A 10-b 20-Msample/s Analog-to-Digital Converter

Stephen H. Lewis, *Member, IEEE*, H. Scott Fetterman, George F. Gross, Jr.,
R. Ramachandran, *Member, IEEE*, and T. R. Viswanathan, *Senior Member, IEEE*

Abstract—This paper describes a 10-b 20-Msample/s analog-to-digital converter fabricated in a 0.9-μm CMOS technology. The converter uses a pipelined nine-stage architecture with fully differential analog circuits and achieves a signal-to-noise-and-distortion ratio (SNDR) of 60 dB with a full-scale sinusoidal input at 5 MHz. It occupies 8.7 mm² and dissipates 240 mW.

I. INTRODUCTION

TRADITIONAL designs of video-rate analog-to-digital converters (ADC's) have used flash architectures and bipolar technologies to obtain 8-b resolution at a conversion rate of 20 Msamples/s [1]. For more than 8-b resolution, however, flash architectures require such large die areas and power dissipations that interest in using multistage conversion architectures has arisen. Although bipolar technologies still make the fastest converters [2], [3], BiCMOS technologies have also been used to build video-rate multistage ADC's because they provide both high enough conversion rates and the required sample-and-hold capability [4], [5]. The cost of these ADC's, however, is increased by the more complex process technology, and the required power dissipations are still large (at least 750 mW). Thus, reducing the cost and power dissipation with the same or even better performance in CMOS technologies is an important objective.

One previous CMOS implementation of a 10-b video-rate ADC has been reported [6]. It reduced the required power dissipation to 250 mW; however, its conversion rate was limited to 15 Msamples/s and its signal-to-noise ratio (SNR) to 52 dB with a low-frequency input and to 45 dB with a 7.5-MHz input. In this paper, a nine-stage pipelined ADC with the same resolution, almost the same power dissipation (240 mW), a higher conversion rate (20 Msamples/s), and an increased SNR (60 dB with a 5-MHz input) is presented. It uses an improved error-correction algorithm, a reduced resolution per stage, and an unfolded-cascode operational amplifier. The converter occupies 8.7 mm² (13 400 mils²) in a 0.9-μm CMOS technology. Also, with a 5-MHz input, the signal-to-distortion ratio (SDR) is 74 dB and the signal-to-noise-and-distortion ratio (SNDR) is 60 dB. The key innovation in

this ADC is the improved correction algorithm, which requires one less comparator per stage than used in traditional architectures.

This paper is divided into four main parts. In Section II, pipelined ADC's are reviewed. Section III introduces the new digital-correction algorithm. In Section IV, the circuits in the ADC are described. Finally, experimental results are given in Section V.

II. REVIEW—PIPELINED ADC's

Since the characteristics of several pipelined ADC's have been extensively described [7]–[11], only a brief review is presented here for convenience. Fig. 1 shows a block diagram of a general pipelined ADC with k stages. Each stage contains a sample-and-hold amplifier (SHA), a low-resolution analog-to-digital subconverter (ADSC), a low-resolution digital-to-analog converter (DAC), and a subtracter. In operation, each stage initially samples and holds the output from the previous stage. Next, the held input is converted into a low-resolution digital code by the ADSC and back into an analog signal by the DAC. Finally, the DAC output is subtracted from the held input, producing a residue that is sent to the next stage for further conversion.

The main advantages of pipelined ADC's are that they can provide high throughput rates and occupy small die areas. Both advantages stem from the concurrent operation of the stages; that is, at any time, the first stage operates on the most recent sample while all other stages operate on residues from previous samples. (The associated latency is not a limitation in many applications.) If the analog-to-digital subconversions are done with flash converters, pipelined architectures require only two main clock phases per conversion; therefore, the maximum throughput rate can be high. Also, since the stages operate concurrently, the number of stages used to obtain a given resolution is not constrained by the required throughput rate. Therefore, under some constraints (such as the total resolution), the number of stages may be chosen to minimize the required die area [12].

III. REDUNDANCY AND DIGITAL CORRECTION

To build pipelined ADC's with a large tolerance to component nonidealities, redundancy is introduced by making the sum of the individual stage resolutions greater than the total resolution. When the redundancy is eliminated by a digital-correction algorithm, it can be used to

eliminate the effects of ADSC nonlinearity and interstage offset on the overall linearity [9], [12]–[15]. In many previous implementations, digital-correction algorithms have used both addition and subtraction to correct errors. There are two problems with this approach. First, this approach is difficult to test because the correction logic has three options at each stage (to add, subtract, or do nothing) and because none of these options is forced to occur for any output codes [16]. As a result, the option used by each stage cannot be determined by examining the corrected ADC output. Thus, satisfactory performance during a functional test of the ADC does not guarantee that the correction logic is fault free. For example, during functional testing, the ADC could contain a set of comparator offsets that force the correction logic to do only addition. Then the presence of faults in the correction logic inhibiting subtraction would be undetected by the functional test. If these comparator offsets were to change after functional testing so that subtraction were required, the undetected faults could cause the ADC output to be incorrect. Therefore, to test the correction logic thoroughly for faults, test vectors must be injected directly into the correction logic, bypassing the ADC. Second, this approach is unnecessarily complex because subtraction is equivalent to addition after a negative offset; therefore, the need for the correction logic to do subtraction can be eliminated by building offsets into the ADSC and DAC [14].

