A 400-MHz Input Flash Converter with Error Correction

CHRISTOPHER W. MANGELSDORF, MEMBER, IEEE

Abstract—An 8-bit, 200-megasample/second (Ms/s) flash converter with 400-MHz analog bandwidth and new error correction circuitry is described. A unique cascoded input stage and a dense bipolar process make the wide bandwidth possible. Three different strategies are employed to reduce errors arising from high input slew rate and comparator metastability.

I. Introduction

THE CURRENT trend in flash converters is to identify performance with sampling rate. This is misleading because the purpose of faster converters is to digitize higher frequency waveforms, not to generate more data in a given amount of time. Bandwidth, distortion, and error rate are of paramount importance.

Although flash converters are usually designed to handle signals up to the Nyquist frequency (half the sampling frequency), bandwidth beyond this is desirable for three reasons. First, the input signal suffers gain compression far below the −3-dB point of the response, and there is invariably serious distortion by the time the amplitude starts to roll off. Second, many instrumentation applications drive flash converters with sample-and-hold circuits where the ADC does not see a continuous signal but rather a series of step-like inputs. In these situations, fast step response—implying wide analog bandwidth—is crucial. Finally, there is a significant class of RF applications where signals above the Nyquist rate are intentionally under-sampled and aliased to a lower frequency. Digital oscilloscopes are perhaps the most common example of this.

Error rate is another aspect of performance which is often overlooked. When the error rate is extremely high, the digitized waveform appears to disintegrate and the signal-to-noise ratio deteriorates rapidly. This situation is easily identified and is universally recognized as unsatisfactory. For many applications involving time-domain analysis, however, error rates down to the parts-per-billion level are required. Often these are systems which attempt to detect narrow pulses and can be confused by ADC glitches. This issue has been ignored in converter specifications largely because of the difficulty in verifying error performance.

A design for an 8-bit, 200-megasample/second (Ms/s) flash converter with 400 MHz of small signal bandwidth and three different methods of error prevention will be described. After an explanation of the overall architecture, the design of the comparator preamp and its bandwidth-enhancing features will be discussed. Errors due to comparator metastability are reduced by the latching stages as described in Section II-C, and errors due to rapidly slewing inputs are corrected by the circuitry described in Section II-D. The final defense against errors is the "second rank error suppression," presented in Section II-E. Section III elaborates on the bipolar process and the particular characteristics that make it appropriate for this device. Finally, measured frequency and error rate performance are examined in Section IV.

II. Circuit Description

A. Overall Architecture

The architecture of the ADC is the well-known flash design with several refinements. Fig. 1 shows how the 257 comparators are arranged in four blocks for convenience of layout and encoding. The reference ladder, made of first-layer metal, has force and sense taps for Kelvin connection and allows for accurate setting of the reference voltage. Midpoint and quarter-point taps can be used for linearity adjustment but find their biggest application in decoupling high-frequency transients from the reference chain. The output code corresponding to overrange can be changed from all ONE's to all ZERO's under the control of the "ORZ" input. This allows the outputs of two converters to be wire-ORed together in a 9-bit converter arrangement.

Perhaps the most unusual feature evident in Fig. 1 is the inclusion of an underrange output as well as the more common overrange output. These two outputs are useful in dynamically adjusting the front-end AGC stages that are ubiquitous in RF systems. Since the outputs are ECL, they can simply be wired together to create a DATA-VALID signal and used to flag invalid data in a high-speed digital signal processor.

Manuscript received August 18, 1989; revised October 30, 1989.
The author is with Analog Devices Semiconductor, Wilmington, MA 01887.
IEEE Log Number 8933138.

