A 2.5-V, 12-b, 5-MSample/s Pipelined CMOS ADC

Paul C. Yu, *Member, IEEE*, and Hae-Seung Lee, *Fellow, IEEE*

***Abstract*—A set of power minimization techniques is proposed for pipelined ADC's. These techniques include commutating feedback-capacitors, sharing of the op-amp between the adjacent stages of the pipeline, reusing the first stage of the op-amp as comparator pre-amp, and exploiting parasitic capacitors for common-mode feedback. This set of low-power design techniques is incorporated in an experimental chip fabricated in a 1.2-μm, double-poly, double-metal CMOS process. At 12-b 5-Msample/s, the chip dissipates 33 mW of power from a 2.5-V analog supply while achieving a maximum differential nonlinearity (DNL) of -0.78 and $+0.63$ least-significant bits (LSB) with a peak signal-to-noise ratio (SNR) of 67.6 dB.**

I. Introduction

MINIMIZATION of power in analog-to-digital converters (ADC's) is a challenging task due to the strong interdependent tradeoffs involved. In this paper, a 12-b, 5-Msample/s ADC is used to demonstrate a set of low-power design techniques for a pipelined architecture. These techniques include commutating feedback-capacitors (Section III), sharing of the op-amp between the adjacent stages of the pipeline (Section IV), reusing the first stage of the op-amp as comparator pre-amp (Section V), and exploiting parasitic capacitors as common-mode feedback capacitors (Section VI). Each of these techniques is presented in terms of its motivation, principle of operation, and relevant design considerations. Whenever appropriate, a comparison with similar or competing techniques is also given. These power minimization techniques can be implemented separately, or together as is the case of the prototype described in Section VII.

Before describing this set of low-power design techniques, a description of the pipelined architecture is first given.

II. General Description of a Pipeline

A typical pipelined architecture uses a number of similar pipelined stages labeled Stage 1, Stage 2, etc., as shown in Fig. 1. All of the pipelined stages are similar in construction, consisting of a sample-and-hold amplifier (SHA), a digital-to-analog-subconverter (DASC), an analog-to-digital-subconverter (ADSC), a subtractor, and a multiply by 2^{m_i} amplifier. The symbol m_i denotes the number of bits the ith stage of the pipeline resolves.

The input $V_{\rm IN}$ is first sampled-and-held and then digitized by the ADSC to arrive at the first m_0 most significan bits

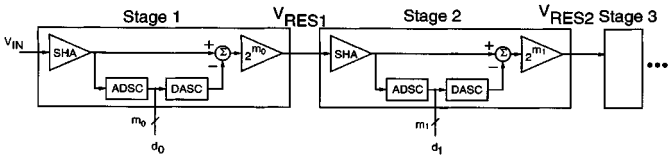

Fig. 1. A general pipelined architecture.

(MSB's), represented by d_0. This digital code d_0 is applied to the DASC to produce an analog voltage which is subtracted from the sampled-and-held input. The difference represents the residue and is amplified by 2^{m_0} to scale it back to the full scale. This amplified residue, V_{RES1}, is passed to Stage 2 as an input. After performing a similar set of operations as described for Stage 1, Stage 2 resolves the next m_1 MSB's. In a typical implementation, the stage gain is reduced by a factor of two and the digital codes d_0, d_1, d_2, etc., are overlapped by 1 b to perform digital error correction [1].

In this manner, as the input signal is processed by Stage 1, Stage 2 concurrently processes the residue signal V_{RES1} from the previous sample. The concurrency of operation by each stage in the pipeline allows the pipelined ADC to achieve high throughput suitable for video applications [1], [2].

In a pipelined architecture, the growth of the hardware is linear with the number of bits resolved. This linear dependence is in contrast with an exponential dependence for a flash architecture. However, a pipelined ADC typically needs high-performance analog components such as op-amps to perform the functions indicated in Fig. 1. These high-performance components can consume large amounts of power. In the following sections, a set of power minimization techniques is described for the pipelined ADC's.

III. The Commutated Feedback-Capacitor Switching (CFCS) Technique

A. Motivation

In a wide range of imaging applications, good differential nonlinearity (DNL) is required. However, integral nonlinearity (INL) is not very critical in these applications. As the INL requirement is relaxed, the resulting distortion requirement can also be relaxed. Therefore, the SNR as opposed to the signal-to-noise plus distortion (SNDR) will be the primary target. Given the relaxed INL, the DNL still needs to be satisfied. In a conventional technique, DNL specification alone requires stringent capacitor matching requirement [1]. The commutated feedback-capacitor switching (CFCS) technique relaxes the capacitor matching requirement to the point that it is easy to satisfy in most modern process technologies. Consequently, the technique allows the capacitors to be scaled down to the

Manuscript received June 18, 1996; revised July 22, 1996. This work was supported by National Science Foundation and ARPA Contract MIP-91-17724.