Fig. 2 shows a block diagram of one stage in a pipelined ADC with offsets in series with both the ADSC and the DAC. A 2-b stage is used as a representative example. The magnitudes of the offsets are both equal to $1/2$ least significant bit (LSB) at a 2-b level ($V_r/4$, where $\pm V_r$ is the full-scale range of the ADC). Fig. 3(a) shows a plot of the ideal residue versus held input without the offsets, and Fig. 3(b) shows the same plot with the offsets. The ADSC offset uniformly shifts the locations of the decision levels to the right, and the DAC offset shifts the entire plot down. Because there are no decision levels at half-scale in Fig. 3(b), a multistage ADC using stages with this transfer characteristic will inherently have excellent linearity at half-scale [16].

Let the correction range be defined as the amount of decision-level movement that can be tolerated without error. If the DAC and SHA are ideal and the interstage gain is 2, the amplified residue from Fig. 3(b) remains within the conversion range of the next stage when ADSC nonlinearity shifts the decision levels by no more than $\pm 1/2$ LSB at a 2-b level. Under these conditions, errors caused by the ADSC nonlinearity can be corrected; therefore, the correction range here is $\pm 1/2$ LSB at a 2-b level or $\pm V_r/4$, which means that the ADSC linearity must only be commensurate with the stage resolution instead of with the entire ADC resolution. Furthermore, because the offset introduced into the ADSC in Fig. 2 shifts the decision levels to the right by the amount of the offset, the digital output is always less than or equal to its ideal value if ADSC nonlinearity can shift the decision levels back to the left by no more than this amount. Thus, the correction
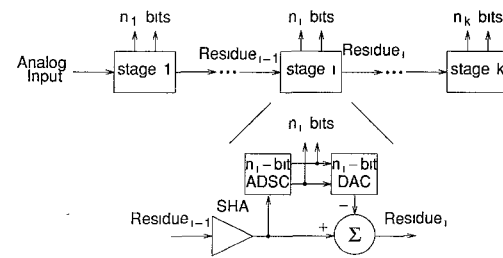

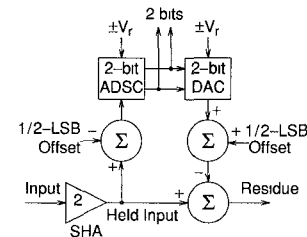Fig. 1. Block diagram of a general pipelined ADC.

Fig. 2. Block diagram of one stage with offsets in both the ADSC and the DAC.
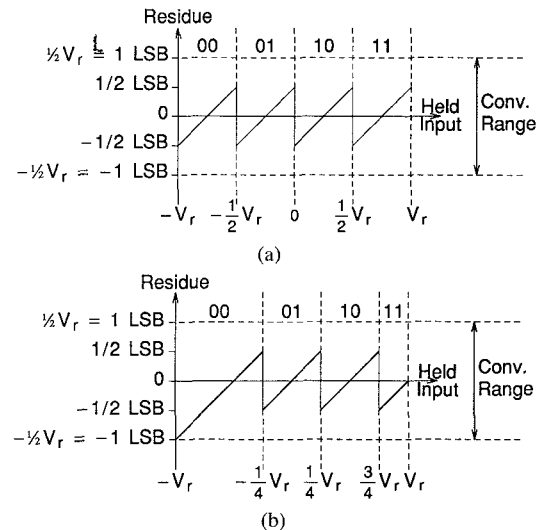

Fig. 3. (a) Ideal residue versus held input (a) without offsets and (b) with offsets.

here requires either no change or addition. Because the associated digital-correction logic need not do subtraction, it is easier to test than conventional correction logic; however, since neither of the remaining two options is forced to occur for any output codes, functional testing of the ADC still does not thoroughly test the correction logic.

While the ideal residue in Fig. 3(a) is always between $\pm V_r/4$, the range of ideal residues in Fig. 3(b) is $-V_r/2$ to $V_r/4$. With identical stages and an interstage gain of 2, the minimum residue in Fig. 3(b) occurs on the left end of the plot and rests on the lower conversion-range boundary of the next stage. Although movement of the decision levels has no effect on the value of this left-end residue, interstage offset or gain error may cause the left-end residue to lie below the conversion range of the next stage. The effects of such interstage offset and gain errors on the ADC linearity are presented next.

An interstage offset error shifts the entire residue plot vertically and is equivalent to the combination of two offsets in the preceding stage: one in the input branch and the other of opposite polarity in the ADSC branch. If the correction range is not exceeded by the combination of this ADSC offset and the nonlinearity present in the ADSC, the effect of the ADSC offset is eliminated by the digital correction, leaving the input-referred offset as the only effect of an interstage offset [9]. Since the correction range is the amount of decision-level movement that can be tolerated without error, and since the value of the left-end residue is independent of the decision-level movement, the correction range is independent of the value of the left-end residue. Therefore, the enlarged left-end residue in Fig. 3(b) (compared to that in Fig. 3(a)) does not increase the sensitivity of the conversion linearity to interstage offset errors.

An interstage gain error scales the entire residue plot vertically and causes an error in the analog input to the next stage when applied to any nonzero residue. If the error in the analog input to the next stage is more than one part in 2^r (where r is the resolution remaining after the interstage gain error), a conversion error that is not removed by digital correction results. Since all nonzero residues are affected by interstage gain errors, the conversion-range boundary has no special significance from a gain-error standpoint. However, the effect of a gain error is largest for the residue with the largest magnitude, which is the left-end residue in the ideal case. In practice, ADSC nonlinearity increases the magnitude of the residue on one side of the affected decision levels, but the magnitude of the left-end residue is still greater than that of any other residue if the ADSC nonlinearity is within the correction range. Therefore, the effect of an interstage gain error on ADC linearity is worsened by the enlarged left-end residue in both the ideal case and in practice. Furthermore, the enlarged left-end residue also increases the swing requirement on the output of the interstage amplifier.

