

A Floating-Gate Trimmed, 14-Bit, 250 Ms/s Digital-to-Analog Converter in Standard 0.25 μ m CMOS

John Hyde, Todd Humes, Chris Diorio, Mike Thomas, Miguel Figueroa

Impinj, Inc.

501 N. 34th Street, Suite 100, Seattle, WA 98103

Phone: +1 (206) 517–5300, Fax: +1 (206) 517–5262, Email: {john.hyde, todd.humes, diorio, mike.thomas, miguel}@impinj.com

Abstract

We describe a floating-gate trimmed, 14-bit, 250Ms/s current-steered DAC fabricated in a 0.25 μ m CMOS logic process. We trim the static INL to ± 0.3 LSB using analog charge stored on floating-gate p FETs. The DAC occupies 0.44mm² of die area, consumes 53mW at 250MHz, allows on-chip electrical trimming, and achieves 72dB SFDR at 250Ms/s.

Introduction

Emerging standards for communications systems require digital-to-analog converters (DAC) with sample rates in the hundreds of Ms/s, 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 p FETs, 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 ± 0.3 LSB INL (an order of magnitude improved over the DAC in [6]), and dynamic performance that benefits from using small transistors.

DAC Architecture

Fig. 1 shows the DAC architecture, and Fig. 2 a die plot. The DAC is segmented as 5 thermometer-decoded MSBs, 9 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 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




Fig. 1. DAC block diagram. The DAC comprises a 5-bit thermometer-decoded MSB section and a 9-bit binary LSB section. The RZ switch reduces output glitch energy. The trim current sources adjust each bit over a ± 5 LSB range. The extra LSB b_{trim} (also trimmable) and the CDS comparator are used only during the trim process.

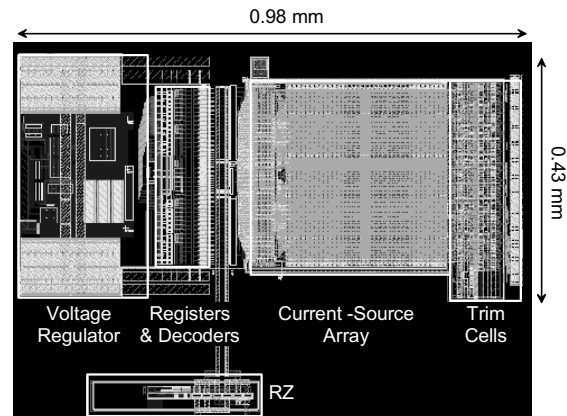


Fig. 2. DAC layout. The total die area is 0.44mm²; the trim circuits account for 14% of this area. We used a 0.25 μ m, 5-metal, single-poly standard CMOS process.

in a single current-source array. We optimized the transistor placement to minimize errors from both linear and quadratic bus-drop gradients. 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.




Fig. 3. A floating-gate *p*FET and its associated tunneling junction, showing the electron tunneling and injection locations. Both *p*FETs share a common floating gate.

Electrically Trimmable Current Source

The heart of our current-source trim is a floating-gate *p*FET. Fig. 3 shows the structure, comprising 2 *p*-channel MOSFETs with a shared floating gate. We use the first transistor for hot-electron injection and a second, with shorted drain and source, 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 standard CMOS processing.

Charge stored on the floating gate determines the gate potential, and, consequently, the channel current of a current-source *p*FET 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 ~4.8V drain-to-source across the injection transistor to cause IHEI at the drain, and ~10V to the shorted transistor’s drain, source, and well to tunnel electrons off the floating gate. The tunneling transistor’s *n*⁻ well acts as a high-voltage implant, allowing us to apply 10V (causing FN tunneling through the shorted transistor’s gate oxide) without incurring *pn*-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 *p*FETs are 3.3V devices with 70Å gate oxides. Accelerated leakage experiments on the floating-gate trim cell shows that leakage-induced changes in the output current will not cause the DACs INL to exceed 1LSB after 10 years. However, to ensure peak performance, we have designed the DAC to allow periodic self-calibration, either at power-up or on command.