0018-9200/90/0200-0184$01.00

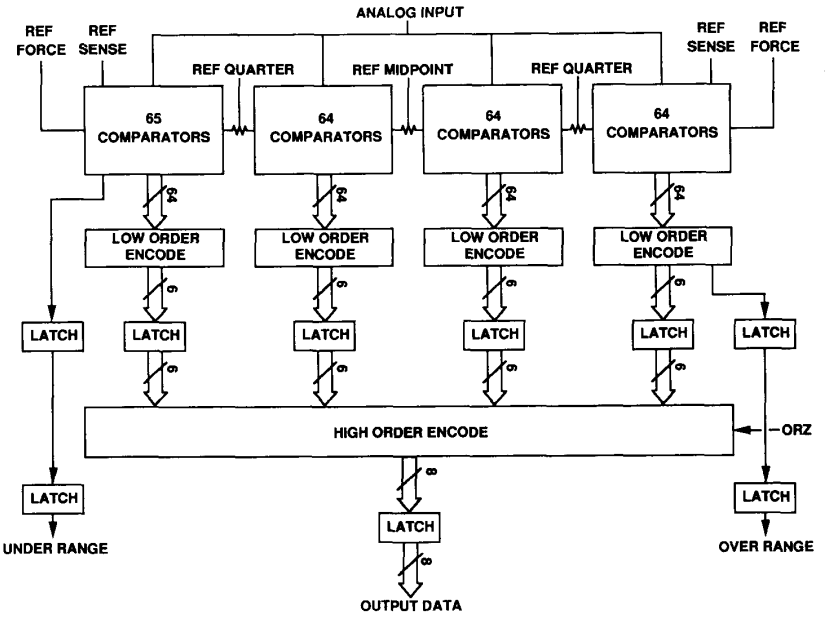

Fig. 1. Block diagram.

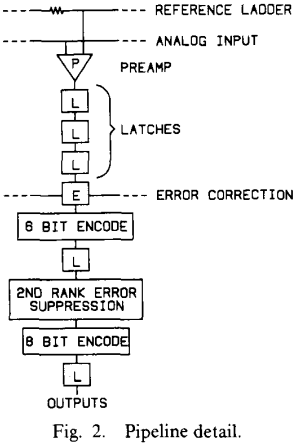

Fig. 2. Pipeline detail.

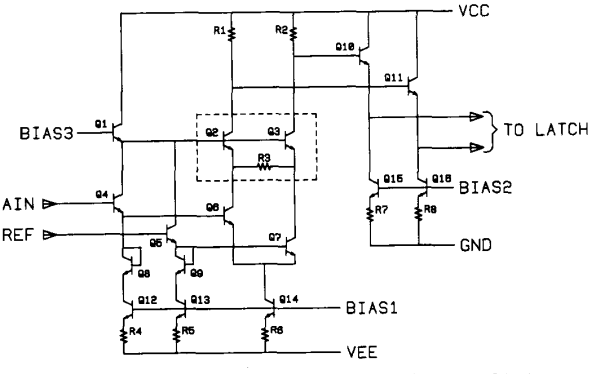

Fig. 3. Preamp schematic. V_{cc} is +5 V and V_{ee} is −5.2 V.

Fig. 2 provides more detail of the serial path from input to output. The analog signals are buffered from the digital circuitry by a wide-band preamp which serves to isolate the digital switching transients or "kickback" of the latches from the input and reference lines. Wide gain bandwidth is essential in this stage to switch the state of the subsequent latch quickly when tracking a rapidly moving analog signal. Following the latch stages of the comparator is the error correction circuitry intended to minimize spurious codes. At this point in the pipeline, each block of 64 comparators acts like an individual flash converter and generates a 6-bit output code. A second rank of error suppression selects which of the four blocks has the correct data and the appropriate 6-bit output is converted to an 8-bit code and passed to the chip outputs. The five latches in the pipeline create a latency of one and a half cycles in addition to the propagation delay of the output circuitry.

B. Preamp Design

The comparator preamp, shown in Fig. 3, was designed for dual supplies (+5 and −5.2 V) for three reasons. First, dual supplies allow more flexibility of input range, including the possibility of a convenient ground-centered full-scale span. Second, high-speed circuits often produce unruly transients, such as overshoot, which are difficult to clamp. The extra headroom afforded by dual supplies permits input excursions beyond full scale and insures a graceful overload recovery. Finally, the large input range means that at least 1 V of reverse bias is maintained across the base–collector junctions of the input transistors ($Q4$ in Fig. 3) under normal conditions (±1-V full-scale range). High reverse bias leads to less signal-dependent variation

in junction capacitance and mitigates a significant source of distortion. [1]

A unique feature of the preamp is the use of cascode transistors $Q2$ and $Q3$. These reduce Miller capacitance loading and protect the differential pair $Q6$ and $Q7$ from excessive collector voltage. Switching speed is increased by resistor $R3$ between the cascode emitters because it prevents the discharging of emitter-junction capacitance when signal current is switched entirely to one side. Without $R3$, $Q2$ or $Q3$ could turn off completely and any change of state would entail recharging the junction so the transistor could turn on again.