To overcome these problems, an extra comparator can be added to the ADSC to reduce the maximum magnitude of the ideal residue. Fig. 4 shows the resulting ideal residue plot for the 2-b example. (The comparator offsets here are assumed to equal zero.) The threshold of the extra comparator is $-(3/4)V_r$, and this comparator is not used to change the digital output of the stage but only to limit the magnitude of the ideal residue to $V_r/4$. Although this is shown only for a 2-b example, it works for any stage resolution. An example is a cyclic ADC with one extra comparator [17]. The main advantage of this technique is that it reduces the magnitude of the maximum ideal residue to $1/2$ LSB at the level of the stage resolution, reducing the swing requirement on the output of the interstage amplifiers and the sensitivity to gain errors to the same as those in Fig. 3(a). The importance of this advantage is diminished, however, by the presence of ADSC nonlinearity, which increases the maximum magnitude of the residue around the affected decision levels. Furthermore, functional testing of the ADC still does not
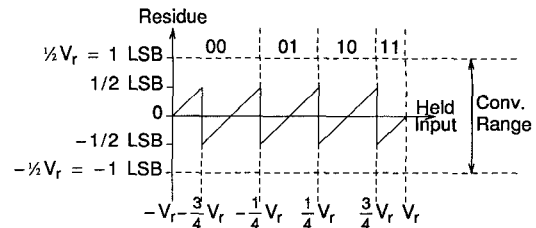

Fig. 4. Ideal residue versus held input with an extra comparator.

thoroughly test the correction logic because no corrected output codes reveal whether an addition or no change was made to the corresponding uncorrected input code. Therefore, the main disadvantages of this technique are that it requires extra comparators and that the associated digital-correction logic is still difficult to test. To overcome these problems, instead of using one more comparator in the ADSC, one less comparator can be used [18].

Because the correction range in Fig. 3(b) is $\pm 1/2$ LSB at a 2-b level, and because the top decision level is $1/2$ LSB at a 2-b level below full scale, Fig. 3(b) also shows that the top comparator in each stage except the last is not needed. Its removal forces a correctable error to occur in the uncorrected ADSC output for a full-scale input. Fig. 5 shows the new ideal residue plot for the 2-b example. (The comparator offsets here are assumed to equal zero.) Without the top comparator, the digital output never reaches code 11 and the residue continues to rise for increasing inputs greater than $1/4$ of the reference. Because the resulting residue on the right side of Fig. 5 has the same magnitude as the left-end residue in Figs. 3(b) and 5, removal of the top comparator does not increase the magnitude of the maximum residue. To obtain code 11 out of this stage after correction, the correction logic must increment the output of this stage. Furthermore, to obtain code 00 out of this stage after correction, since the correction logic cannot subtract, it must do nothing. So the ability of the correction logic to do both nothing and addition can be verified simply by testing the complete ADC for the presence of all its output codes. This simplifies the testing of the digital-correction logic.

Furthermore, after removal of the top comparator, the correction range is still $\pm 1/2$ LSB at a 2-b level because the remaining decision levels can move by up to this amount before the resulting residue exceeds the conversion range of the next stage. Therefore, in this example, only two comparators are needed in each stage except the last. The last stage still needs three comparators because its output cannot be corrected. In general, if n is the number of digital output bits per stage, while the last stage needs $2^n - 1$ comparators, every other stage only needs $2^n - 2$ comparators. As a result, the resolution of each stage except the last is $\log_2 (2^n - 1)$ bits. If $n = 2$ b, as in this example, the resolution per stage is about 1.5 b, and the stage architecture is identical to one proposed by Jusuf [17] because adding a comparator to a 1-b ADSC is the same as removing a comparator from a 2-b ADSC. The decrease in the number of comparators here decreases
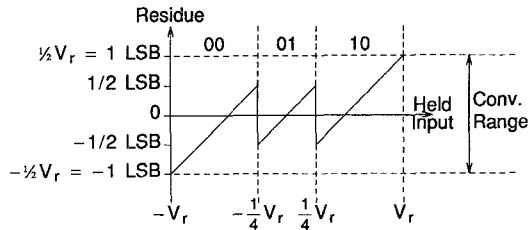
Fig. 5. Ideal residue versus held input without top comparator.

Fig. 6. Block diagram of the prototype.

the total area and power dissipation as well as the capacitive load on each SHA. It also reduces the number of required levels in the DAC by one, which is important because the resulting DAC is faster and less sensitive to capacitor ratio error than its conventional counterpart.

