Copyright © 1987, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

A PIPELINED 5 Ms/s 9-BIT ANALOG TO DIGITAL CONVERTER

by

Stephen H. Lewis and Paul R. Gray

Memorandum No. UCB/ERL M87/46

19 June 1987

COVER PAGE

A PIPELINED 5 Ms/s 9-BIT ANALOG TO DIGITAL CONVERTER

by

Stephen H. Lewis and Paul R. Gray

Memorandum No. UCB/ERL M87/46

19 June 1987

ELECTRONICS RESEARCH LABORATORY

College of Engineering
University of California, Berkeley
94720

TITLE PAGE

A PIPELINED 5 Ms/s 9-BIT ANALOG TO DIGITAL CONVERTER

by

Stephen H. Lewis and Paul R. Gray

Memorandum No. UCB/ERL M87/46

19 June 1987

ELECTRONICS RESEARCH LABORATORY

College of Engineering
University of California, Berkeley
94720

A Pipelined 5 Ms/s 9-bit Analog to Digital Converter*

Stephen H. Lewis
Paul R. Gray

Electronics Research Laboratory
University of California, Berkeley
Berkeley, California
(415) 642-9399 or (415) 642-5179

ABSTRACT

A pipelined, 5 Ms/s, 9-bit analog-to-digital converter with digital correction has been designed and fabricated in a 3-micron, CMOS technology. It requires 8500 square mils, consumes 180 mW, and has an input capacitance of 3 pF. A fully differential architecture is used; only a 2-phase, nonoverlapping clock is required, and an on-chip, sample-and-hold amplifier is included.

May 13, 1987

* Research sponsored by DARPA under contract number N00039-C-0107 and National Science Foundation under contract number DCI-8603430.

A Pipelined 5 Ms/s 9-bit Analog to Digital Converter*

Stephen H. Lewis

Paul R. Gray

Electronics Research Laboratory

University of California, Berkeley

Berkeley, California

(415) 642-9399 or (415) 642-5179

I. INTRODUCTION

Traditional designs of high-speed, CMOS analog-to-digital (A/D) converters have used flash architectures [1-13]. While flash architectures make the fastest converters, they tend to require large silicon areas because they need many comparators. An important objective is the realization of high-speed A/D converters in much less area than flash converters require so that these converters can be integrated on the same chips with complex, high-speed, image-processing functions. Multistage conversion architectures reduce the required area by reducing the number of comparators [14-19]. Using pipelining in these architectures allows the stages to operate concurrently and makes the maximum conversion rate independent of the number of stages. Also, digital correction reduces the sensitivity of the architecture to component nonidealities. In this paper, an experimental, 4-stage, pipelined A/D converter with digital correction that has 9-bit resolution and 5 Ms/s conversion rate in a 3-micron, CMOS technology is described.

This paper is divided into 4 additional parts. In section II, pipelined A/D architectures are described conceptually, and their advantages over flash and two-step subranging architectures [17] are explained. In section III, the error sources present in pipelined A/D converters are identified, and the

* Research sponsored by DARPA under contract number N00039-C-0107 and National Science Foundation under contract number DCI-8603430.

way in which digital correction eliminates the effects of some of these errors is shown. In section IV, the circuits in an experimental prototype are described. Finally, experimental results from the prototype converters are given in section V.

II. CONCEPTUAL DESCRIPTION

A block diagram of a general pipelined A/D converter with k stages is shown in Figure 1. Each stage contains a sample-and-hold (S/H) circuit, a low-resolution A/D subconverter, a low-resolution digital-to-analog (D/A) converter, and a differencing, fixed-gain amplifier. First, each stage samples and holds the output from the previous stage. Next, each stage does a low-resolution A/D conversion on the held input. Then, the code just produced is converted back into an analog voltage by a D/A converter. The D/A converter output is subtracted from the held input, producing a residual voltage that is amplified and sent to the next stage.

The primary, potential advantages of the pipelined architecture are high throughput rate and low hardware cost. The high throughput rate of the pipelined architecture stems from concurrent operation of the stages. At any time, the first stage operates on the most recent sample, while the next stage operates on the residual from the previous sample, and so forth. If the A/D subconversions are done with flash converters, a pipelined architecture only needs two clock cycles per conversion. Flash architectures also require two clock cycles per conversion, one each for sampling and A/D conversion, and use pipelining to do the digital encoding to binary operation. However, flash architectures make the fastest converters because their pipelined information is entirely digital and can be transferred to 1-bit accuracy in less time than it takes to generate and transfer the analog residual in a pipelined multistage architecture. The area and consequent manufacturing cost required for pipelined converters is small compared to flash converters because pipelined converters require fewer comparators than flash converters. For example, the 9-bit prototype pipelined converter described in section IV requires a core area of 8500 square mils in a 3-micron, CMOS technology. In the same technology, a 9-bit flash converter would be more than ten times larger than the pipelined prototype. Not only is the area small, but also the area of pipelined converters is linearly related to the resolution because, if the necessary accuracy can be achieved through calibration or trimming, the resolution can be increased by adding stages to the

end of the pipeline without increasing the number of clock cycles required per conversion. In contrast, if the necessary accuracy can be achieved with calibration or trimming, flash and subranging architectures need exponential, rather than linear, increases in area to increase their resolution.

The differences between subranging and pipelined architectures are caused by the presence or absence of operational amplifiers (opamps). The opamps in pipelined architectures are used to design dedicated S/H amplifiers that allow both concurrent stage operation and accurate high input frequency sampling. In contrast, subranging architectures do not use opamps or S/H amplifiers and thus have sequential stage operation and reduced accuracy sampling at high input signal frequencies. Flash converters also do not use a dedicated, on-chip S/H amplifier because of the difficulty in building an opamp in CMOS technologies that is fast enough to drive the inherently large input load; therefore, flash architectures suffer reduced performance at high input signal frequencies.

III. ERROR SOURCES

The primary error sources present in a pipelined A/D converter are offset errors in the S/H circuits and amplifiers, gain errors in the S/H circuits and amplifiers, A/D subconverter nonlinearity, D/A subconverter nonlinearity, and opamp settling time errors. With digital correction, as shown below, the effects of offset, gain, and A/D subconverter nonlinearity are reduced or eliminated; therefore, the D/A converter nonlinearity and opamp settling time errors limit the performance of pipelined A/D converters. To begin the error analysis, the effects of offset and gain errors are considered next.

A block diagram of a 2-stage, pipelined A/D converter with offset and gain errors in each of the S/H circuits and interstage amplifier is shown as a representative example in Figure 2. The nonideal S/H circuits and interstage amplifier are replaced by ideal elements in series with gain and offset errors, and each of these replacements is surrounded by a dotted line. The gain error in the first-stage S/H circuit changes the conversion range of the A/D converter and does not affect linearity. The gain errors in the interstage amplifier and second-stage S/H circuit can be combined into one equivalent error that does affect linearity. However, because the interstage gain only has to be accurate enough to preserve the linearity of the stages after the first stage, the effect of this gain error on linearity is small. For example, if both stages in Figure 2 have 4-bit resolution, the interstage amplifier gain should be equal to 16

and must be accurate to within ±3%.

The offset error in the first-stage S/H circuit causes an input-referred offset but does not affect linearity. The offset errors in the interstage amplifier and second-stage S/H circuit can be combined into one equivalent offset that does not affect linearity if digital correction is used. Because addition is commutative, the equivalent offset can be pushed to the left of the first-stage subtractor. To move the equivalent offset to the input branch, where it causes an input-referred offset, an equal but opposite offset must be inserted in the first-stage A/D subconverter branch. As shown below, the effect of the offset in the first-stage A/D subconverter is eliminated by the digital correction.

Next, the effect of nonlinearity in the first-stage A/D subconversion is considered. A block diagram of one stage in a pipelined A/D converter is shown in Figure 3a. A 2-bit stage is used as a representative example. Nonlinearity in the A/D subconverter is modeled as an input-referred linearity error. The effect of this nonlinearity is studied by examining plots of residual voltage vs. input voltage. Two such plots are shown in Figures 3b and 3c.

In Figure 3b, both the A/D subconverter and the D/A converter are assumed to be ideal. The plot has a sawtooth shape because when the input is between the decision levels determined by the A/D subconverter, the A/D subconverter and D/A converter outputs are constant; therefore, the residual rises with the input. When the input crosses a decision level, the A/D subconverter and D/A converter outputs increase by one least significant bit (LSB) at a 2-bit level, so the residual decreases by 1 LSB. Here, the residual is always between ±1/2 LSB in magnitude and consists only of the part of the input that is not quantized by the first stage. With the interstage gain equal to 4, the maximum residual is amplified into a full-scale input to the next stage; therefore, the conversion range of the next stage is equal to the maximum residual out of the first stage.