Preamp bias levels are set to permit input operation anywhere between $+2$ and -2 V. Diodes $Q8$ and $Q9$ insure that the widest possible input range is achieved by reducing the collector voltage seen by current sources $Q12$ and $Q13$ to the same level seen by $Q14$. The small physical size of the input devices (5.4×1.0-μm emitters) leads to a low input capacitance of 20 pF, and contributes to wide bandwidth.

C. Metastability Prevention

Flash converter latches have traditionally been arranged in a master/slave configuration. In the current design, a three-latch cascade was found to provide a lower error rate for the same total power consumption. The error mechanism of concern is the occasional inability of a comparator to resolve a small differential input into a valid logic level, a phenomenon known as metastability. (The name refers to the ability of a comparator to balance right at its threshold for a short period of time.) This mechanism has long been understood as a limitation in the design of high-speed flash converters. [2]

When the latch is switched from transparent mode to regenerative "latched" mode, the transient appearing at the outputs can be approximated by a simple expression [3]:

$$V_o = AV_i \exp(t/T)$$

where V_o is the output of the latch, A is the gain in "transparent" mode, V_i is the input voltage, t is the time since the onset of positive feedback, and T is the time constant of the latch in positive feedback. The output voltage roughly follows this growing exponential until it reaches the valid logic level. During this period we can consider the regenerative circuit to have an effective gain which is a function of time:

$$A_{\text{eff}} = V_o/V_i = A\exp(t/T).$$

As the clock rate of the ADC is increased, the amount of time the latch spends in positive feedback is reduced and the effective gain reached at the end of half a clock cycle is lower. Lower effective gain makes it more likely that an input voltage will be too small to be amplified to a full logic level and thus increases the chances for an error. Previous converters have used two cascaded latching stages strobed on opposite phases of the clock, making the total time in positive feedback one whole clock cycle. Improvements to error rate must then be made through the time constant T by increasing the gain bandwidth of the latch. (Changing the linear term A has comparatively little effect on the error rate.) In this design, however, it was found that a three-latch design was more effective than a two-latch design for the same total power consumption.

TABLE I
EXAMPLE PARAMETERS FOR TWO- AND THREE-LATCH DESIGNS

Parameter	Two-Stage	Three-Stage
A_P	8.1	8.1
A_1	4.3	4.3
T_1	300 ps	350 ps
A_2	5.7	5.7
T_2	290 ps	340 ps
A_3	-	5.7
T_3	-	340 ps

In the simplest analysis, positive feedback time is increased by 50 percent in going to three latches because now the regeneration goes on for three half clock cycles rather than two. The effective gain, A_{e2}, of a two-stage design is:

$$A_{e2} = A_p[A_1\exp(t_1/T_1)][A_2\exp(t_2/T_2)].$$

For a three-stage design this becomes

$$A_{e3} = A_p[A_1\exp(t_1/T_1)][A_2\exp(t_2/T_2)][A_3\exp(t_3/T_3)]$$

where

A_p	preamp gain,
A_1, A_2, A_3	transparent-mode gain of latches,
t_1, t_2, t_3	time duration of latch mode for each latch generally equal to half a clock cycle,
T_1, T_2, T_3	latch-mode time constants of latches.

Example parameters for two- and three-latch designs with the same power consumption are shown in Table I. (These parameters were taken from simulations.) The latches in the three-stage design run at 33 percent less power, so their time constants are larger but not 33 percent larger. Consequently, the effective gain of the three-stage design, A_{e3}, is over 700 times greater than the effective gain of the two-stage design, A_{e2}, at 200 MHz. Since the chance of getting an error due to metastability is inversely proportional to effective gain, error rate will be improved by the same ratio.

Projected error rates for the two example designs are shown in Fig. 4. It can be seen that the increase in the latch time constant in the three-latch design is more than offset by the increase in positive feedback time. The only drawbacks of the third latch are another stage of pipeline delay and the increase in die area. Delay is inconsequential in most applications, and because of the density of the process, the extra rank of latches comprises less than 5 percent of the total die area.