P. C. Yu was with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139 USA. He is now with Texas Instruments Inc., Dallas, TX 75243 USA.

H.-S. Lee is with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139 USA.

Publisher Item Identifier S 0018-9200(96)08220-0.

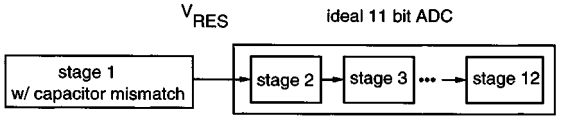

Fig. 2. A conceptual pipelined ADC with a capacitor mismatch in the first stage.

kT/C noise limit. With a reduced capacitive load, op-amp power consumption is also reduced.

B. Principle of Operation

For simplicity, we consider the case of a single-ended 1-b-per-stage pipelined ADC, where the first stage has a capacitor mismatch while the remaining 11 stages are ideal. The overall 12-b pipeline is schematically represented in Fig. 2.

We now examine the effect of the capacitor mismatch with $C_1 > C_2$ in the MSB stage. For comparison purposes, the conventional technique [3] is drawn in Fig. 3(a). During the sampling phase, both the input voltage $V_{\rm IN}$ and the offset voltage of the op-amp are sampled onto C_1 and C_2. For simplicity, $V_{OS} = 0$ is assumed. The case of $V_{OS} \neq 0$ will be dealt with in the next section on op-amp sharing technique. During the amplification phase, C_1 is selected as the feedback capacitor. Depending on the digital decision, d, $V_{\rm REF}$ or ground (GND) is subtracted from the sampled input. The amplified residue voltage $(V_{\rm RES})$ is shown in Fig. 4(a). Note that the slopes of the residue voltage are given by $(C_1 + C_2)/C_1$, regardless of whether $d = 0$ or 1. This is because a dedicated capacitor C_1 is used as the feedback capacitor. At the comparator decision boundary, as shown at the bottom of Fig. 4(a), the residue drop $(V_{\rm drop})$, given by $V_a - V_b$, depends on the matching of the capacitors.

The significance of $V_{\rm drop}$ becomes clear when two analog inputs $V_{\rm IN} = V_{\rm REF}/2 + \delta_{\rm in}$ and $V_{\rm IN} = V_{\rm REF}/2 - \delta_{\rm in}$ on either side of the decision point are considered. The quantity $\delta_{\rm in}$ represents a very small positive voltage compared with 1 LSB at 12-b level. Since the analog input barely changes, in order for the ADC to have a 12-b DNL, the two digital codes resulting from the two different analog inputs must not differ by more than 1 LSB. Referring to Fig. 4(a), when $V_{\rm IN} = V_{\rm REF}/2 - \delta_{\rm in}$, $d = 0$ and the residue V_a should make the following 11-b ADC to produce an output code $= \{1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1\}$. When $V_{\rm IN} = V_{\rm REF}/2 + \delta_{\rm in}$, $d = 1$ and the residue V_b ideally should make the following 11-b ADC to produce an output code $= \{0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0\}$. Therefore, to achieve 12-b DNL, the 11-b ADC must produce output codes that differ by the exact full scale. This means that $V_{\rm drop}$ from the first stage must be exactly $V_{\rm REF}$. Since the percentage mismatches ϵ_1 and ϵ_2 are random errors, the two cannot be expected to cancel each other. As a result, with conventional capacitor switching scheme as shown in Fig. 3(a), to satisfy the above condition requires 12-b capacitor matching for a 12-b DNL.

We now examine the proposed CFCS technique as shown in Fig. 3(b). During the sampling phase, the input is sampled on both C_1 and C_2 as in the conventional case. But during the amplifying phase, when $d = 0$, C_1 is selected as the feedback capacitor; on the other hand, when $d = 1$, C_2 is selected as the feedback capacitor. The corresponding residue plot is shown in Fig. 4(b). Since $C_1 > C_2$, we see that when $d = 0$, the residue voltage will be slightly smaller than ideal, resulting in a smaller slope. When $d = 1$, the residue voltage will be slightly larger than ideal, hence a steeper slope. Focusing on the decision point at $V_{\rm IN} = V_{\rm REF}/2$, it can be seen that, to either side of this boundary, the slopes of the residue voltage are changed in the opposite direction, yielding a residue drop $V_{\rm drop}$. This residue drop is nearly exactly $V_{\rm REF}$, the full scale of the following 11-b ADC. The error is only second-order, and a 12-b DNL requires only 6–7 b capacitor matching.