A similar curve is shown in Figure 3c for a case when the A/D subconverter has some nonlinearity, but the D/A converter is still ideal. In this example, two of the A/D subconverter decision levels are shifted, one by -1/2 LSB and the other by +1/2 LSB. When the input crosses a shifted decision level, the residual decreases by one LSB. If the decision levels are shifted by less than 1/2 LSB, the residual is always between ±1 LSB in magnitude. Here, the residual voltage consists of both the

unquantized part of the input and the error caused by the A/D subconverter nonlinearity. Because the D/A converter is assumed to be ideal, these increased residuals are accurate for the codes to which they correspond; therefore, at this point, no information is lost. If the interstage gain is still 4, however, information is lost when the larger residuals saturate the next stage and produce missing codes in the conversion. Therefore, if the conversion range of the second stage is increased to handle the larger residuals, they can be encoded and the errors corrected. This process is called digital correction [20-21] and is described next.

A block diagram of a 2-stage pipelined A/D converter with digital correction is shown in Figure 4. The new elements in this diagram are the pipelined latches, the digital correction logic circuit, and the amplifier with a gain of 0.5. The amplifier with a gain of 0.5 is conceptual only and is drawn to show that the interstage gain is reduced by a factor of 2 so that nonlinearity error less than ±1/2 LSB at a **n1**-bit level in the first-stage A/D subconversion does not produce residuals that saturate the second stage. If the first stage is perfectly linear, only half the conversion range of the second stage is used. Therefore, 1 bit from the second stage is saved to digitally correct the outputs from the first stage; the other **n2-1** bits from the second stage are added to the overall resolution. After the pipelined latches align the outputs in time so that they correspond to one input, the digital correction block detects overrange in the outputs of the second stage and changes the output of the first stage by 1 LSB at a **n1**-bit level if overrange occurs. Digital correction improves linearity by allowing the converter to postpone decisions on inputs that are near the first-stage A/D subconverter decision levels until the residuals from these inputs are amplified to the point where similar nonlinearity in later-stage A/D subconverters is insignificant.

To do the digital correction, a correction logic circuit is required. Also, if flash converters are used in the stages, all stages after the first require twice as many comparators as without digital correction. The logic is simple, however, and the comparators no longer need to be offset cancelled.

It is shown above that, with digital correction, if the D/A converter is ideal, nonlinearity in the A/D subconverters can be corrected. Therefore, the D/A converter in the first stage determines the linearity of the entire A/D converter. For about 9-bit integral linearity, such a D/A converter can be

made with a resistor string. For integral linearity greater than 9 bits, the design of such a D/A converter is not trivial and probably requires calibration. Also, fast settling opamps are required to do analog subtraction and amplification at the sampling rate of the A/D converter. The 3-micron, CMOS prototype described in section IV is able to do these functions at 5 Ms/s. The maximum speed of such processing increases in scaled technologies, and video conversion rates should be achievable in 1.5-2-micron CMOS technologies.

IV. PROTOTYPE

The high-level design problems and solutions for the prototype, pipelined A/D converter are now considered. First because the goal of this project is to make the A/D converter small enough that it can be incorporated within a primarily digital chip, the A/D converter must be able to operate in the presence of large power supply noise caused by the digital circuits. To reduce the sensitivity of the converter to this noise, all analog signal paths in the prototype are fully differential. Second, to operate at as high a speed as possible, fast opamps, flash subconverters, and a small resolution per stage should be used. The resolution per stage should be small so that the interstage gain is small and the corresponding bandwidth is large. Third, for about 9-bit integral linearity, digital correction, resistor string D/A converters, and a large resolution per stage should be used. The resolution per stage should be large so that interstage amplifier gain error has little effect on the overall linearity. Note that the speed and linearity requirements conflict in determining the optimum resolution per stage. Finally, to keep the area small, a compact A/D, D/A subconverter structure, and a medium resolution per stage are used; this compromise in the resolution per stage keeps both the number of opamps and the number of comparators small.

To meet these requirements, the prototype is divided into 4 stages with 3 bits produced per stage. A block diagram of one stage is shown in Figure 5. The A/D subconversions are done with flash converters, so each stage needs 7 comparators. The S/H amplifier block replaces both the S/H circuit and interstage gain shown in earlier figures. Because the interstage gain is 4 instead of 8, half the range and one bit from each of the last three stages are saved to digitally correct the outputs of the previous stages. Thus, instead of obtaining 3 bits of resolution from each of these stages, only 2 bits of

resolution are obtained from each. In total, 9 bits of resolution are produced, using 28 comparators and 4 opamps.

The S/H amplifier block is expanded in Figure 6a. Figure 6b shows that the clock is divided into two nonoverlapping phases. On clock phase ϕ_1, the input is sampled onto the $4C_I$ capacitors, and the integrating (C_I) and common-mode feedback (C_{CM}) capacitors are reset. On ϕ_2, the two inputs are connected together so the difference between the two inputs is amplified by the ratio of the sampling to integrating capacitors. To the extent that the opamp in a closed-loop configuration drives its differential input to zero, the gain is insensitive to parasitic capacitances on either the top or bottom plates of any of these capacitors. Meanwhile, the common-mode feedback (CMFB) capacitors are connected to the outputs of the opamp to start the CMFB circuit. Switched-capacitor CMFB is useful in pipelined A/D converters because pipelined converters inherently allow a clock phase needed to reset the capacitor bias.

Because of digital correction, the offsets of all the opamps are referred to the input of the A/D converter each in an amount diminished by the combined interstage amplifier gain preceding the offset; therefore, the opamps do not have to be offset cancelled and do not have to be placed in a unity-gain feedback configuration. Since the opamps do not have to be unity-gain stable, their speed can be optimized for a closed loop gain of 4. The opamp, shown in Figure 7, uses a fully differential, class A/B configuration with dynamic bias. The class A/B structure gives both high slew rate and high gain after slewing.

The opamp is similar to one reported by Castello and Gray [22], and its operation is now described. Transistors M1-M4 form the input stage and generate the class A/B action. Source followers M5-M8 are used to bias the input stage so that it conducts some current even for zero differential input. For an increase in the voltage on the positive input and a corresponding decrease on the negative input, the gate-to-source voltages of both M1 and M4 increase while those of M2 and M3 decrease; therefore, the current in M1 and M4 increases and that in M2 and M3 decreases from their standby values. Transistors M9 and M13, M10 and M14, M11 and M15, and M12 and M16 form current mirrors that reflect and amplify current from the input branches to the output branches.

Cascode transistors M17-M20 increase the gain of the opamp by increasing the output resistance of the output nodes to ground. A high-swing, dynamic bias circuit composed of transistors M31-M38 adjusts the gate bias on the cascode transistors so that the output branches can conduct large currents during slewing and have high swings during settling. Transistors M41-M44, together with the C_{CM} capacitors and associated switches in Figure 6a, form the CMFB circuit. Because the gates of M41 and M42 are tied to a constant bias voltage, these transistors are constant-current sources. The gates of M43 and M44 are connected to the CMBIAS terminal shown in Figure 6a. This point is alternatively switched from a bias voltage on ϕ_1 to a capacitively coupled version of the output on ϕ_2. During ϕ_2, the CMBIAS point rises and falls with changes in the common-mode output voltage. This change adjusts the current drawn through M43 and M44 so that the common-mode output voltage is held constant near zero volts. Note that if the two halves of the differential circuit match perfectly, changes in the differential output voltage do not change the CMBIAS point.

Because the speed of this opamp is limited by the speed of its current mirrors, wide-band current mirrors are used to increase the speed. To this end, transistors M9-M12 are not simply diode connected, but instead are buffered by source followers MB1-MB4. Because of this change, the currents needed to supply the capacitive short circuit between the gates and sources of the current mirrors at high frequencies come from the power supplies instead of from the input branch. The drawback to this approach is that the drain-to-source voltages of transistors M9-M12 are increased by the gate-to-source voltages of transistors MB1-MB4, respectively. Therefore, input stage transistors M1-M4 operate with less drain-to-source voltage than if M9-M12 were diode connected. As a result, M1-M4 enter the triode region for smaller differential inputs than with diode-connected loads, and the amount of current that the input stage can produce while slewing is limited. Because a high-swing, dynamic bias circuit is used, this is not a problem for ±5 V operation; however, for +5 V operation, these wide-band current mirrors probably would limit the slew rate of the opamp.

A block diagram of an A/D, D/A subsection is shown in Figure 8. To save area, one resistor string is shared for both the A/D and D/A functions. The resistor string divides the reference into equal segments and provides the boundaries between these segments as thresholds for a bank of comparators.