Detailed scrutiny of the data propagation shows the foregoing analysis to be a vast oversimplification because 1) it neglects the time required to propagate the new data

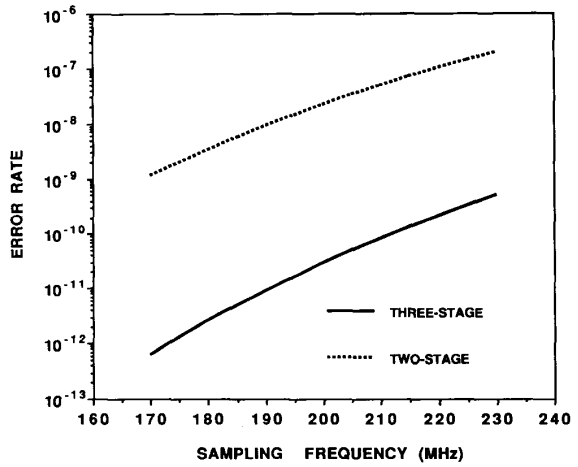

Fig. 4. Theoretical error rates for two- and three-latch designs from the parameters of Table I.

through to the next stage after the metastable state is resolved, and 2) the delay in resolving the state may mean data no longer propagate in synchronism with the clock edges which may result in timing errors both on- and off-chip. If these effects are taken into account, the available regeneration time is greatly reduced, and the addition of the third latch provides an even larger percentage increase in effective gain.

D. Error Correction

Comparators in a flash converter generate what is commonly known as a "thermometer code." If a particular comparator's reference point is below the level of the input signal, the comparator's output is high; if the reference point is above the input, its output is low. When everything is working ideally, the collection of comparator outputs should resemble a thermometer: all ZERO's above the input level, all ONE's below. The ZERO-to-ONE transition point rises and falls within the input level. The thermometer code is translated to a final binary output word using the ZERO-to-ONE transition point to address a ROM.

Under extremely high input slew rate conditions, timing differences between signal paths or even slight differences in comparator response time can cause the effective strobe point of one comparator to be quite different from another. Consequently, a ONE may be found above a ZERO in the thermometer code even though this cannot happen at dc. Errors of this type are sometimes referred to as "bubbles" because they resemble bubbles in the "mercury" of the thermometer code. Examples of different types of bubbles are shown in Fig. 5.

Early flash converters used a simple two-input gate to detect the ZERO-to-ONE transition. If bubbles exist, there will be at least two ZERO-to-ONE transition points and the circuitry will address two ROM lines. The resulting binary output thus becomes the bitwise OR of the two ROM outputs. Since comparators are usually arranged in several

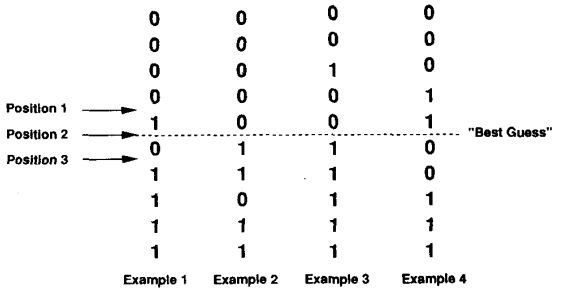

Fig. 5. Examples of thermometer-code bubbles.

rows, this problem is accentuated at the boundaries between rows where the signals to adjacent comparators may take vastly different routes. And because row boundaries usually occur at major carry points, the error resulting from ORed codes can be substantial.

A common method of suppressing bubbles is to use what amounts to a three-input gate to address the ROM [4]. This would, for example, require two ZERO's and a ONE to cause the ROM line to go high, and errors would be avoided in Examples 1 and 2 of Fig. 5 but not in Example 3. There is also a secondary problem concerning accuracy under such conditions. In Example 1, a gate requiring two ZERO's and a ONE will select position 1 as the top of the thermometer. A gate set for two ONE's and a ZERO will select position 3. Since it cannot be known which comparator is in error, the best choice is position 2.

Recently, encoding and error suppression have become more sophisticated. Gray coding is sometimes used, because if done correctly, it will fail benignly in the presence of bubbles or even some of the invalid logic states that result from metastability. However, re-encoding the data in binary form is slow. Akazawa *et al.* [5] struck a compromise with a "quasi-Gray" code which retains much of the advantage of the Gray code but which is far simpler to re-encode to binary. Actual error correction has been implemented by Garuts *et al.* [6] in the form of a hardware bubble sort referred to as the "bitswap" technique.