Fig. 5(a) and (b) shows the overall 12-b transfer curves of the conventional and the CFCS techniques, respectively.

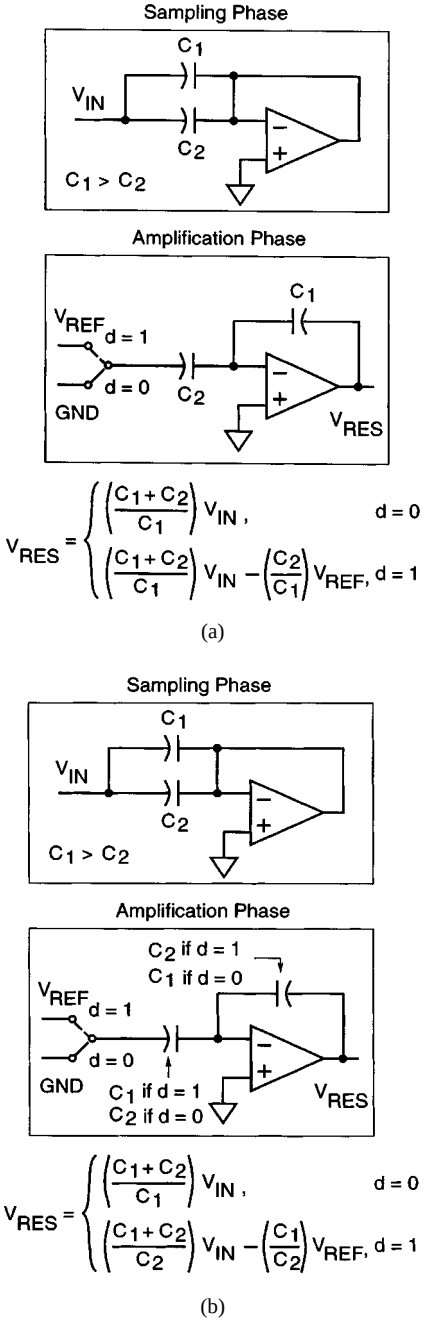

Fig. 3. Comparison between: (a) conventional and (b) CFCS techniques.

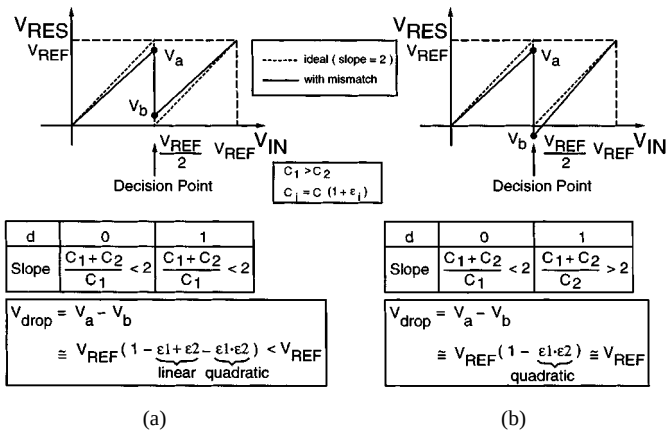

(a) (b)

Fig. 4. Residue plot comparison between (a) conventional and (b) CFCS techniques.

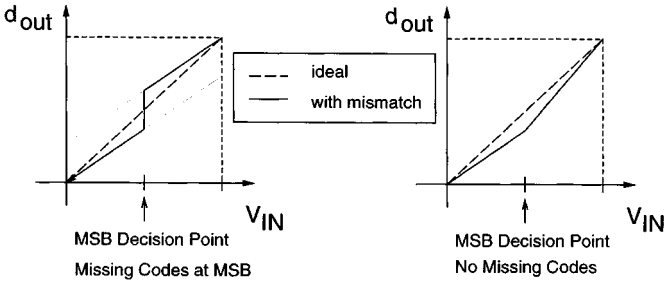

(a) (b)

Fig. 5. Transfer curve comparison between (a) conventional and (b) CFCS.