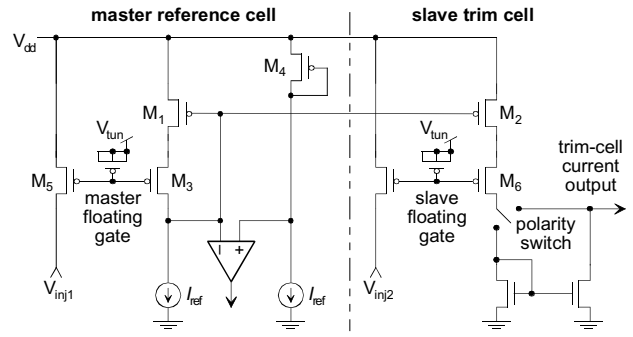


Fig. 4. Trim cell. A single master reference controls 41 slave cells. The master cell adjusts *M*₂’s gate voltage to ensure that floating-gate transistor *M*₆ provides a constant output current despite temperature variations in transconductance and mobility. We trim the floating-gate voltages using tunneling and injection (see Fig. 3).

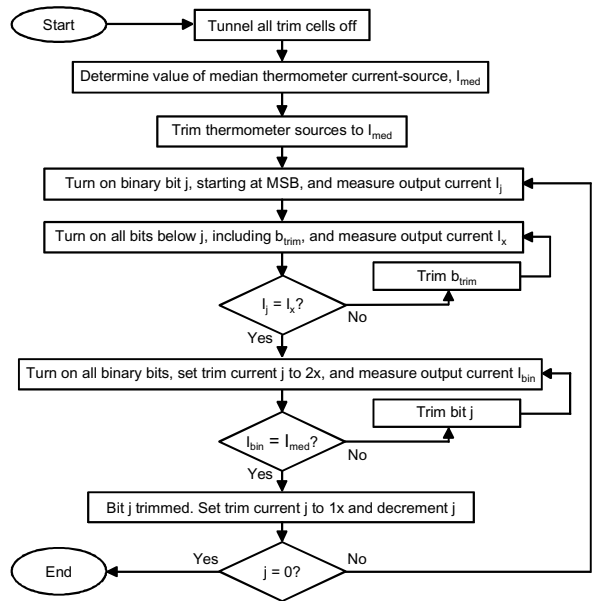


Fig. 5. Trim procedure. We first turn off all the current sources, then successively turn on and trim individual sources, proceeding from MSB to LSB. The algorithm needs only an on-chip correlated-double-sampling comparator (i.e. no off-chip instruments).

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 *p*FET’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 *p*FET *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.

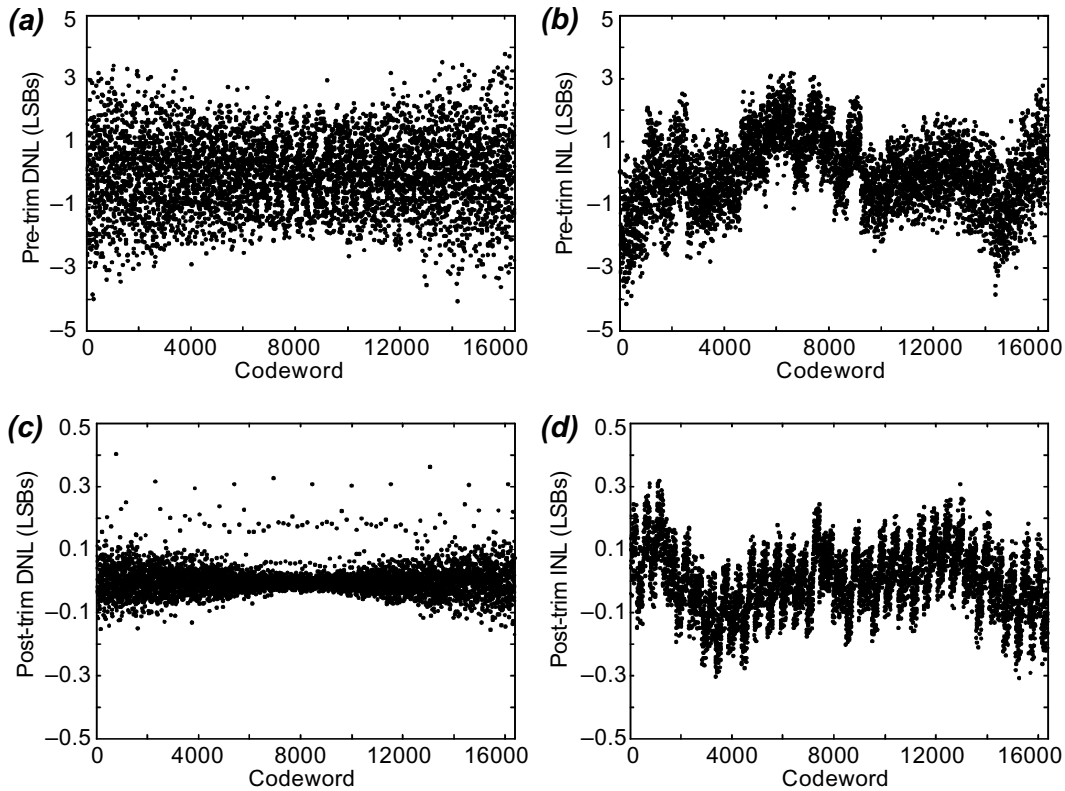


Fig. 6. Static performance. Parts (a) and (b) show the DAC's pretrim differential nonlinearity (DNL) and integral nonlinearity (INL); parts (c) and (d) show the post-trim DNL and INL. Trimming improves both DNL and INL by roughly a order of magnitude.