The comparators are clocked at the end of ϕ_2. On ϕ_1, eight D/A converter outputs are enabled and one is selected based on control signals generated from the comparator outputs (y_1, ... , y_8). Although Figure 8 shows a single-ended representation of both the A/D subconverter and D/A converter functions, on the prototype, both functions are fully differential. Therefore, instead of just one D/A converter output, equal and opposite D/A converter outputs are used. Also, each comparator compares a differential input to a differential reference instead of a single-ended input to a single-ended reference.

The connection of the comparator within an A/D, D/A subsection is shown in Figure 9. The points labeled VR+ and VR- are connected to taps on the resistor string that depend on which comparator in the bank is under consideration. For example, for the top comparator, VR+ is connected to the most positive A/D subconverter tap, and VR- is connected to the most negative A/D subconverter tap. On clock ϕ_1, the comparator inputs are grounded, and the capacitors sample the differential reference. On ϕ_2, the left sides of the capacitors are connected to the differential input. Ignoring parasitic capacitance, the input to the comparator is then the difference between the differential input and the differential reference. The parasitic capacitances on the inputs to the comparator attenuate the input slightly, but the decision is not affected if the comparator has enough gain. As mentioned in section III, because of digital correction, no offset cancellation on the comparator is required. Therefore, the comparator is never placed in a feedback loop and does not have to be stable in a closed-loop configuration.

The comparator, shown in Figure 10, uses a conventional, latched-differential-amplifier configuration. Transistors M1 and M2 are source followers. Transistors M3-M8 form a differential amplifier, and ML1 and ML2 form a latch. Transistors MCS1 and MCS2 form a current switch that allows the bias current from MB2 to flow through either the differential amplifier or the latch. With the latch signal low, the inputs are amplified. Because M7 and M8 are biased in the triode region, the gain of the amplifier is only about 20 dB. When the latch signal is raised, the bias current is switched from the amplifier to the latch. During the transition, the parasitic capacitances on the inputs to the latch hold the amplified input. Finally, the latch switches, and the comparison is completed.

V. EXPERIMENTAL RESULTS

The prototype has been tested in two ways [23-24]: first with a code density test, and second with a signal-to-noise ratio (SNR) test. Both tests have used high and low frequency input signals. Results of the code density test are considered first.

In Figure 11, differential nonlinearity (DNL) is ploued on the y-axis vs. code on the x-axis for all 512 codes. The conversion rate is 5 Ms/s, and the input frequency is 2 MHz. Because the DNL never goes down to -1 LSB, there are no missing codes. The maximum DNL is less than 0.6 LSB.

In Figure 12, integral nonlinearity (INL) is plotted on the y-axis vs. code on the x-axis. Again, the conversion rate is 5 Ms/s and the input frequency is 2 MHz. The maximum INL is 1.1 LSB. The nonideality in the curve is caused by both nonlinearity in the first-stage D/A converter and incomplete settling of the first-stage opamp output.

SNR measurements were made by taking fast Fourier transforms on 1024 samples from the A/D converter at the down-sampled rate of 20 kHz while the converter was running at 5 Ms/s. In Figure 13, SNR is plotted on the y-axis vs. input level on the x-axis for five input frequencies: 2 kHz, 22 kHz, 202 kHz, 2.002 MHz, and 5.002 MHz. The curve for 5.002 MHz represents a beat frequency test on the converter when compared to the curve for 2 kHz because the converter is running at the difference between these two frequencies or 5 Ms/s. An ideal 9-bit curve is also shown. The peak SNR is around 50 dB instead of 56 dB, as would be expected with a 9-bit converter; this difference is accounted for by distortion generated from the INL for large input signals. When the input signal is reduced in amplitude, the distortion is reduced and the real curves approach the ideal 9-bit curve. Note that there is little difference in the curves for different input frequencies, showing that the first-stage S/H amplifier is able to accurately sample high-frequency input signals.

The results of the code density and SNR tests for variations in the input frequency are summarized in Table 1.

Table 1

Data Summary over Input Frequency Variation

9-bit Resolution
5 Ms/s Conversion Rate
±5 V Power Supplies

Input Frequency	2 kHz	2 MHz	5.002 MHz
Peak DNL (LSB)	0.5	0.6	0.5
Peak INL (LSB)	1.0	1.1	1.2
Peak SNR (dB)	50	50	49

Peak DNL, INL and SNR are shown for three input frequencies, and the performance is almost constant. This is important because it shows that the first-stage S/H amplifier is able to accurately sample high-frequency input signals.

A photograph of the core of a prototype chip is shown in Figure 14. The core is about 50 mils high by 150 mils wide. The stages follow one after another and are identical except that the fourth stage does not have a D/A converter or a subtractor and the 2-phase, nonoverlapping clock alternates from stage to stage. A test opamp and a test comparator are at the end. The prototype was made by MOSIS in a 3-micron, double-polysilicon, p-well, CMOS process.

SUMMARY

Pipelining is not a new idea, but had not been applied to monolithic CMOS A/D conversion because of the difficulty in building fast enough switched-capacitor amplifiers. This paper reports on a prototype pipelined A/D converter with typical characteristics summarized in Table 2.

Table 2

Typical Performance: 25°C

Technology	3-u CMOS
Resolution	9 bits
Conversion Rate	5 Ms/s
Area*	8500 $mils^2$
Power Supplies	±5 V
Power Dissipation	180 mW
Input Capacitance	3 pF
Input Offset	< 1 LSB
CM Input Range	±5 V
DC PSRR	50 dB

*Does not include clock generator, bias generator, reference generator, digital error correction logic, and pads

In summary, the prototype demonstrates that pipelined architectures and digital correction techniques are of potential interest for high-speed CMOS A/D conversion applications.

ACKNOWLEDGEMENT

We gratefully acknowledge the help of R. Kavaler and J. Doernberg with the testing of the prototype converters,

REFERENCES

1. A. G. F. Dingwall, "Monolithic Expandable 6 Bit 20 MHz CMOS/SOS A/D Converter", IEEE Journal of Solid-State Circuits, pp. 926-932, December, 1979.

2. J. G. Peterson, "A Monolithic Video A/D Converter, "IEEE Journal of Solid-State Circuits, pp. 932-937, December 1979.

3. T. Takemoto, et al, "A Fully Parallel 10-bit A/D Converter with Video Speed", IEEE Journal of Solid-State Circuits, pp. 1133-1138, December, 1982.

4. Y. Fujita, et al, "A Bulk CMOS 20MS/s 7b Flash ADC", Digest of Technical Papers, 1984 International Solid-State Circuits Conference, San Francisco, CA, pp. 56-57, February, 1984.

5. A. G. F. Dingwall and V. Zazzu, "High Speed CMOS A/D and D/A Conversion", Digest of Technical Papers, 1984 IEEE International Symposium on Circuits and Systems, pp. 420-424, May, 1984.

6. M. Inoue, et al, "A Monolithic 8-Bit A/D Converter with 120 MHz Conversion Rate", IEEE Journal of Solid-State Circuits, pp. 837-841, December, 1984.

7. T. Tsukada, et al, "CMOS 8b 25MHz Flash ADC", Digest of Technical Papers, 1985 International Solid-State Circuits Conference, New York, N. Y., pp 34-35, February, 1985.

8. A. K. Joy, et al, "An Inherently Monotonic 7-Bit CMOS ADC for Video Applications", IEEE Journal of Solid-State Circuits, pp. 436-440, June 1986.

9. T. Kumamoto, et al, "An 8-bit High-Speed CMOS A/D Converter", IEEE Journal of Solid-State Circuits, pp. 976-982, December, 1986.

10. B. Peetz, et al, "An 8-bit 250 Megasample per Second Analog-to-Digital Converter: Operation Without a Sample and Hold", IEEE Journal of Solid-State Circuits, pp. 997-1002, December, 1986.

11. Y. Yoshii, et al, "An 8b 350MHz Flash ADC", Digest of Technical Papers, 1987 IEEE International Solid-State Circuits Conference, New York, N. Y., pp. 96-97, February, 1987.

12. Y. Akazawa, et al, "A 400MSPS 8b Flash AD Conversion LSI", Digest of Technical Papers, 1987 IEEE International Solid-State Circuits Conference, New York, N. Y., pp. 98-99. February, 1987.

13. J. Corcoran, et al, "A 1GHz 6b ADC System", Digest of Technical Papers, 1987 IEEE International Solid-State Circuits Conference, New York, N. Y., pp. 102-103, February, 1987.

14. R. J. van de Plassche, "A High-Speed 7 bit A/D Converter", IEEE Journal of Solid-State Circuits, pp. 938-943, December, 1979.