In Fig. 5(a), the thin dashed lines extending from the actual transfer curve show that the two segments have the same slope due to the use of a dedicated feedback capacitor in the conventional switching method. This results in a missing-code error. Using CFCS with the same mismatched capacitors, Fig. 5(b) shows that the first and the second segments of the transfer curve have slopes that are, respectively, less than and greater than ideal. The missing code error at the MSB decision is eliminated.

In the previous discussion, the mismatch is such that C_1 is assumed to be slightly greater than C_2. In the case when C_2 is greater than C_1, the CFCS technique is equally valid. In this case, the residue drop $V_{\rm drop}$ is greater than $V_{\rm REF}$ in the conventional case. The following stages of the pipeline will be saturated, resulting in wide codes as shown in the overall transfer curve in Fig. 6(a). With a similar mismatch using the CFCS technique, the wide codes are eliminated as shown in Fig. 6(b).

From Figs. 5(b) and 6(b), it is clear that while the INL is commensurate to the capacitor matching, the DNL is significantly enhanced from the conventional case.

Although a 1-b-per-stage case is considered, the CFCS technique is general and can be applied to pipelined architectures with any number of bits per stage [4]. A 2-b-per-stage architecture is implemented in [5].

C. Digital Error Correction with CFCS

We have focused on the MSB stage in Fig. 2. In practice, the CFCS technique can be applied to each of the following stages.

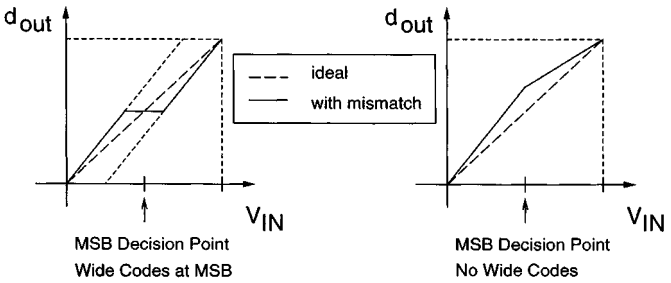

(a) (b)

Fig. 6. Transfer curves when $C_2 < C_1$: (a) conventional and (b) CFCS.

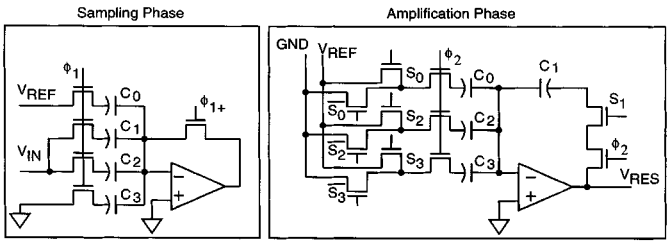

Fig. 7. Implementation of a pipelined stage with digital error correction.

TABLE I
CAPACITOR CONNECTION DURING THE AMPLIFYING PHASE IN A DIGITAL CORRECTION STAGE

V_{IN}	digital code	C_0	C_1	C_2	C_3
$-\frac{V_{REF}}{2} < V_{IN} \le 0$	-1	FB	GND	GND	GND
$0 < V_{IN} \le \frac{V_{REF}}{2}$	0	V_{REF}	FB	GND	GND
$\frac{V_{REF}}{2} < V_{IN} \le V_{REF}$	1	V_{REF}	V_{REF}	FB	GND
$V_{REF} < V_{IN} \le \frac{3V_{REF}}{2}$	2	V_{REF}	V_{REF}	V_{REF}	GND

For these later stages, a digital error correction technique can be used to accommodate the residue voltage, which can rise above or fall below the range of $[0, V_{\rm REF}]$. The standard error correction method described in [1] requires a reduction of the stage gain by a factor of two. The gain reduction is typically achieved by using a dedicated feedback capacitor and reducing the number of input capacitors by a factor of two [1]. As a result, it cannot be used in conjunction with the CFCS technique. Instead, error correction which employs two extra capacitors can be used.

For a 1-b-per-stage case, the circuit incorporating digital error correction is shown in Fig. 7. During the sampling phase, the bottom plates of C_1 and C_2 are connected to the input as in a regular pipelined stage. Two extra capacitors C_0 and C_3 are added with their bottom plates connected to $V_{\rm REF}$ and GND, respectively. During the amplifying phase, the feedback capacitor is commutated from C_0 to C_3, depending on the digital decision which can range from -1 when the input is under-range to 2 when the input is over-range. Table I shows the connection of the four capacitors during the amplifying phase.