We adjust the charge on slave-cell p FET 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. A polarity switch determines whether the trim cell adds or subtracts its output from the corresponding bit's static output current.

We anticipate that our DACs will undergo periodic retrim, either during system idle states or during power-up. This periodic retrim will compensate both MOSFET V_t drift and any small leakage from the floating gate. Consequently, we have designed the DAC for on-chip trimming. The trim algorithm does not require any external signals (other than the DAC reference current); rather, it uses on-chip error signals to trim each current source in turn.

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 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_{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 10V to the 41 well-tunneling junctions. Our present design uses an off-chip high-voltage source; future designs will incorporate a 10V charge pump on chip.

After tunneling, we trim the cells using injection. Electron injection requires that we apply roughly -1.5V to the injection p FET drains. We generate -1.5V using individual single-stage charge pumps for each of the 41 trim cells. Interestingly, because peak injection efficiency in p FET floating-gate devices occurs with channel currents near threshold, we can use small charge-pump capacitors. Therefore, 41 individual charge pumps are smaller than one large charge pump and a decoder.

Performance

Fig. 6 shows measured DNL and INL, both before and after trimming. The DAC uses CMOS matching techniques [9] to obtain a pretrim INL of roughly $\pm 4\text{LSBs}$. Trimming improves the INL to less than $\pm 0.3\text{LSB}$. If we had designed for 0.3LSB 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-