15. R. A. Blauschild, "An 8b 50ns Monolithic A/D Converter with Internal S/H", Digest of Technical Papers, 1983 IEEE International Solid-State Circuits Conference, New York, N. Y., pp. 178-179, February, 1983.

16. R. E. J. van de Grift and R. J. van de Plassche, "A Monolithic 8-Bit Video A/D Converter", IEEE Journal of Solid-State Circuits, pp. 374-378, June, 1984.

17. A. G. F. Dingwall and V. Zazzu, "An 8MHz CMOS Subranging 8bit A/D Converter", IEEE Journal of Solid-State Circuits, pp. 1138-1143, December, 1985.

18. R. E. J. van de Grift and M. van der Veen, "An 8b 50MHz ADC with Folding and Interpolation Techniques", Digest of Technical Papers, 1987 IEEE International Solid-State Circuits Conference, New York, N. Y., pp. 94-95, February, 1987.

19. S. H. Lewiş and P. R. Gray, " A Pipelined 5MHz 9b ADC", Digest of Technical Papers, 1987 IEEE International Solid-State Circuits Conferences, New York, N. Y., pp. 210-211, February 1987.

20. O. A. Horna, "A 150 Mbps A/D and D/A Conversion System", Comsat Technical Review, Vol. 2, No. 1, pp. 52-57, 1972.

21. S. Taylor, "High Speed Analog-to-Digital Conversion in Integrated Circuits", Ph. D. Thesis, U. C. Berkeley, pp. 30-42, 1978.

22. R. Castello and P. R. Gray, "A High-Performance Micropower Switched-Capacitor Filter", IEEE Journal of Solid-State Circuits, pp. 1122-1132, December 1985.

23. W. A. Kester, "Characterizing and Testing A/D and D/A Converters for Color Video Applications", IEEE Transactions on Circuits and Systems, pp. 539-550, July, 1978.

24. J. Doernberg, H. Lee, D. A. Hodges, "Full-Speed Testing of A/D Converters", IEEE Journal of Solid-State Circuits, pp. 820-827, December 1984.

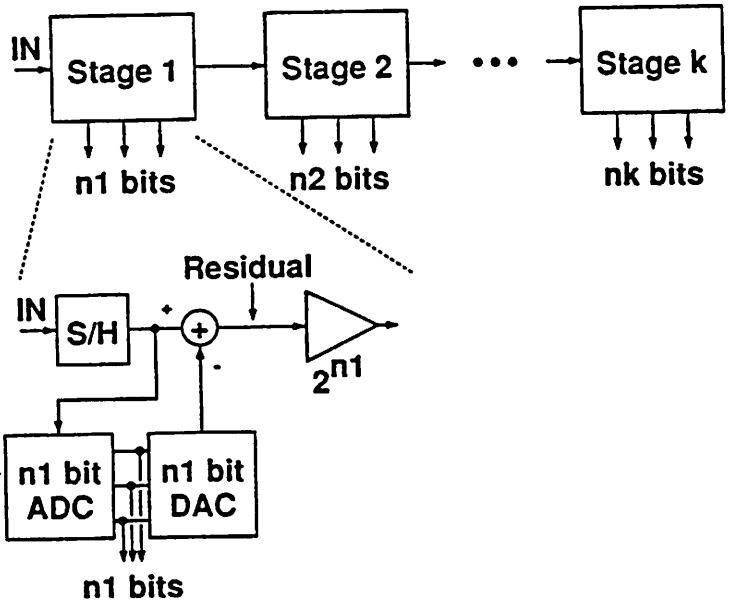

Figure 1 - Block diagram of a general pipelined A/D converter

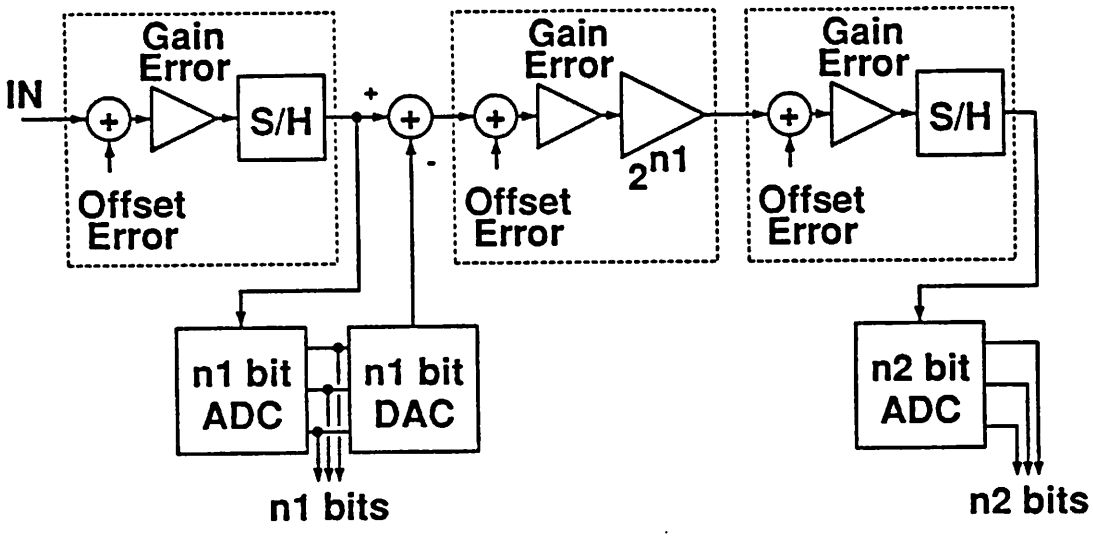

Figure 2 - Block diagram of a 2-stage pipelined A/D converter with offset and gain errors

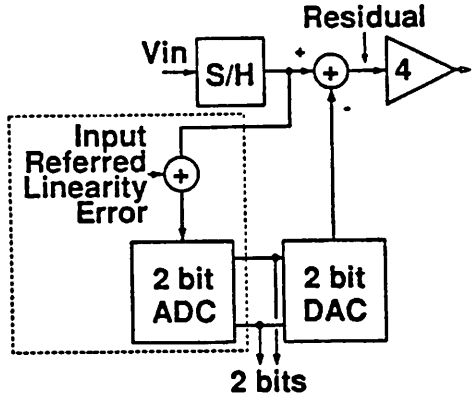

a. Block diagram of one 2-bit stage in a pipelined A/D converter

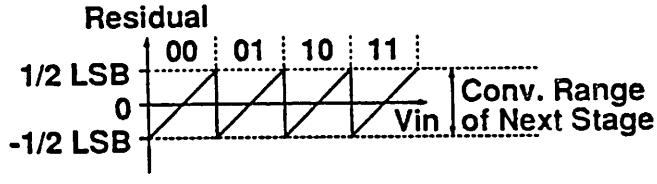

b. Ideal residual vs. input

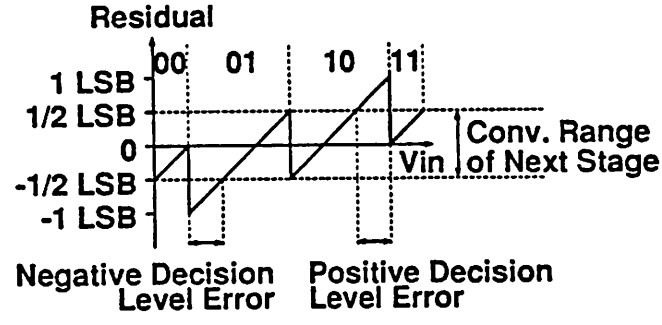

c. Residual vs. input with A/D subconverter nonlinearity

Figure 3

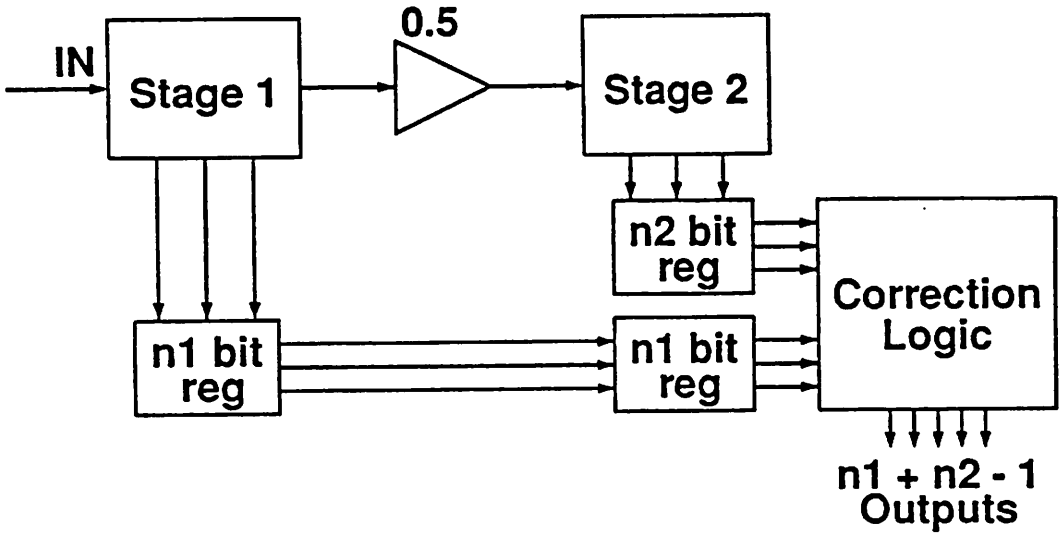

Figure 4 - Block diagram of a 2-stage pipelined A/D converter with digital correction