The error correction scheme in the present circuit may be thought of as a voting process. Each comparator output is examined relative to its two nearest neighbors, and the output is changed if it disagrees with both. For example, the logic equation for the corrected output, $C2^*$, of comparator 2 is

$$C2^* = C1 \cdot C2 + C2 \cdot C3 + C1 \cdot C3$$

where CN is the uncorrected output state of the Nth comparator. Essentially, the corrected output is the majority state of three adjacent comparators: a democratic solution.

Fig. 6 shows how this would be applied to Example 1 (of Fig. 5). COMP1 has a ZERO on one side and a ONE on the other so its output is not changed. COMP2, however, is flanked by two ONE's, so its output is changed to a ONE. Likewise, COMP3 sees a ZERO on both sides (before COMP2 is changed) and reverses its output to ZERO.

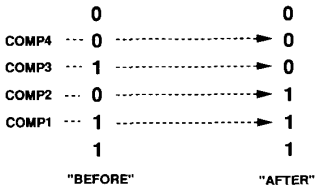

Fig. 6. Bubble correction process.

Looking at the corrected outputs in Fig. 6, we see that the bubble has been removed. This approach will successfully remove the bubbles in Examples 1, 2, and 3 (of Fig. 5) but fails in the extreme case of Example 4. (The error shown in Example 4 represents serious degradation and should only ever happen between different blocks of comparators. This issue is resolved by some less subtle circuitry described in the next section that steps in to squash these large bubbles between blocks.)

After comparator correction, the ONE-to-ZERO transition point must be identified to address the ROM. Conceptually, the circuit looks for a ZERO above a ONE. To continue the previous example,

$$R\{2,3\} = C2^* \cdot \overline{C3^*}$$

where $R\{2,3\}$ is a ROM address line. In practice, the two steps of correction and ONE-to-ZERO identification can be combined into a single operation:

$$R\{2,3\} = C1 \cdot \overline{C4} \cdot (C2 \oplus C3).$$

This last equation leads to very efficient hardware implementation as shown in Fig. 7. (In this schematic "*CN*" refers to the Nth comparator and "*EN*" refers to the Nth correction circuit. There is one correction circuit between each pair of adjacent comparators.) Series gating allows for a simple realization of the XOR term, and the rest of the equation is implemented with only one additional current source at each location. The correction circuit of Fig. 7 makes up about 16 percent of the comparator area or less than 5 percent of the total die. The additional power required is 110 mW or roughly 6 percent of the total ADC power.

The error correction scheme has a number of interesting properties:

1) The "depth" of the bubble (i.e., its distance away from the correct ONE-to-ZERO transition point) makes no difference. In fact, a stuck comparator in either state, which would disable whole sections of a converter with a traditional encoding scheme, will cause only a single missing code with this approach.

2) The logical operation of the correction circuitry may be thought of as a pattern-matching operation. Either of two patterns,

0		0
0		1
1	or	0
1		1

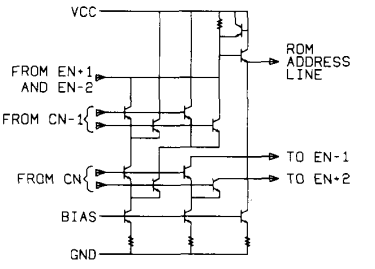

Fig. 7. Error correction circuit. *CN* is the Nth comparator, *EN* is the Nth correction circuit.

will be identified as the ONE-to-ZERO transition point. This is the functional equivalent of two 4-input gates.

3) This scheme always makes the "best guess" of position 2 in Examples 1, 2, and 3 of Fig. 5.

4) If two comparator thresholds become interchanged, as they might when the full-scale range is reduced, this circuitry unscrambles the thermometer code so that monotonicity is maintained and no output codes are missed. (The significance of this is discussed in Section IV.)

E. Second Rank Error Suppression

The outputs of the error correction circuits address a 6-bit wired-OR ROM in each of the four blocks of 64 comparators. However, only one of the blocks will have the correct data corresponding to the lower order bits of the final 8-bit word. Some sort of arbitration circuitry is required to decide which block has the right data and what the two most significant bits should be. Timing differences between adjacent comparators are most pronounced at block boundaries because the signals to different blocks take different paths. Consequently, thermometer-code bubbles are often most severe when the analog input is moving from one block to another. Bubbles can cause two blocks to have nonzero data, and if allowed to propagate to the output, the ORing of the high-order bits makes these errors dramatic.