IV. CIRCUIT DESCRIPTION

Fig. 6 shows a block diagram of the prototype. It consists of the ADC core, digital-correction logic, and a clock generator. To maximize the PSRR and minimize the even harmonic distortion, all analog signals are fully differential. Assume that all stages are identical to minimize the design time. Then the most basic architectural characteristic is the individual stage resolution. It determines the number of stages required to obtain 10-b resolution and sets the value of the interstage gain. The choice of the optimum stage resolution is determined by two factors: the conversion rate and linearity. To cover most video-rate applications, a conversion rate of 20 Msamples/s is required. This corresponds to a conversion period of 50 ns, which is divided into two nonoverlapping phases of equal duration by the clock generator. As a result, the op-amp settling time must be less than 25 ns. Furthermore, the op-amp open-loop gain must be more than 2000 to get 10-b linearity in the ADC. To meet these requirements, the minimum stage resolution should be selected because this minimizes the required interstage gain, which, in turn, maximizes the bandwidth since the gain–bandwidth product is limited in any technology. At the same time, however, some redundancy should be included to eliminate the effects of ADSC nonlinearity and interstage offset on overall linearity. To balance these concerns, a stage resolution of 1.5 b is chosen here; that is, there are three possible outputs from each stage. With an interstage gain of 2, each stage contributes 1 b toward the overall resolution. The other $1/2$ b in each stage is redundant. The digital-correction logic eliminates this redundancy and produces a 10-b output. Each stage contains one op amp and two comparators except the last stage, which uses three comparators. Since there are nine stages, nine op amps and 19 comparators are used in all.

In [9], each core stage contained an ADSC and a DAC that shared a common resistor string. While this configuration reduces area, it also increases resistor matching requirements. Here, capacitor-based DAC's are used instead. As a result, the resistors here only determine the ADSC decision levels. Because redundancy and digital correction make the conversion linearity insensitive to
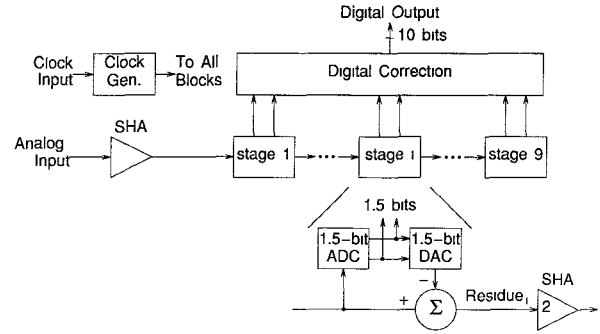
these levels, the conversion linearity no longer depends on resistor matching but only on capacitor matching and op-amp gain. The DAC, subtracter, and SHA all share a common capacitor array, and their functions are combined in a multiplying DAC (MDAC). This is the key circuit in the ADC and is described next.

A full 2-b MDAC with a gain of 2 requires six equal-value capacitors (four for sampling and two for integrating) and is capable of generating five DAC levels. One way to increase the closed-loop bandwidth and speed of such a MDAC is to increase the feedback factor. To do this without changing the gain, two sampling capacitors can be removed, and the MDAC can sample the inputs onto both the sampling and integrating capacitors [19]. Fig. 7(a) shows the schematic of the resulting MDAC. It consists of an op amp, four equal-sized capacitors, and several switches. Fig. 7(b) shows a timing diagram for the clock signals. The two main clocks, ϕ_1 and ϕ_2, are nonoverlapping. To reduce the sample-to-hold transition error, two extra clocks, ϕ_1' and ϕ_1'', are also used [20]. While ϕ_1, ϕ_1' and ϕ_1'' are high, the op-amp inputs are connected to each other and *bias* 6, and the SHA inputs are each connected to both a sampling capacitor C_S and an integrating capacitor C_I. When ϕ_1'' goes low, the op-amp inputs are disconnected from *bias* 6 but remain connected to each other until ϕ_1' goes low, which causes the inputs to be sampled onto all the capacitors. During ϕ_2, the integrating capacitors are connected to the op-amp outputs and the sampling capacitors are connected to each other, the positive reference, or the negative reference depending on the state of digital inputs X–Z. The resulting output consists of two parts: one arising from the feedforward of the integrating capacitors and another arising from the charge transfer between the sampling and integrating capacitors. Since only the second part is ratio dependent, the feedforward reduces the effect of capacitor mismatch on the interstage gain. This is important because the accuracy of the interstage gain of 2 determines the linearity of the ADC. To minimize the gain error without trimming, fully differential, common-centroid capacitor arrays with dummy capacitors surrounding the arrays are used in the MDAC's [21].

Since digital inputs X–Z control only two capacitors, this DAC generates only three levels. Ignoring parasitics, the feedback factor here is $C_I/(C_I + C_S)$, or $1/2$ when C_I
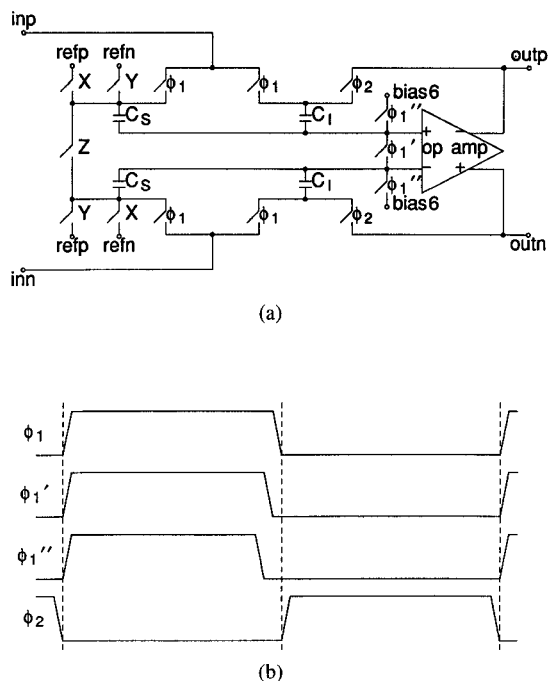
Fig. 7. (a) Simplified schematic of MDAC. (b) Timing diagram.

Fig. 8. Operational-amplifier schematic.

Fig. 9. Comparator schematic.

