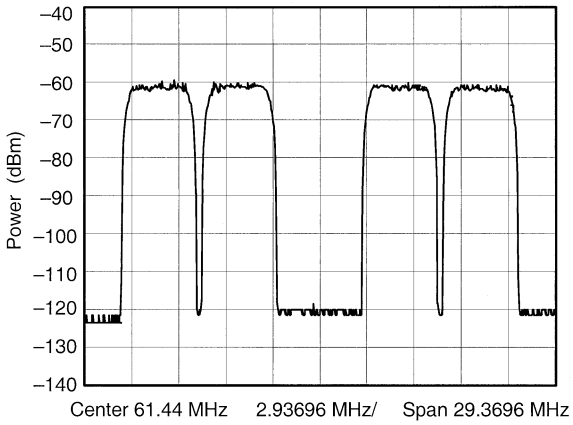

Fig. 9. ACPR performance on a four-channel WCDMA waveform. The ACPR, after subtracting the spectrum analyzer noise, measured 62 dB, exceeding the 3GPP WCDMA requirement of 45 dB for an immediately adjacent channel or 50 dB for other channels.

To demonstrate DAC performance, we measured the adjacent channel power ratio (ACPR) on a four-carrier wide-band code division multiple access (WCDMA) waveform. ACPR measures the power in a single carrier relative to the power in an adjacent unused carrier slot, and stresses a DAC's static and dynamic performance. INL and DNL affect the DAC's ACPR by raising the noise floor in adjacent channels. We synthesized four carriers, centered at 61.44 MHz with 5-MHz spacing, clocking the DAC at 245.76 MS/s. Spectrum-analyzer noise dominates the noise floor shown in Fig. 9. After correcting