The "democratic" correction scheme described in the previous section is used to establish which block contains the valid data. A corrected comparator output corresponding to the highest comparator of the block is generated. This signal indicates if the block is "full" in the sense that the input has exceeded the maximum reference voltage within the block or that the "meniscus" of the thermometer code lies in a higher block. If an unfilled block detects a full signal from the block below, it places its data on the output bus and creates the appropriate high-order bits. The advantage of using corrected comparators is that they yield a better estimate of the input level when bubbles appear.

III. PROCESS

The converter was built in the self-aligned polysilicon bipolar process described in Table II. Of all the features listed in the table, density is perhaps the most important.

TABLE II
B.I.T.1 PROCESS

Minimum Mask Feature Size	2μm
Active Transistor Area	14 μm^2
Transistor Pitch	8μm x 16μm
Transistor F_t	> 5 GHz @ 100μA I_c
Metal Layers	2 Au
Metal Pitch	2μm / 2μm
Power Delay	90 fJ
Density	1 Equiv. Gate/Mil^2

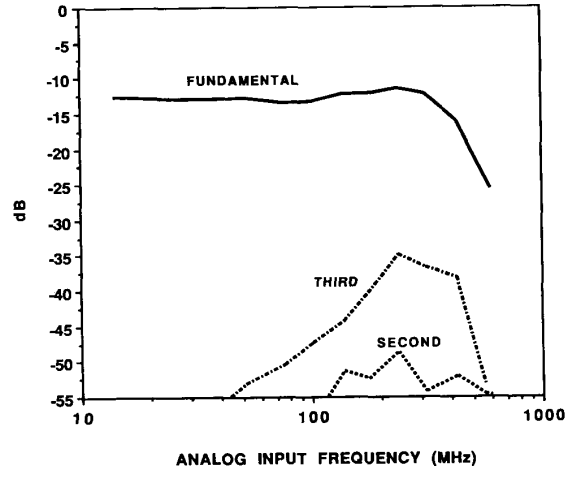

Fig. 8. Dynamic performance at one quarter of full scale.

The wide input bandwidth of the ADC is due in part to the low parasitic capacitance of the transistors which results from their small physical size. High packing density made it possible to add the extra latch and the error correction circuit to each of the 257 comparators.

Although the F_t is modest by today's standards (5 GHz), the low current (100 μA) required to achieve this speed is significant. A higher F_t would be useless for this application if it dictated an operating level of several milliamperes. As a figure of merit, gigahertz of F_t per milliampere of collector current (= 50 GHz/mA in this case) is a better measure of process suitability than peak F_t.

IV. PERFORMANCE

Fig. 8 shows the performance of the ADC versus input frequency at 200-MHz sampling rate with the nominal $\pm$1-V full-scale range. The input was driven directly from a 50-Ω cable terminated in a 6-dB pad. This figure was compiled from DFT analysis of the output data. The 3-dB bandwidth of the quarter-scale signal is about 400 MHz. For full-scale signals, the bandwidth degrades to 250 MHz. The harmonic distortion shown is a consequence of the large input slew rate and declines rapidly with decreasing input amplitude. Since the input follower of the preamp can follow a rising signal better than a falling signal, a slight amount of "pumping" occurs at the highest frequency, and a dc offset is observed.

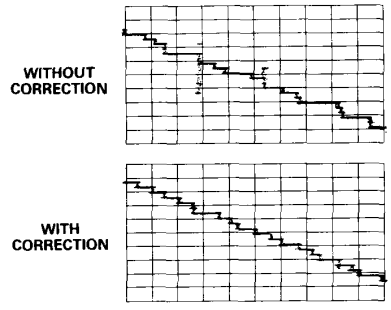

Fig. 9. Effect of error correction. Top trace: Similar converter without error correction. Bottom trace: present converter.

The success of the error correction can be seen in a low-frequency cross-plot when the full-scale range is reduced. For smaller ranges, comparator offsets become large relative to an LSB and thresholds can become scrambled, resulting in thermometer-code bubbles even for slowly varying inputs. Fig. 9 shows a portion of such a cross-plot when the range is reduced to 500 mV for both the present design (with correction circuitry) and a similar converter on the same process which has no error correction. The improvement can be seen by the absence of glitches in the lower trace.