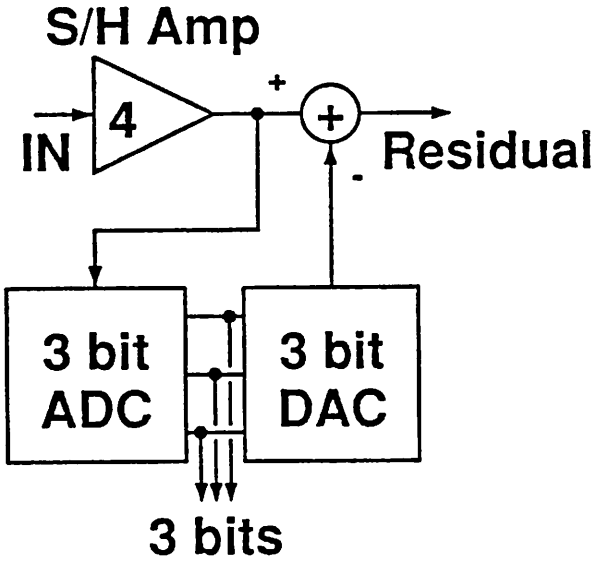

Figure 5 - Block diagram of one stage in the prototype

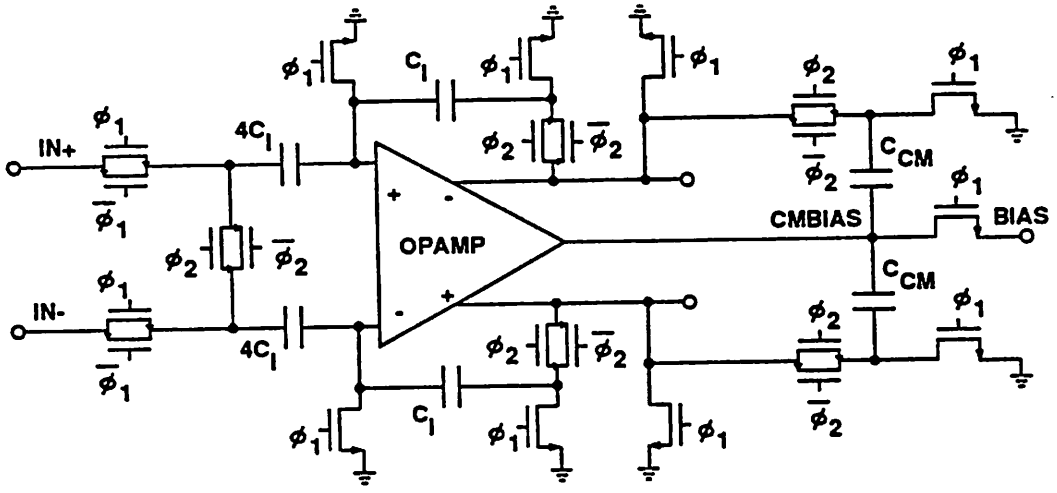

a. Schematic of S/H amplifier

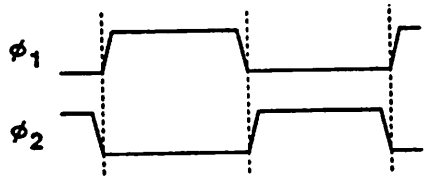