$= C_S$. Under the same conditions, the feedback factor in a conventional 2-b MDAC (with two extra sampling capacitors) is $C_I/(C_I + 2C_S)$, or $1/3$. If the three-level and 2-b MDAC's use the same op amp, this difference in the feedback factors translates into the same difference in loop gains; therefore, ignoring parasitics, the three-level MDAC is 50% faster than its conventional 2-b counterpart. This is important because the MDAC speed limits the conversion rate of the ADC.

Many op-amp architectures can give an open-loop gain of at least 2000; however, few of those can give a settling time of less than 25 ns in a 1-μm CMOS technology. Previous designs of fast switched-capacitor op amps have used class A/B and folded-cascode architectures [9], [22]. To minimize signal-dependent power-supply current, and to eliminate the need for p-channel transistors in the signal path, a class-A unfolded-cascode op amp is used here [23], [24]. Fig. 8 shows the op-amp schematic. It consists of an input differential pair (M_1, M_2), a tail current source (M_3), two levels of n-channel cascodes (M_4-M_7), and p-channel double-cascode current sources (M_8-M_{13}) as loads. Switched-capacitor common-mode feedback [25] and a high-swing bias circuit [26] are used but not shown for simplicity. According to simulation, the op-amp gain is about 80 dB and its output settles in 20 ns to 0.05% with a 4-V differential output into a 3-pF load. This circuit is different from the one shown in [18], in which only a single level of n-type cascodes was used. The extra level of cascodes here was inserted to increase the op-amp open-loop gain by a factor of about 4 in an effort to reduce the ADC nonlinearity reported in [18].

The ADSC consists of a resistor string, comparator bank, and encoder and is constructed as in [9] except that here only two comparators per stage are required. Fig. 9
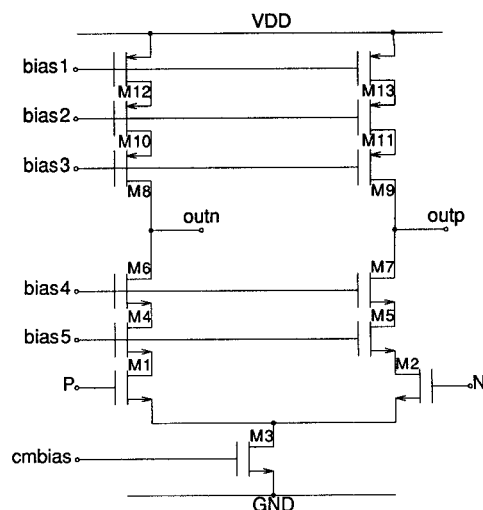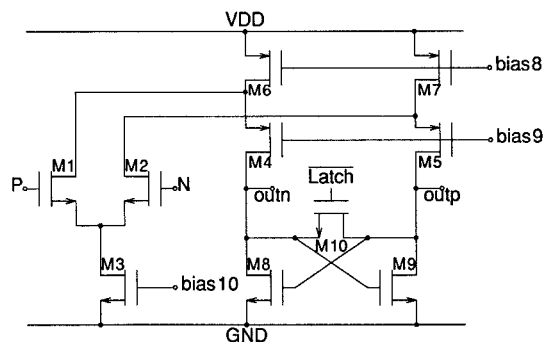
shows the comparator schematic. The circuit consists of a folded-cascode amplifier (M_1-M_7) in which the load has been replaced by a current-triggered latch (M_8-M_{10}). A folded-cascode architecture is chosen so that n-channel transistors can be used in both the differential pair and the latch. When M_{10} is on, the comparator outputs are connected together and to the gates of M_8 and M_9. In this configuration, the differential current flowing from the cascodes flows through M_{10}. When M_{10} turns off, the comparator outputs are separated, leaving M_8 and M_9 connected in a positive-feedback configuration. As a result, differential current coming from the cascodes charges the output-node parasitics and the output is latched; that is, one output is pulled up to the positive supply and the other down toward ground. Because the latch is current triggered, a resistive load is not required. This reduces the output parasitic capacitance by that amount associated with a resistive load and, as a result, speeds up the comparator. According to simulation, the comparator requires about 8 ns to set up and about 4 ns to latch with a 1-mV differential input.

V. EXPERIMENTAL RESULTS

Fig. 10 shows the dc linearity of the ADC at a conversion rate of 20 Msamples/s. In Fig. 10(a), differential nonlinearity (DNL) is plotted versus code, and in Fig.
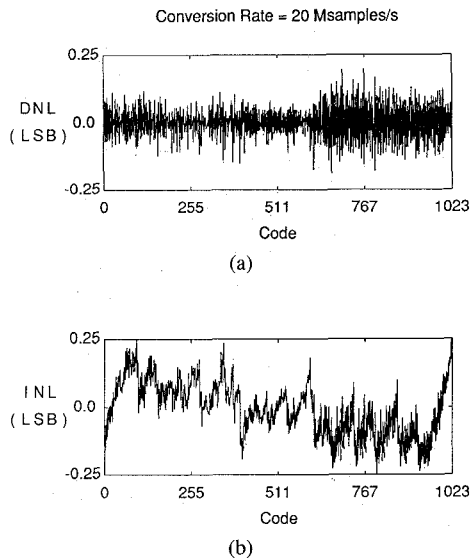
Conversion Rate = 20 Msamples/s

(a)

(b)

Fig. 10. (a) Differential nonlinearity (DNL) versus code. (b) Integral non-linearity (INL) versus code.