The error correction has an unexpected effect on the high-frequency performance of the device. For any flash converter, reducing the full-scale range will improve the bandwidth and reduce the distortion that can be achieved for full-scale signals simply because the maximum slew rate is reduced. However, as described above, this can lead to substantial errors when the comparator offsets become larger relative to an LSB. Although the error correction circuitry of the present converter cannot prevent a degradation of differential linearity, it does stop large glitches and allows the converter to fail gracefully. Thus, for a given frequency input, the effective number of bits (at full scale) may actually increase as the full-scale range is reduced. This is illustrated in Fig. 10. For the low-frequency input, slew rate is not an issue and reducing the range degrades the linearity, reducing the effective number of bits. At high frequency, the improvement in distortion outweighs the loss of differential linearity and a net enhancement is seen. The effective number of bits actually becomes more constant with frequency if the full-scale range is lowered.

An application which makes use of this reduced input range performance is shown in Fig. 11. Here, a 400-MHz amplitude modulated carrier is demodulated directly by the ADC set for a 500-mV full scale. (Again, the input was driven directly from a 50-Ω cable terminated in a 6-dB pad.) The bottom trace is the RF input and the top trace is the baseband signal detected by phase locking the 200-MHz sample clock to the input signal and reconstructing the ADC output with a DAC. The center trace is 10-kHz IF output generated by clocking the ADC at 199.995 MHz.

A plot of error rate versus sample frequency is shown in Fig. 12. (The test method is described in [7].) For these

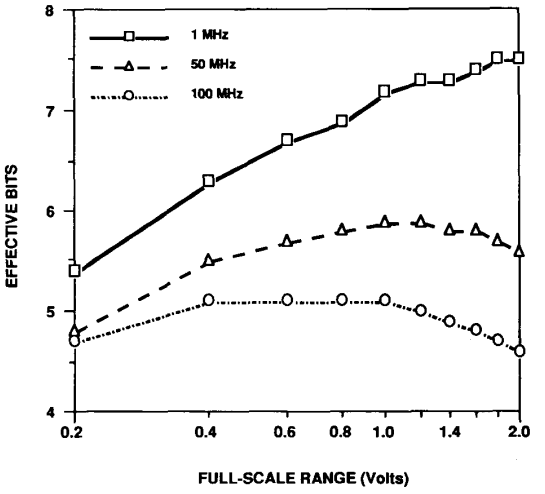

Fig. 10. Effective number of bits versus full-scale range for 1-, 50-, and 100-MHz analog input frequencies, all at 200 Ms/s. Input amplitude was adjusted for each range.

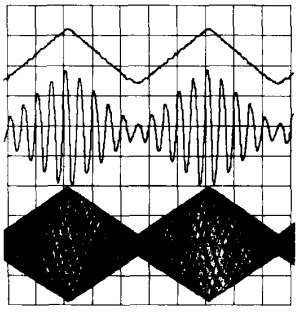

Fig. 11. Direct demodulation of a 400-MHz amplitude modulated carrier. Top trace: baseband output, $F_s = 200.000$ Ms/s. Center trace: 10-kHz IF output, $F_s = 199.995$ Ms/s. Bottom trace: 400-MHz RF input at 200 mV/div. Time axis is 200 μs/div.

measurements, only errors greater than 4 LSB's were tallied. The input was maintained within one block of 64 comparators to avoid confusion with inter-block dynamics and second rank error suppression. The theoretical estimate for metastability-induced errors, plotted along with the data, is the result of a fairly sophisticated analysis of the internal timing, and takes account of the fact that only errors greater than 4 LSB's were recorded. There are still enough parameters of unknown or debatable value to justify a wide range of predictions. Considering the inherent difficulty in making these theoretical projections, the good fit here is not to be taken too seriously. However, the shape of the theoretical curve indicates that the functional form of the error rate is correct even if the absolute value of the predictions is suspicious.