b. Timing diagram of a 2-phase nonoverlapping clock

Figure 6

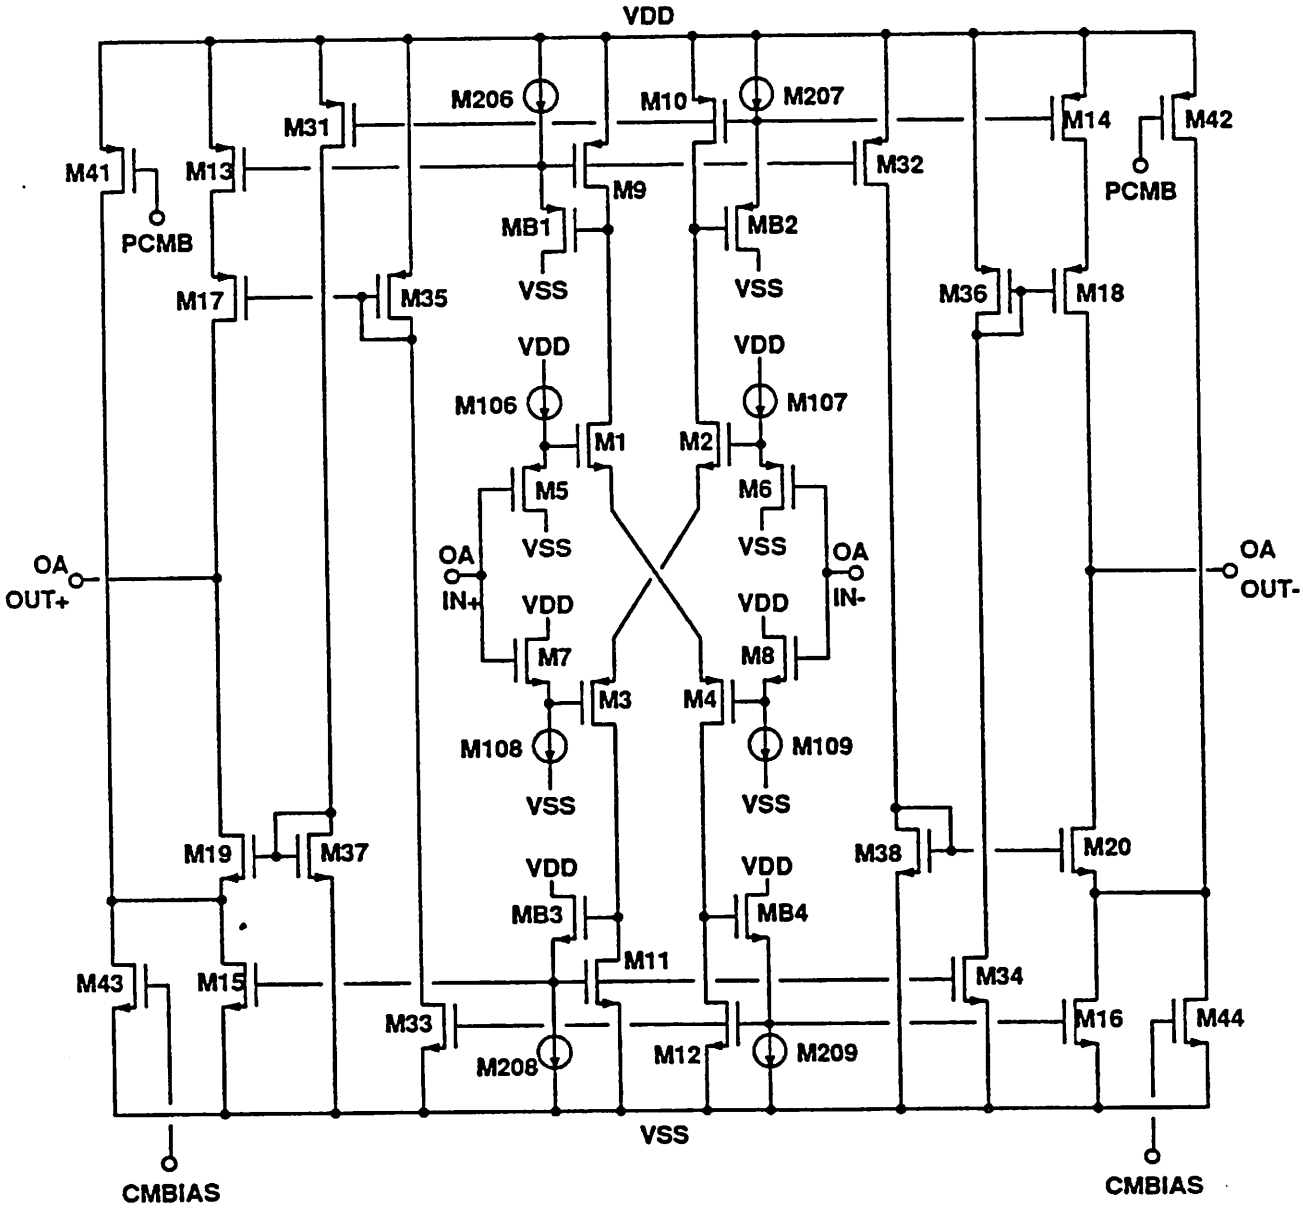

Figure 7 - Opamp schematic

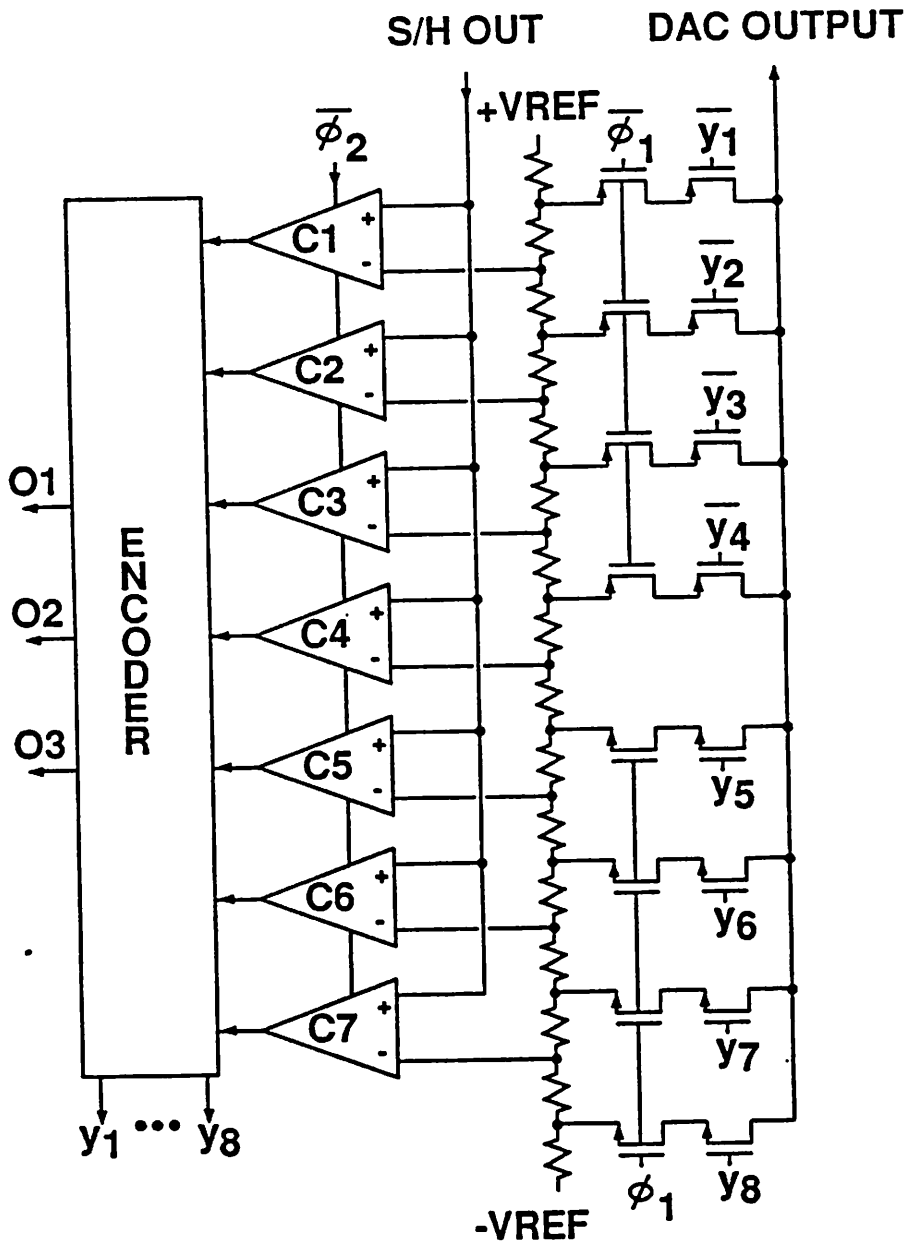

Figure 8 - Block diagram of A/D, D/A subsection