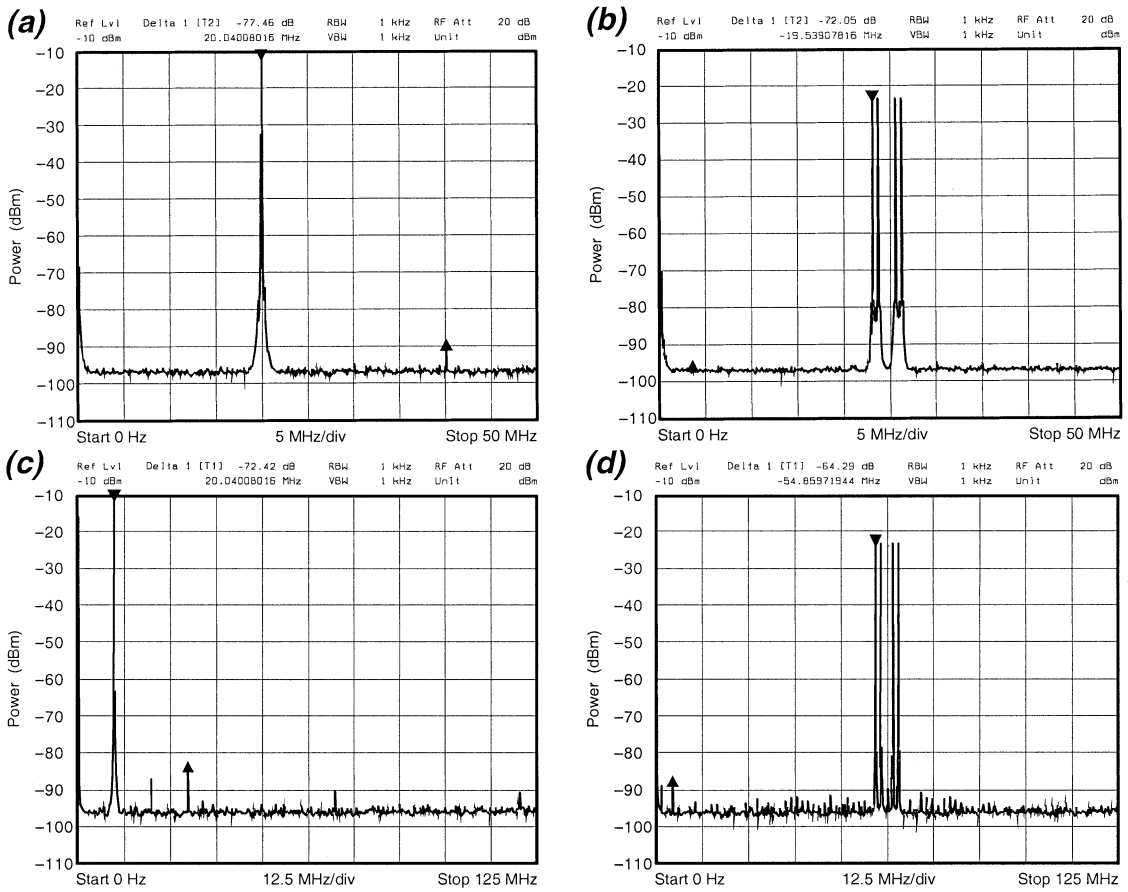

Fig. 10. Dynamic performance. (a), (b) Single and multitone spurious performance at 100 MS/s. For single tones, the SFDR exceeds 77 dB; the four-tone intermodulation performance is better than 72 dB (limited by the SNR of the measurement system rather than by any discrete spur). Parts (c) and (d) show single and multitone spurious performance at 250 MS/s. For single tones, the SFDR exceeds 72 dB; the four-tone intermodulation performance is better than 64 dB.

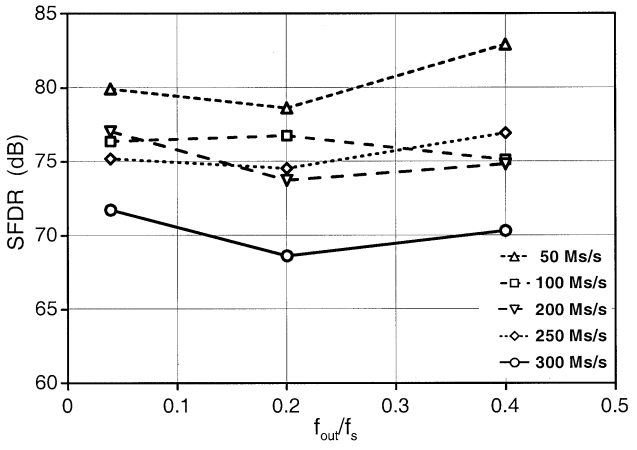

Fig. 11. SFDR versus output frequency for different sample rates. The plot shows SFDR versus the normalized DAC output frequency, defined as the ratio of the output sinusoid frequency divided by the sample rate, for sample rates from 50 to 300 MS/s.

for the analyzer noise, the DAC ACPR measures 62 dB. This result exceeds the Third-Generation Partnership Project (3GPP) WCDMA specification of 45 dB for an immediately adjacent channel and 50 dB for other channels.

We use an RZ circuit at the DAC output to improve high-frequency spurious performance. Fig. 10 shows the DAC's spur-free dynamic range (SFDR) to be about 72 dB at a sample rate of 250 MS/s and 77 dB at a sample rate of 100 MS/s. Fig. 10 also shows multitone performance at 100 and 250 MS/s. Fig. 11 shows that the SFDR does not degrade significantly over the full Nyquist band for sample rates up to 300 MS/s.

VII. Process Scalability

To demonstrate the benefits of floating-gate trimming, we ported the 0.25-μm DAC to a 0.18-μm process merely by resizing the contacts and vias. Processing changed the digital-decoder and polarity-register transistors from the 50-Å devices in the 0.25-μm process to 32-Å devices in the 0.18-μm process. Our trim cells continued to use 70-Å devices to ensure long-term retention. Most major foundry processes maintain a 70-Å gate-oxide option down to 90-nm technologies. Fig. 12(a) shows the post-trim INL of the ported DAC to be $< \pm 0.35$ LSB. Fig. 12(b) shows the SFDR to be > 76.5 dB at 250 MS/s. By relying on trimming rather than intrinsic device matching, we can rapidly port the DAC both to smaller geometries and to different foundries.

VIII. Conclusion

Our DAC dissipates 53 mW with a 10-mA output at a 250-MHz clock rate, comprising 39 mW from a 3.3-V analog supply and 14 mW from a 2.5-V digital supply. The floating-gate trim is crucial to the DAC performance; it