Using charge conservation principle, it can be shown that the two added capacitors in Fig. 7 produce $V_{\rm drop}$ at all the decision points of GND, $V_{\rm REF}/2$, and $V_{\rm REF}$ that match with

the full-scale of the following stages in the pipeline. Hence, the input voltage outside the range of [GND, V_{REF}] can be folded back into the normal range and the digital codes produced in the process can be used to correct for the over-range [4].

D. Monte Carlo Simulation Results

The CFCS technique described above is general and can be applied to any number of bits per stage. A 2-b-per-stage case was analyzed both mathematically and via a Monte Carlo simulation in [4]. It can be shown that comparator offsets produce only second-order errors. The result of the Monte Carlo simulation is shown in [4] with a capacitor mismatch of 0.78% and a comparator offset of 1.56% of V_{REF}, both at the 3-σ level. A yield of 95.2% is obtained for the 12-b 0.5-LSB DNL. In a conventional converter, only 7-b DNL would be achievable with this degree of matching.

IV. Op-Amp Sharing Technique

Whereas the CFCS technique described in the previous section minimizes power by minimizing the capacitive load of the op-amps, the op-amp sharing technique minimizes power by reducing the effective number of the op-amps used in the entire pipeline. This technique is described in this section.

A. Motivation

In a pipelined ADC, during the odd phases, the odd stages are sampling while the even-stages are amplifying. During the even phase, the roles of the odd and the even stages interchange. During the sampling phase, both the analog input and the op-amp offset are sampled. In a conventional switching scheme [3], the op-amp is in a unity-gain configuration for the offset to be sampled. The op-amp power during the sampling phase is used solely for the purpose of offset cancellation.

If the op-amp does not need to be offset-canceled, adjacent stages of the pipeline can share one op-amp, resulting in the use of half the number of op-amps. In the technique described below, a large saving in the power consumption is achieved, while offset cancellation is maintained.

B. Principle of Operation

The proposed op-amp sharing method is shown in Fig. 8. During the odd phase, the op-amp labeled a_{1o}, used for the odd stage of the pipeline, is in the sampling mode. At the same time, the even stage is in the amplifying mode. A two-stage op-amp design is used, consisting of a_{1e}, which is identical to a_{1o}, and a_2. This second stage of the op-amp, a_2, is shared between adjacent stages of the pipeline so that when the phase changes to even, a_2 is switched from the even stage to the odd stage. To push out the nondominant pole, a_2 is typically designed to consume more current than a_{1o} and a_{1e}. Assuming a_2 consumes four times as much current as a_1, the op-amp sharing technique will achieve a power saving of 40%.

During the odd phase, the op-amp a_{1o} is in a unity feedback configuration where the output is directly connected to the inverting input of the op-amp at virtual ground. Therefore, very little voltage swing is required. As a result, a cascode op-amp topology can be used to achieve high dc open-loop gain without being penalized for output voltage swing reduction. The offset of a_{1o} that is sampled during the odd phase is offset-canceled during the even phase. The offset of a_2, although not sampled, is divided by the high gain of a_{1o} when referred to the input of the two-stage op-amp. Since a_{1o} and a_{1e} already produce large dc open-loop gain, the dc gain of a_2 can be small. This allows a_2 to use a noncascoded topology which achieves high swing.

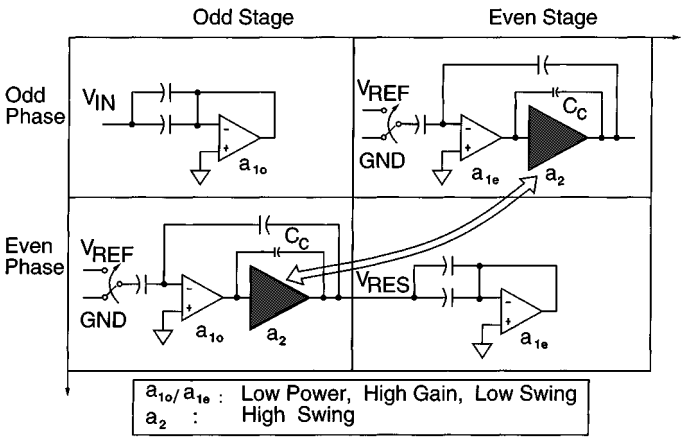

Fig. 8. Op-amp sharing technique results in about 40% saving in power.