10(b), integral nonlinearity (INL) versus code is plotted. The magnitudes of the maximum DNL and INL are less than 0.2 and 0.25 LSB, respectively. The nonlinearities here are about 4 times smaller than those shown in [18]. This improvement is caused by the increase in the op-amp open-loop gain described in Section IV and affects other performance measures as well, as shown next.

Fig. 11 shows the output of a fast Fourier transform (FFT) on a block of 4096 consecutive codes. The conversion rate is 20 Msamples/s, and the input is a full-scale sine wave at 4.97 MHz. The SNR is about 60 dB; the SDR is about 74 dB and the SNDR about 60 dB. The SDR is about 14 dB higher than in [18]. Because the SNR is still limited by quantization noise, however, the SNDR here is also limited by the quantization noise. Next, FFT outputs such as in Fig. 11 have been used to generate plots of SNDR versus input level, input frequency, and conversion rate.

Fig. 12 shows plots of SNDR versus input level for two input frequencies: 100 kHz and 19.9 MHz. The 19.9-MHz input beats with the 20-Msamples/s conversion rate to form a 100-kHz ADC output. An ideal 10-b curve is also shown. Both data curves are close to ideal except at high input signal levels. For the 100-kHz input, the peak SNDR is about 60 dB instead of 62 dB in the ideal case. For the 19.9-MHz input, the peak SNDR is about 55 dB. This difference is caused by distortion generated by the input SHA.

Fig. 13 shows a plot of SNDR versus input frequency. The conversion rate is constant at 20 Msamples/s. The SNDR is down by 3 dB when the input frequency is 15 MHz. Eventually, the SNDR decreases at 6 dB per octave because it is limited by jitter in the sampling instant. This performance is much better than for ADC's without SHA's.

Fig. 14 shows two plots of the SNDR versus the conversion rate for different values of I_b, which is the pri-

Conversion Rate = 20 Msamples/s
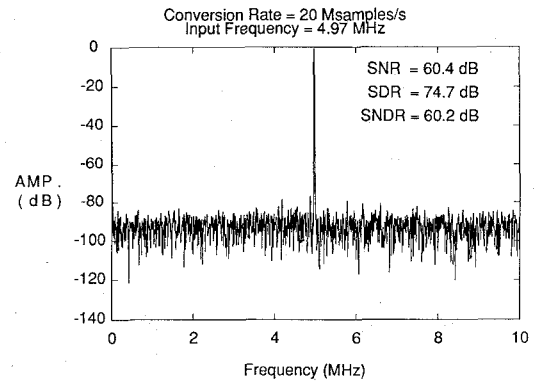Input Frequency = 4.97 MHz

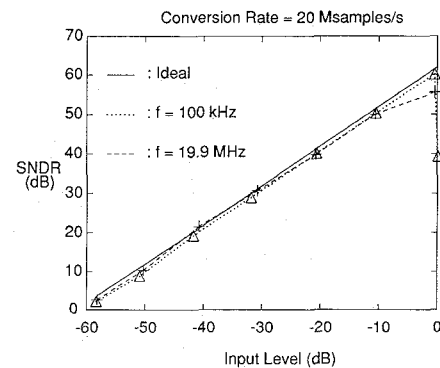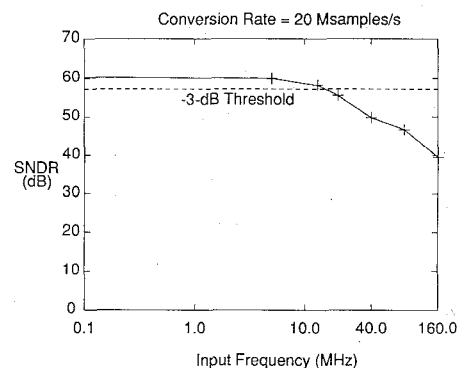Fig. 11. Fast-Fourier-transform output.

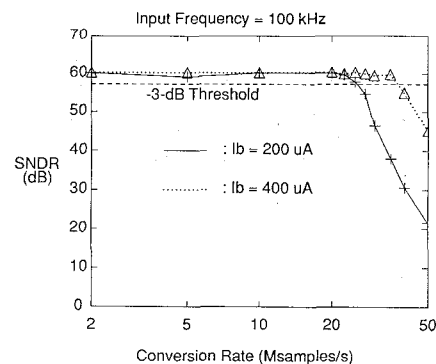Fig. 12. SNDR versus input level.

Fig. 13. SNDR versus input frequency.

Fig. 14. SNDR versus conversion rate.

Fig. 15. Die photograph.