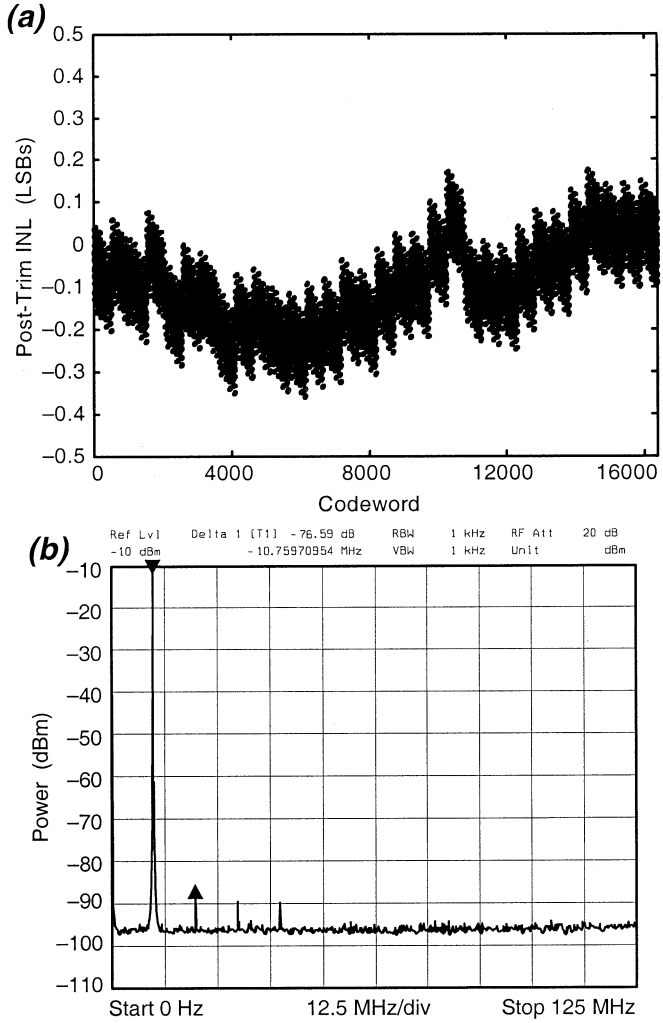

Fig. 12. Performance of a 0.18-μm DAC. We ported the DAC layout from a 0.25-μm logic process to a 0.18-μm logic process simply by scaling contacts and vias. During testing, we reduced the digital power supply from 2.5 to 1.8 V to accommodate the lower breakdown voltages of digital transistors in a 0.18-μm process. Part (a) shows the resulting post-trimmed INL to be $< \pm 0.35$ LSB. Part (b) shows the SFDR at a 250-MS/s clock rate shown to be 76.5 dBc, which is about 4 dB better than the 0.25-μm DAC. The performance improvement resulted from the faster transistors and lower capacitive loading of the 0.18-μm process. Using trimming rather than inherent matching leads to rapid design portability between processes as well as foundries.

vastly reduces die size, improves linearity, eliminates continuous-calibration trim spurs, and reduces power consumption. Furthermore, floating-gate trim allows fabrication in standard CMOS logic processes with no additional process masks (i.e., the DAC design uses *only* nFETs and pFETs). Finally, floating-gate trimming facilitates rapid portability between processes and foundries. Because our DAC is compatible with logic CMOS, we anticipate integrating it with embedded logic circuits to enable precision mixed-signal systems-on-a-chip.

REFERENCES

[1] J. Sevenhans and Z. Chang, "A/D and D/A conversion for telecommunications," *IEEE Circuits Devices Mag.*, vol. 14, pp. 32–42, Jan. 1998.

[2] A. Van den Bosch, M. Borremans, M. Steyaert, and W. Sansen, "A 12-b 500-Msample/s current-steering CMOS D/A converter," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, 2001, pp. 366–367.

[3] A. R. Bugeja and B.-S. Song, "A self-trimming 14-b 100-MS/s CMOS DAC," *IEEE J. Solid-State Circuits*, vol. 35, pp. 1841–1852, Dec. 2000.

[4] G. A. M. Van-Der-Plas, J. Vandenbussche, W. Sansen, M. S. J. Steyaert, and G. G. E. Gielen, "A 14-bit intrinsic accuracy Q^2 random walk CMOS DAC," *IEEE J. Solid-State Circuits*, vol. 34, pp. 1708–1718, Dec. 1999.

[5] C. Diorio, S. Mahajan, P. Hasler, B. A. Minch, and C. Mead, "A high-resolution nonvolatile analog memory cell," in *Proc. IEEE Int. Symp. Circuits and Systems*, vol. 3, 1995, pp. 2233–2236.

[6] M. Figueroa, J. Hyde, T. Humes, and C. Diorio, "A floating-gate trimmable high-resolution DAC in standard 0.25-μm CMOS," in *Proc. IEEE Nonvolatile Semiconductor Memory Workshop*, 2001, pp. 46–47.