The behavior of error rate as a function of analog input frequency is complex and, as yet, unexplained. It does not show any monotonic increase with frequency that would be expected if, at some point, the error rate became

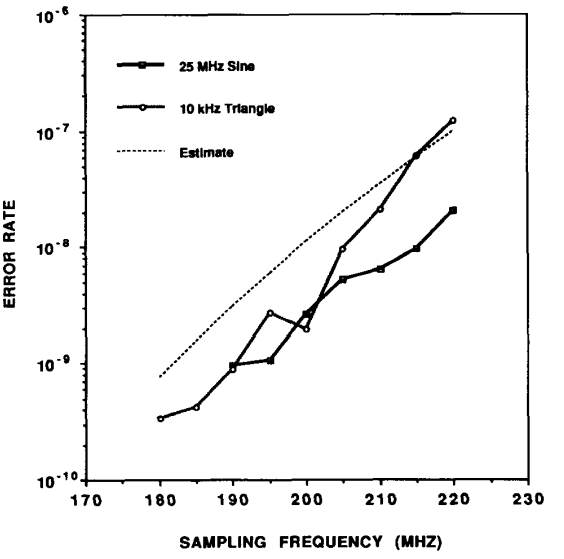

Fig. 12. Measured error rate (in errors/sample) versus sampling frequency.

TABLE III
ADC PERFORMANCE SUMMARY

Resolution	8 Bits
Max. Sample Rate	200 MHz
Large Signal Bandwidth	250 MHz
Small Signal Bandwidth	400 MHz
Input Capacitance	19 pF
Power Dissipation	2 Watts
Die Size	3.30 mm × 4.64 mm

dominated by slew-induced effects. This is further evidence for the success of the error correction circuits.

V. CONCLUSION

A summary of typical characteristics for the ADC is shown in Table III. Design and process combined to yield usable bandwidth up to four times the Nyquist frequency. Although the improvement was not quantified, a clear reduction of thermometer-code errors resulting from high slew rate or low full-scale range was demonstrated. Remaining errors due to metastability were dramatically lowered by the addition of an extra latch in each comparator. Finally, the success of the error correction admits the possibility of low-amplitude applications and allows reduced distortion through the lowering of the full-scale range.

ACKNOWLEDGMENT

The author would like to thank P. Sheppard for the layout, and T. Delmuro for his extensive characterization of the chip and his assistance in preparing the figures for this paper.

REFERENCES

[1] B. Peetz, B. Hamilton, and J. Kang, "An 8-bit 250 megasample per second analog-to-digital converter: Operation without a sample and hold," *IEEE J. Solid-State Circuits*, vol. SC-21, no. 6, pp. 997–1002, Dec. 1986.

[2] B. Zojer, R. Petschacher, and W. Luschnig, "A 6-bit/200-MHz full Nyquist A/D converter," *IEEE J. Solid-State Circuits*, vol. SC-20, no. 3, pp. 780–786, June 1985.

[3] H. Veendrick, "The behavior of flip-flops used as synchronizers and prediction of their failure rate," *IEEE J. Solid-State Circuits*, vol. SC-15, no. 2, pp. 169–176, Apr. 1980.

[4] J. Peterson, "A monolithic video AD converter," *IEEE J. Solid-State Circuits*, vol. SC-14, no. 6, pp. 932–937, Dec. 1979.

[5] Y. Akazawa *et al.*, "A 400 Msps 8b flash AD conversion LSI," in *ISSCC Dig. Tech. Papers*, vol. 30, 1987, pp. 98–99.

[6] V. Garuts, Y. Yu, E. Traa, and T. Yamaguchi, "A dual 4-bit 2-Gs/s full Nyquist analog-to-digital converter using a 70-ps silicon bipolar technology with borosenic-poly process and coupling-base implant," *IEEE J. Solid-State Circuits*, vol. 24, no. 2, pp. 216–222, Apr. 1989.

[7] C. Mangelsdorf, "Improving bandwidth and error rate in flash converters," in *1989 Symp. VLSI Circuits Dig. Tech. Papers*, pp. 53–54.

Christopher W. Mangelsdorf (S'79–M'83) was born in Chester, PA, in 1955. He received the B.S. degree in physics, magna cum laude, from Davidson College, Davidson, NC, in 1977. In 1980 and 1984 he received the M.S. and Ph.D. degrees in electrical engineering from the Massachusetts Institute of Technology, Cambridge, where he held the Analog Devices Fellowship.

He has been associated with Analog Devices Semiconductor, Wilmington, MA, since summer employment in 1980, and has held the position of Senior Scientist since 1983. His current research centers on high-speed converter architectures, sample and holds, and dynamic test techniques.

Dr. Mangelsdorf is a member of Phi Beta Kappa and Sigma Pi Sigma (physics), and serves on the Technical Program Committee for the IEEE Bipolar Circuits and Technology Meeting.