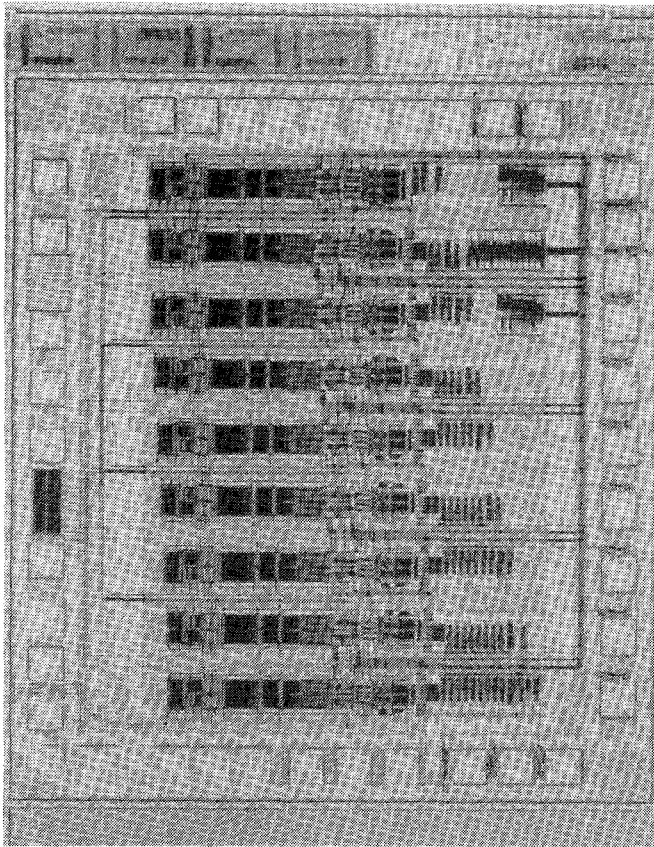TABLE I
ADC PERFORMANCE: +5 V AND 25°C

Technology	0.9-μm CMOS
Resolution	10 b
Conversion Rate	20 Msample/s
Area	8.7 mm^2
Power Dissipation	240 mW
Input Offset	6 LSB
DNL	0.20 LSB
INL	0.25 LSB
SNDR (f_{in} = 5 MHz)	60 dB
Differential Phase	0.2° p-p
Differential Gain	0.2% p-p
PSRR (5 kHz)	60 dB
CMRR (5 MHz)	55 dB

mordial bias current that sets the value of all other constant current sources in the ADC. The nominal value of I_b is 200 μA and sets the analog power dissipation to about 200 mW. The input frequency is constant at 100 kHz. When I_b = 200 μA, the solid plot results. The SNDR is down by 3 dB at 25 Msamples/s. This limit is set by the 20 ns required to settle the internal SHA's. When the bias current is doubled, the dotted plot results, and the -3-dB conversion rate becomes 38 Msamples/s, corresponding to a settling time of less than 13 ns.

Fig. 15 shows a photograph of the chip. It is about 2.8 mm by about 3.1 mm (110 mils by 120 mils). The stages follow one after another and are identical except that the ninth stage does not have a DAC or a subtractor and the two-phase nonoverlapping clock alternates from stage to

stage. The layout has been autorouted except in the analog blocks. This is important because the structure is modular and autorouting should allow the rapid implementation of a family of pipelined ADC's.

VI. SUMMARY

This paper describes a pipelined ADC, with typical characteristics summarized in Table I, and shows that pipelined conversion architectures can now be used in CMOS technologies for video-rate ADC applications.

ACKNOWLEDGMENT

REFERENCES

[1] J. G. Peterson, "A monolithic video A/D converter," IEEE J. Solid-State Circuits, vol. SC-14, pp. 932–937, Dec. 1979.
[2] B. Zojer, B. Astegher, H. Jessner, and R. Petschacher, "A 10 b 75 MHz subranging ADC," in ISSCC Dig. Tech. Papers, Feb. 1990, pp. 164–165.
[3] J. Marsh, K. Lofstrom, H. J. Engert, and B. Price, "A 10 bit 75 mega-sample per second A/D converter," in Proc. Custom Integrated Circuits Conf., 1991, pp. 26.6.1–26.6.4
[4] P. Real, D. H. Robertson, C. W. Mangelsdorf, and T. L. Tewksbury, "A wide-band 10-b 20-Ms/s pipelined ADC using current-mode signals," IEEE J. Solid-State Circuits, vol. 26, pp. 1103–1109, Aug. 1991.
[5] A. Matsuzawa et al., "A 10 b 30 MHz two-step parallel BiCMOS ADC with internal S/H," in ISSCC Dig. Tech. Papers, Feb. 1990, pp. 162–163.
[6] B.-S. Song, S.-H. Lee, and M. F. Tompsett, "A 10-b 15-MHz CMOS recycling two-step A/D converter," IEEE J. Solid-State Circuits, vol. 25, pp. 1328–1338, Dec. 1990.
[7] K. Martin, "A high-speed, high accuracy pipeline A/D converter," in Proc. Asilomar Conf. Circuits Syst., Computers, 1981, pp. 489–492.
[8] G. C. Temes, "High-accuracy pipeline A/D converter configuration," Electron. Lett., pp. 762–763, Aug. 1985.
[9] S. H. Lewis and P. R. Gray, "A pipelined 5-Msamples/s 9-bit analog-to-digital converter," IEEE J. Solid-State Circuits, vol. SC-22, pp. 954–961, Dec. 1987.
[10] S. Sutarja and P. R. Gray, "A pipelined 13-bit, 250-ks/s, 5-V analog-to-digital converter," IEEE J. Solid-State Circuits, vol. 23, pp. 1316–1323, Dec. 1988.
[11] B.-S. Song, M. F. Tompsett, and K. R. Lakshmikumar, "A 12-bit 1-Msample/s capacitor error-averaging pipelined A/D converter," IEEE J. Solid-State Circuits, vol. 23, pp. 1324–1333, Dec. 1988.
[12] S. H. Lewis, "Video-rate analog-to-digital conversion using pipelined architectures," Univ. California, Berkeley, Memo. UCB/ERL M87/90, Nov. 1987.
[13] G. G. Gorbatenko, "High-performance parallel-serial analog to digital converter with error correction," in IEEE Nat. Conv. Rec. (New York), Mar. 1966.
[14] O. A. Horna, "A 150 Mbps A/D and D/A conversion system," Comstat Tech. Rev., vol. 2, no. 1, pp. 52–57, 1972.
[15] S. Taylor, "High speed analog-to-digital conversion in integrated circuits," Ph.D. dissertation, Univ. California, Berkeley, pp. 30–42, 1978.
[16] M. Snelgrove, private communication.
[17] G. Jusuf, "A 1-bit/cycle algorithmic analog-to-digital converter without high-precision comparators," Univ. California, Berkeley, Memo. UCB/ERL M90/69, Aug. 1990
[18] S. H. Lewis, H. S. Fetterman, G. F. Gross, Jr., R. Ramachandran, and T. R. Viswanathan, "A pipelined 9-stage video-rate analog-to-digital converter," in Proc. Custom Integrated Circuits Conf., 1991, pp. 26.4.1–26.4.4
[19] S. Sutarja, "High-resolution pipelined analog-to-digital conver-