[7] C. Diorio, "Neurally inspired silicon learning: From synapse transistors to learning arrays," Ph.D. dissertation, Calif. Inst. Technol., Dept. Elec. Eng., Pasadena, CA, 1997.

[8] M. Lenzlinger and E. H. Snow, "Fowler–Nordheim tunneling into thermally grown SiO_2," *J. Appl. Phys.*, vol. 40, no. 6, pp. 278–283, 1969.

[9] K. Rahimi, C. Diorio, C. Hernandez, and M. D. Brockhausen, "A simulation model for floating-gate MOS synapse transistors," *Proc. IEEE Int. Symp. Circuits and Systems*, vol. 2, pp. 532–535, 2002.

[10] M. J. M. Pelgrom, L. Aad, C. J. Duinmaijer, and A. P. G. Welbers, "Matching properties of MOS transistors," *IEEE J. Solid-State Circuits*, vol. 24, pp. 1433–1440, May 1989.

John Hyde (M'95) received the B.S. degree in electrical engineering from the Massachusetts Institute of Technology, Cambridge, in 1972.

He joined Impinj, Inc., Seattle, WA, in June 2000, as Chief Engineer. Prior to joining Impinj, he worked as both a Consultant and an employee at TRW, Inc., Redondo Beach, CA. As a Consultant, from 1978, he specialized in designing high-speed (400 Mb/s to $>$ 10 Gb/s) digital and mixed-signal integrated circuits in silicon bipolar as well as GaAs and InP HBT technologies. As an employee, from 1972 to 1978, he worked primarily designing high-speed hybrid and monolithic circuits including analog-to-digital converters, multiplexers and demultiplexers, and frequency synthesizers

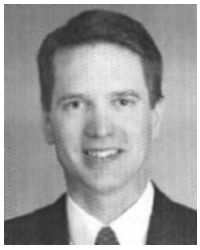

Todd Humes (M'03) received the B.S. degree in electrical engineering from California Polytechnic State University, San Luis Obispo, in 1987 and the M.S. degree in electrical engineering from the California Institute of Technology, Pasadena, in 1989.

He joined Impinj, Inc., Seattle, WA, as Chief Technical Officer in June 2000. Prior to joining Impinj, he was with TRW, Inc., Redondo Beach, CA. While at TRW, he served as Product Line Manager for the frequency synthesizers and digital receivers. His specialities include frequency synthesizers and clock recovery loops, optical communications, and high-speed circuit design.

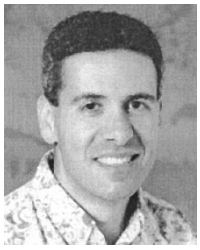

Chris Diorio (M'88) received the B.A. degree in physics from Occidental College, Los Angeles, CA, in 1983, and the M.S. and Ph.D. degrees in electrical engineering from the California Institute of Technology, Pasadena, in 1984 and 1997, respectively.

He is currently an Associate Professor with the Department of Computer Science and Engineering, University of Washington, Seattle, and is a cofounder of Impinj, Inc., Seattle. He has worked as a Senior Staff Engineer at TRW, Inc., Redondo Beah, CA, as a Senior Staff Scientist at American Systems Corporation, and as a Technical Consultant at The Analytic Sciences Corporation. He has published more than 50 journal and conference articles.

Dr. Diorio has received numerous awards including a Sloan Fellowship, a Packard Fellowship, a Presidential Early Career Award, an ONR Young Investigator Award, the University of Washington Distinguished Teaching Award, and the IEEE EDS Paul Rappaport Award.

Mike Thomas received the B.S. and M.S degrees in electrical engineering from Virginia Polytechnic Institute and State University, Blacksburg.

He joined Impinj, Inc., Seattle, WA, in June 2000, as Lead RF Engineer. From 1996 until 2000, he worked for Allen Telecom, Inc., developing commercial cellular infrastructure products. He designed low-power wireless communications systems and circuits for government applications for EVI, Inc. from 1991 to 1996. From 1986 to 1991, he developed high-speed digital and analog-to-digital data acquisition systems and hardware for government customers at American Systems Corporation.

Miguel Figueroa received the B.Sc., Professional Engineering, and M.S. degrees, all in electrical engineering, from the University of Concepción, Concepción, Chile, in 1988, 1991, and 1997, respectively. He received the M.S. degree in computer science from the University of Washington, Seattle, in 1999, where he is currently working toward the Ph.D. degree.

His research interests include mixed-signal VLSI design, silicon implementations of machine learning and signal processing algorithms, and reconfigurable architectures.