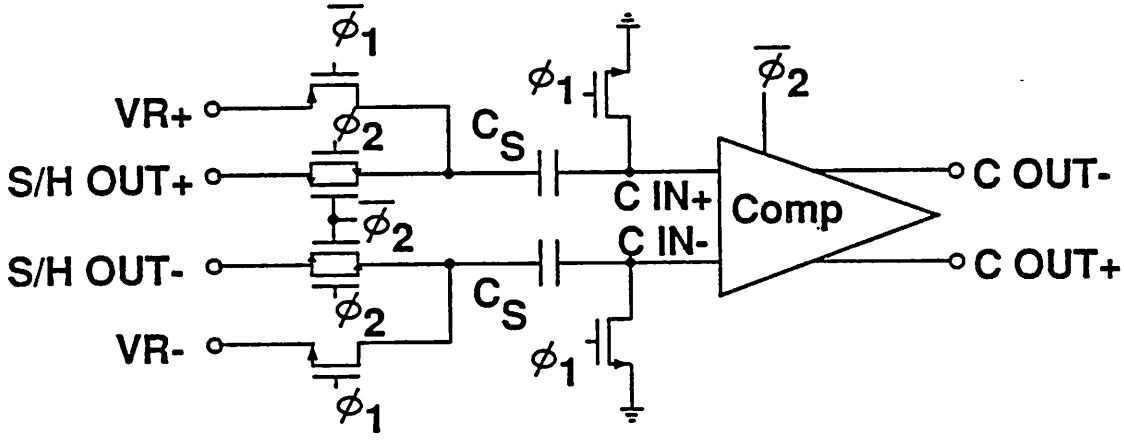

Figure 9 - Connection of comparator within A/D, D/A subsection

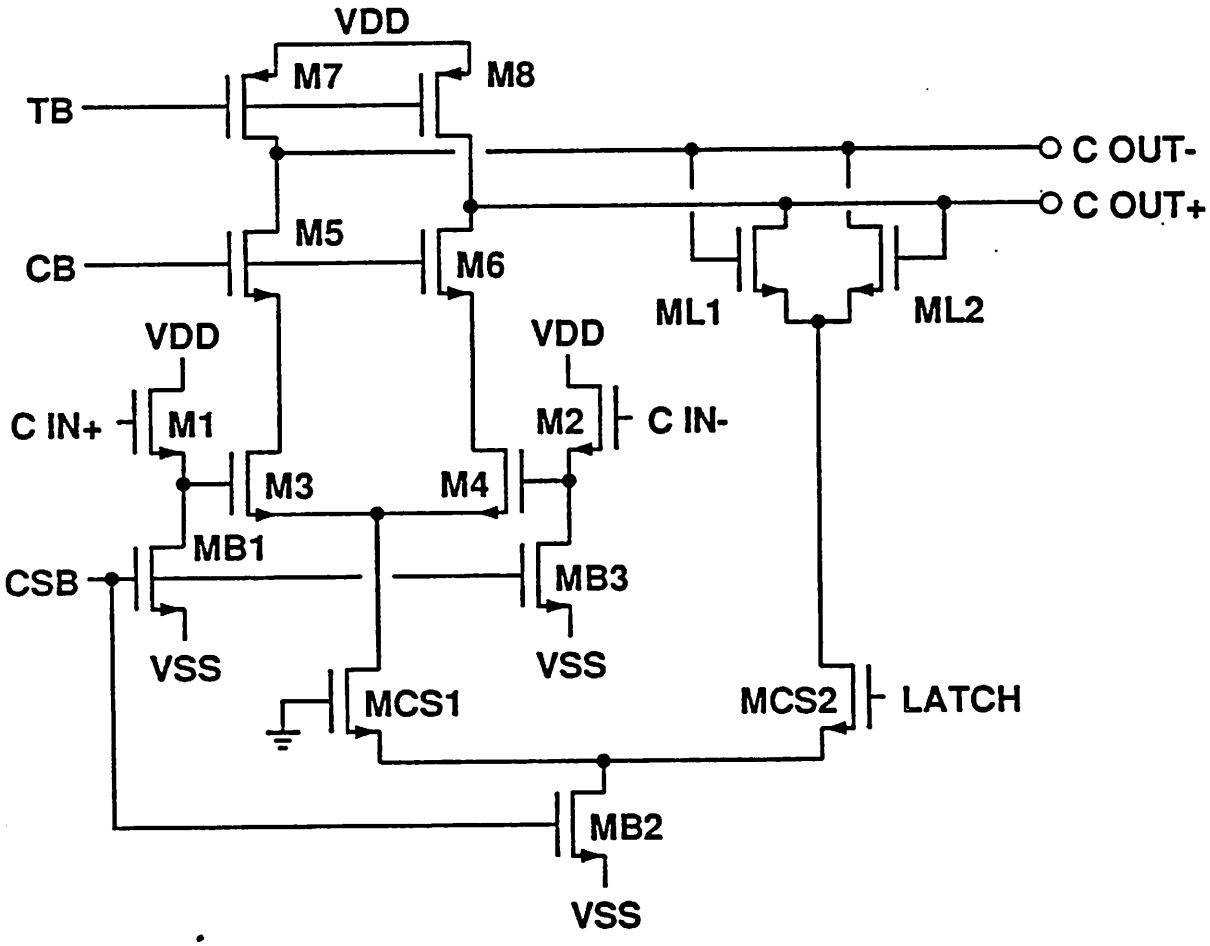

Figure 10 - Comparator schematic

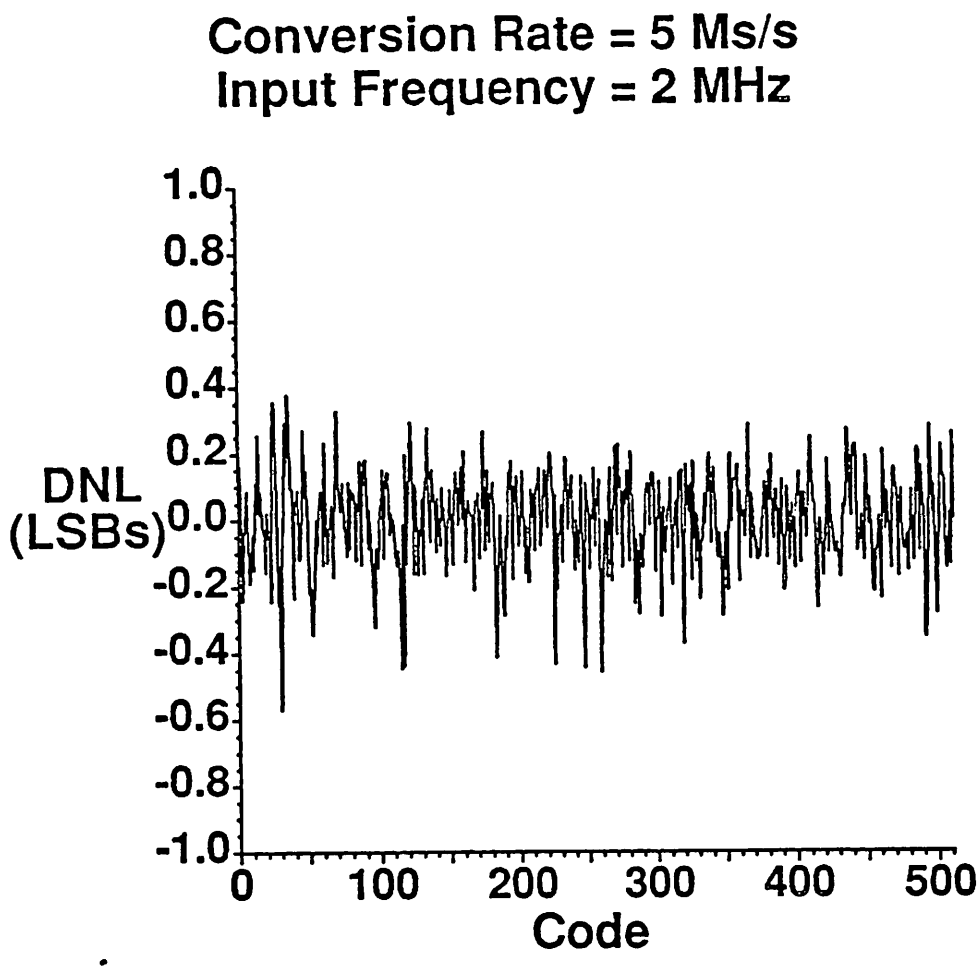

Figure 11 - DNL vs. code