sion," Univ. California, Berkeley,, Memo. UCB/ERL M88/27, May 1988, pp. 40–42.

[20] Y.-M. Lin, B. Kim, and P. R. Gray, "A 13-b 2.5-MHz self-calibrated pipelined A/D converter in 3-μm CMOS," *IEEE J. Solid-State Circuits*, vol. 26, pp. 628–636, Apr. 1991.

[21] J. H. Fischer, private communication.

[22] T. C. Choi et al., "High-frequency CMOS switched-capacitor filters for communications application," *IEEE J. Solid-State Circuits*, vol. SC-18, pp. 652–664, Dec. 1983.

[23] P. R. Gray, private communication.

[24] G. Nicollini, P. Confalonieri, and D. Senderowicz, "A fully differential sample-and-hold circuit for high-speed applications," *IEEE J. Solid-State Circuits*, vol. 24, pp. 1461–1465, Dec. 1989.

[25] D. Senderowicz, S. F. Dreyer, J. H. Huggins, C. F. Rahim, and C. A. Laber, "A family of differential NMOS analog circuits for a PCM codec filter chip," *IEEE J. Solid-State Circuits*, vol. SC-17, pp. 1014–1023, Dec. 1982.

[26] N. S. Sooch, "MOS cascade current mirror," U.S. Patent 4 550 284, Oct. 1985.

Stephen H. Lewis (S'86–M'87) received the B.S. degree from Rutgers University, New Brunswick, NJ, in 1979, the M.S. degree from Stanford University, Stanford, CA, in 1980, and the Ph.D. degree from the University of California at Berkeley in 1987, all in electrical engineering.

From 1980 to 1982 he was with Bell Laboratories in Whippany, NJ, where he was involved in circuit design for magnetic recording. In 1988 he rejoined Bell Laboratories in Reading, PA, where he concentrated on the design of analog-to-digital converters. In 1991 he joined the Department of Electrical Engineering and Computer Science, University of California, Davis. His research interests include data conversion and analog circuit design.
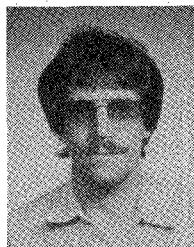
H. Scott Fetterman was born in Tamaqua, PA, in 1958. He received the B.S. degree in electrical engineering from the Pennsylvania State University, Middletown, PA, in 1983. He received the M.S. degree in electrical engineering from Lehigh University, Bethlehem, PA, in 1990.

When he joined AT&T Bell Laboratories, Allentown, PA, in 1983 he was involved in the design and testing of both CMOS and complimentary bipolar telecommunications integrated circuits. Since 1988 his work has concentrated mainly on high-speed CMOS A/D converters and data acquisition systems.
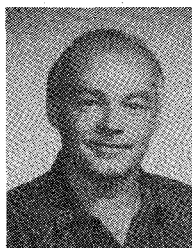
George F. Gross, Jr. received the Associate degree in electrical engineering from Pennsylvania State University in 1964.
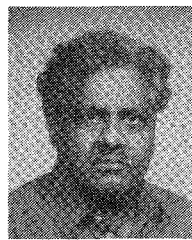
Upon graduation he joined AT&T Bell Laboratories, Allentown, PA, where he is presently employed as a Member of the Technical Staff. He has been involved in the design of traveling-wave tube amplifiers, cathode-ray tubes, high-power transistors, and most recently CMOS VLSI mixed signal circuits with concentration in the data converter area.

R. Ramachandran (S'75–M'82) was born in Madras, India, in 1952. He received the B.E. (Hons.) degree in electronics engineering from the Birla Institute of Technology and Science, Pilani, India, in 1974, the M.E. degree in electronics and communications engineering from the Indian Institute of Science, Bangalore, in 1976, and the Ph.D. degree in electrical and computer engineering from Clarkson University, Potsdam, NY, in 1981.

Since 1983 he has been with AT&T. During this time he has worked on several digital CMOS IC designs. He is currently a Member of the Technical Staff in the CMOS IC design group at AT&T Bell Laboratories, Allentown, PA. His research interests include fault-tolerant computing, finite state machine theory, digital system design, and BIST and DFT techniques.

Dr. Ramachandran is a member of Sigma Xi.

T. R. Viswanathan (S'62–M'64–SM'76) was Professor of Electrical Engineering at the Indian Institute of Technology, Kanpur, and the University of Waterloo, the University of Michigan, and Carnegie-Mellon University. He joined AT&T Bell Laboratories in 1985 as Supervisor of the CMOS analog IC design group. The main focus of this group is the design of IC's for data conversion and communication. He is also an Adjunct Professor at the Moore School of Electrical Engineering at the University of Pennsylvania, Philadelphia.