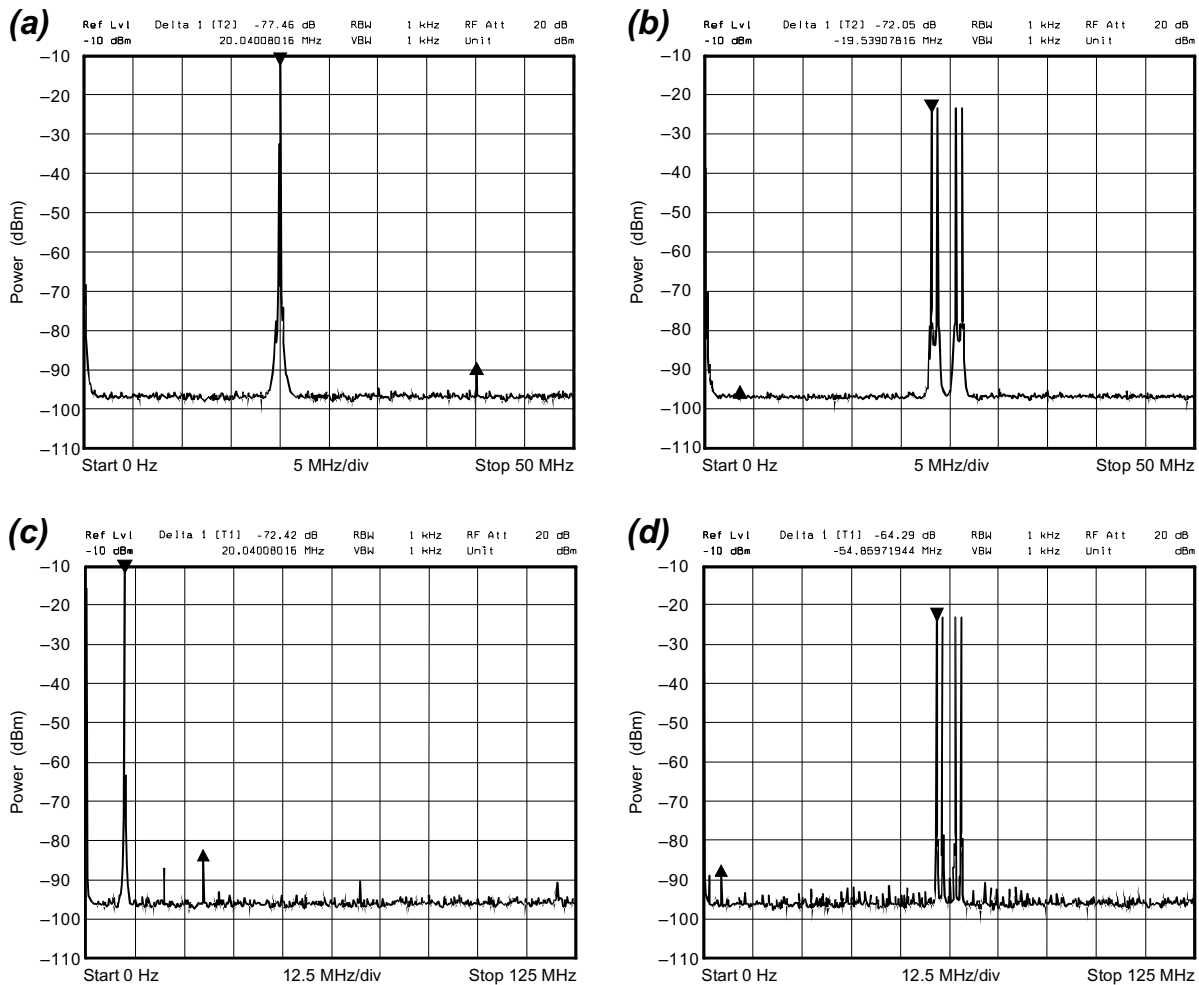


Fig. 7. Dynamic performance. Parts (a) and (b) show single and multi-tone spurious performance at 100 Msps. For single tones, the SFDR exceeds 77 dB; the four-tone intermodulation performance is better than 72 dB (limited by the SNR of the measurement system rather than by any discrete spur). Parts (c) and (d) show single and multi-tone spurious performance at 250 Msps. For single tones, the SFDR exceeds 72 dB; the four-tone intermodulation performance is better than 64 dB.

gate trim, we would have needed trim continuously, increasing die area and degrading dynamic performance.

We use a return-to-zero (RZ) circuit at the DAC output to improve spurious performance. Fig. 7 shows the DAC 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; the SFDR does not degrade significantly over the full Nyquist band. Fig. 7 also shows multi-tone performance at 100 and 250 Ms/s.

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 vastly reduces die size, improves linearity, eliminates continuous-calibration trim spurs, and reduces power consumption (because all the MOSFETs are small). 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). We anticipate integrating our DAC with embedded logic circuits, to enable precision mixed-signal SOCs.

References

- [1] J. Sevenhans and Z. Chang, "A/D and D/A Conversion for Telecommunications," *IEEE Circuits and Devices Mag.*, pp. 32–42, Jan. 1998.
- [2] A. Van den Bosch, M. Borremans, M. Steyaert, W. Sansen, "A 12 b 500 Msample/s current-steering CMOS D/A converter," *Proc. IEEE Solid-State Circuits Conf.*, San Francisco, CA, pp. 366–367, 2001.
- [3] A. R. Bugeja and B-S Song, "A self-trimming 14-b 100MS/s CMOS DAC," *IEEE J. Solid-State Circuits*, vol. 35, no. 12, pp. 1841–1852, 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, no. 12, p. 1708–1718, 1999.
- [5] C. Diorio, S. Mahajan, P. Hasler, B. A. Minch, and C. Mead, "A high-resolution nonvolatile analog memory cell," *Proc. IEEE Intl. Symp. on Circuits and Systems*, Seattle, WA, vol. 3, pp. 2233–2236, 1995.
- [6] M. Figueroa, J. Hyde, T. Humes, and C. Diorio, "A floating-gate trimmable high-resolution DAC in standard 0.25 μ m CMOS," *Proc. IEEE Nonvolatile Semiconductor Memory Workshop*, Monterey, CA, pp. 46–47, 2001.
- [7] C. Diorio, *Neurally Inspired Silicon Learning: From Synapse Transistors to Learning Arrays*, Ph.D. thesis, Department of Electrical Engineering, California Institute of Technology, Pasadena, CA, 1997.
- [8] M. Lenzlinger and E. H. Snow, "Fowler–Nordheim tunneling into thermally grown SiO_2 ," *J. of Appl. Phys.*, vol. 40, no. 6, pp. 278–283, 1969.
- [9] M. J. M. Pelgrom, L. Aad, C. J. Duijnmaier, A. P. G. Welbers, "Matching properties of MOS transistors," *IEEE J. Solid-State Circuits*, vol. 24, no. 5, pp. 1433–1440, 1989.