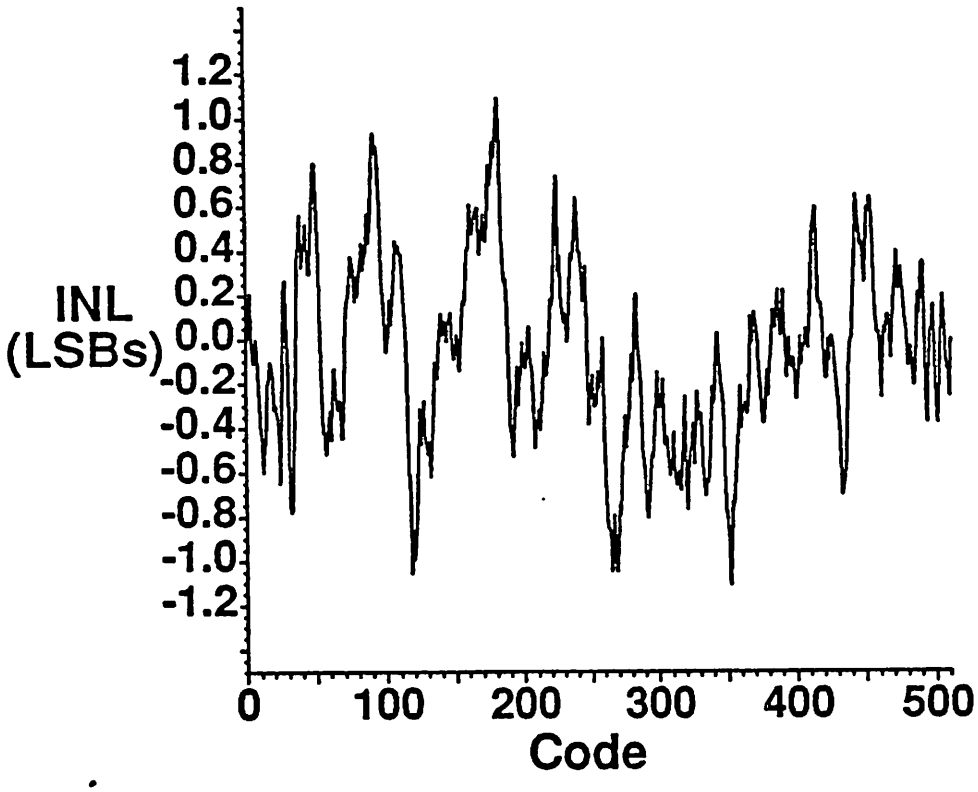

Figure 12 - INL vs. code

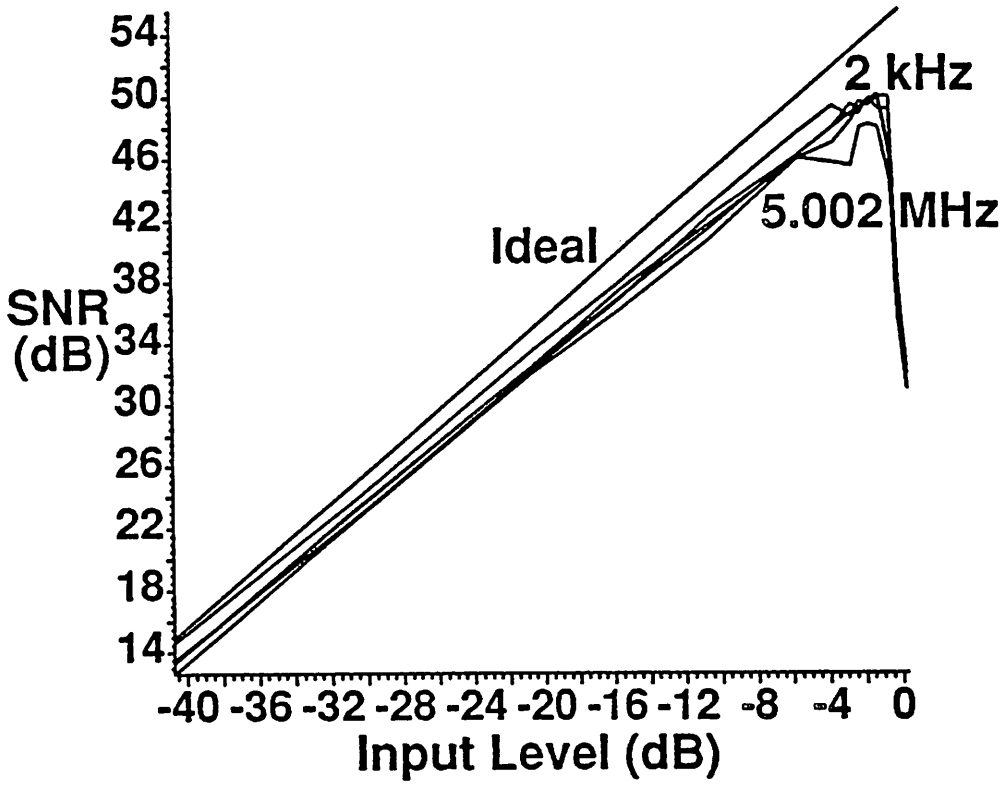

Figure 13 - SNR vs. input level

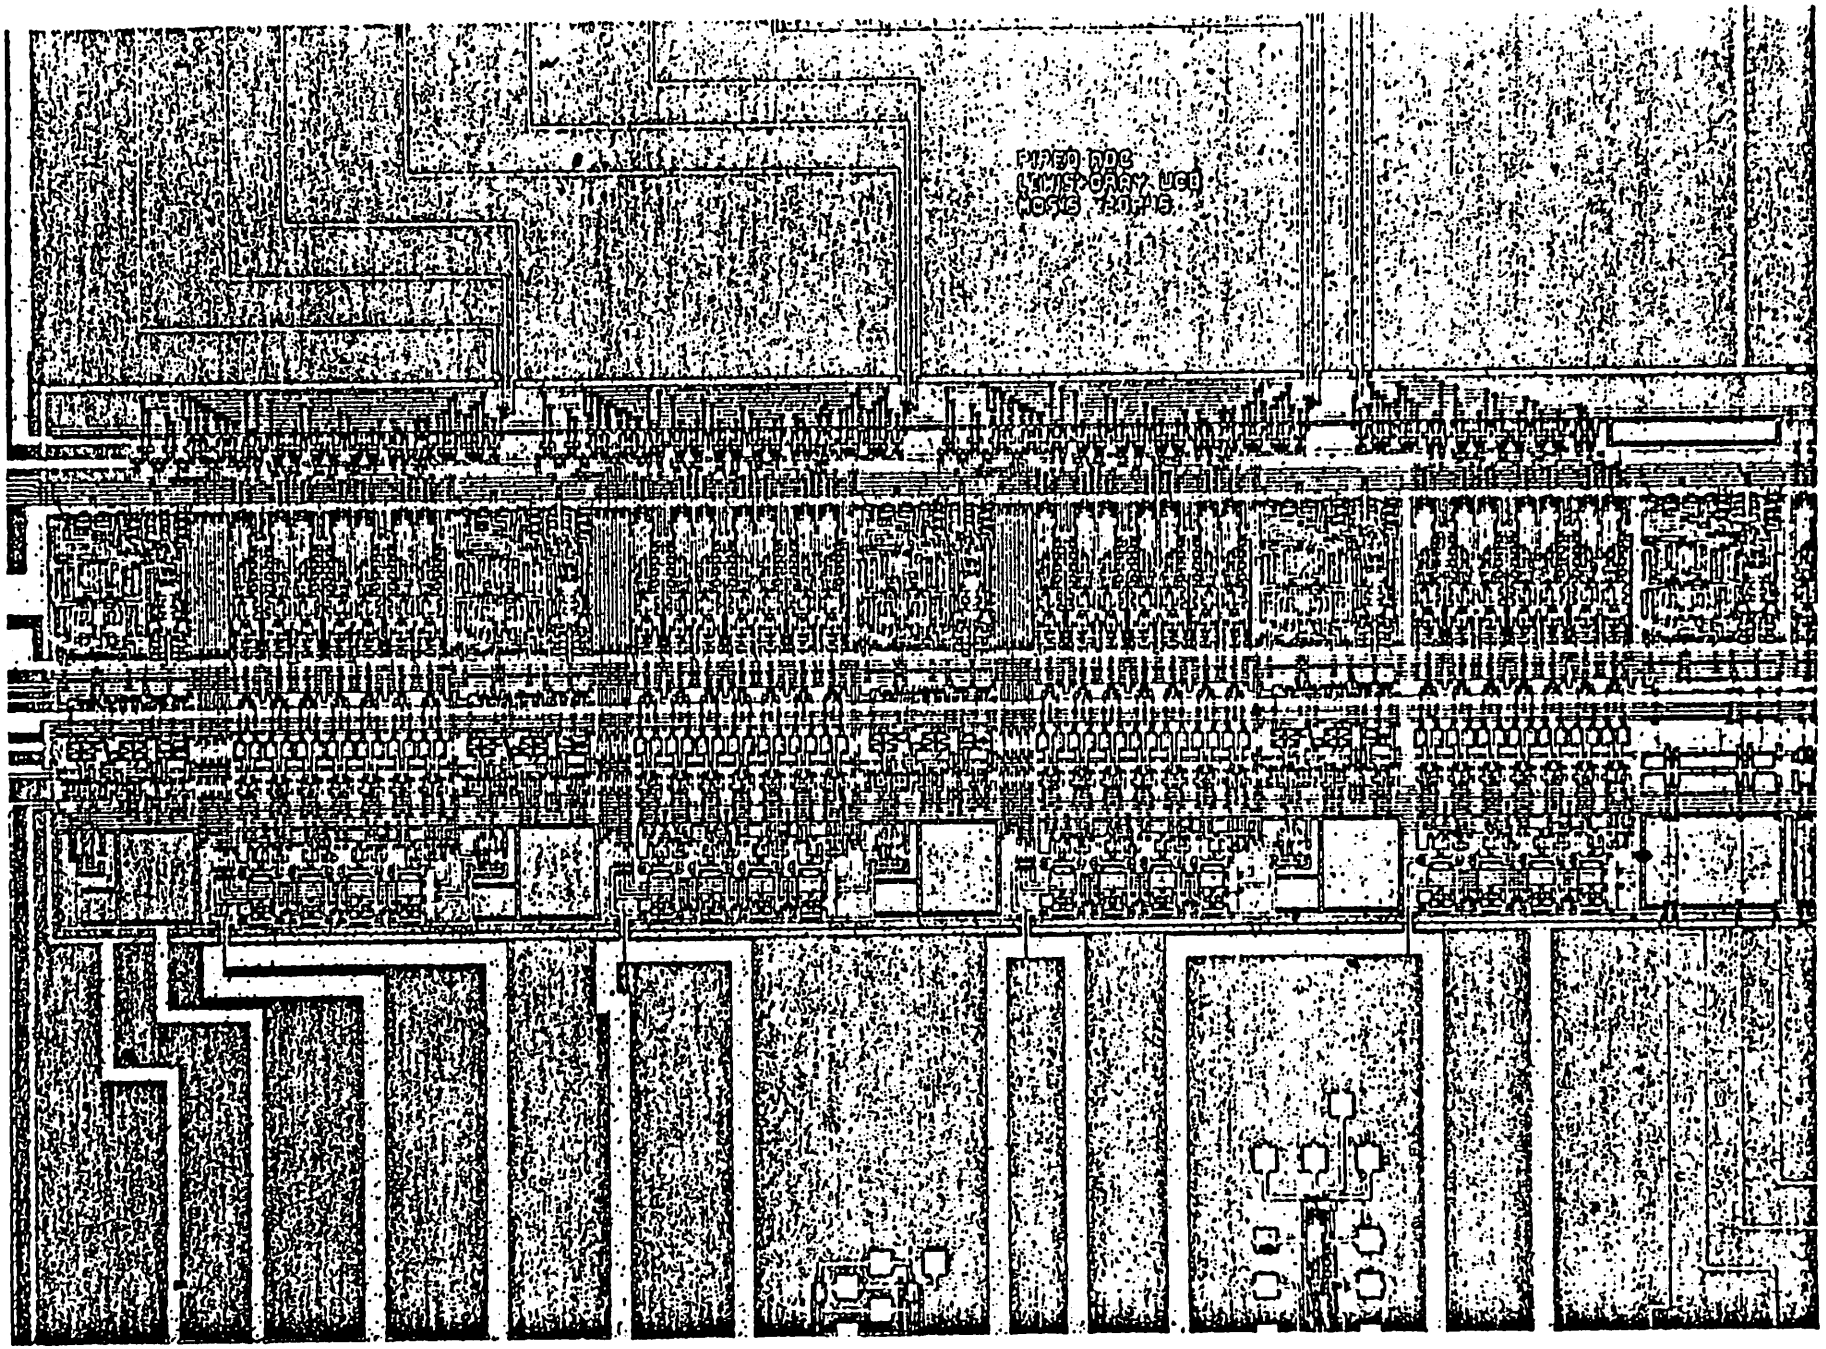

Figure 14 - Photograph of core of prototype.