C. Comparison with Other Techniques

For a cyclic ADC, a similar op-amp sharing idea was proposed by [6]. While an excellent technique for cyclic converters, this technique requires an auxiliary amplifier for the offset cancellation. When applied to a pipelined ADC, it consumes an unnecessarily large amount of power.

V. Reusing the First Stage of an Op-Amp as a Comparator Pre-Amp

In the previous section, power saving is accomplished by sharing the second stage of the op-amp between the adjacent stages of the pipeline. In this section, power saving is accomplished by reusing the first stage of the op-amp as a comparator pre-amp. This technique is described in this section.

A. Motivation

Although the CFCS technique is tolerant to comparator offsets, having small comparator offset allows the residue voltage to exceed the full scale by only a small amount. This reduced over-range requirement allows a larger full scale for a given op-amp output voltage swing. A careful examination of the op-amp topology yields a pre-amp that already exists for free. This pre-amp is used to reduce the comparator offset when referred to the input of the pre-amp.

B. Principle of Operation

As shown in Fig. 8 in Section IV, during the sampling phase, the op-amp has only one stage. The circuit schematic for a_{1o} is shown in Fig. 9(a). The first stage amplifier a_{1e}

for the even stage of the pipeline is identical to a_{1o}. Since the op-amp is fully differential, a common-mode circuit, also shown in Fig. 9(a), is used. Transistors M5 and M6 are used to sense the output common-mode voltage of op-amp a_{1o}, while M7 is used to apply the feedback signal V_1 to the op-amp so as to set the common-mode to the desired V_{CMR}. Used as a common-mode circuit, the drain of M5 and M6 are typically tied to the supply and their output V_2 is wasted. With the addition of two resistors as shown in Fig. 9(a), M5 and M6 double as the second stage of a two-stage differential pre-amp for the comparator. As shown in Fig. 9(b), once the input is sampled, the one-stage op-amp goes into open loop. By grounding the bottom plates of C_1 and C_2, the inverted input is offset-canceled and applied to the pre-amp. The first stage of the pre-amp was previously the op-amp (a_{1o}), while the second stage of the pre-amp was the common-mode circuit (a_{1CM}). The output of this two-stage pre-amp, V_2, is applied as the input to the latch. Since an op-amp in an open loop can be slow, a resistor at the output of a_{1o} is switched in to reduce the op-amp open-loop gain when the op-amp is reused as a comparator pre-amp. Because the resistor is across the output, the differential offset of a_{1o} is not affected. It should be noted that this technique can only be used in a 1-b-per-stage architecture.

C. Advantages

In addition to the obvious power reduction, power is further reduced by the elimination of the first sample-and-hold stage. In many previously reported pipelined ADC's, e.g., [2], the first stage of a pipeline is an explicit sample-and-hold stage, employing a single sampling capacitor. The gain-of-one sample-and-hold stage contributes kT/C noise, requiring larger capacitors both in the sample-and-hold stage and in the following stages of the ADC. Since the configuration shown in Fig. 9(b) performs a sample-and-hold function, the explicit gain-of-one sample-and-hold stage can be eliminated.

To see how much power can be saved when the explicit sample-and-hold is eliminated, consider first a 1-b-per-stage pipeline where all the stages are identical. Let $\overline{v_{n1}^2}$ be the noise contribution from the first stage of the pipeline. In the limit when the number of pipelined stages approaches infinity, the total input referred noise of the ADC is given by

$$\begin{aligned}\overline{v_{tot1}^2} &= \overline{v_{n1}^2}(1+1+\tfrac{1}{4}+\tfrac{1}{16}+\cdots)\\ &= \tfrac{7}{3}\cdot\overline{v}_{n1}^2.\end{aligned} \tag{1}$$

In contrast, when the explicit sample-and-hold is eliminated, the total input referred noise of the ADC is given by

$$\begin{aligned}\overline{v_{tot2}^2} &= \overline{v_{n2}^2}(1+\tfrac{1}{4}+\tfrac{1}{16}+\cdots)\\ &= \tfrac{4}{3}\cdot\overline{v_{n2}^2}.\end{aligned} \tag{2}$$

Equating the noise given by (1) and (2), it can be shown that

$$\frac{\overline{v_{n1}^2}}{\overline{v_{n2}^2}} = \frac{4}{7}. \tag{3}$$

Equation (3) implies that when an explicit sample-and-hold stage is eliminated, the sampling capacitance of each pipelined stage can be 43% smaller. Hence, the widths of the transistors as well as the current can be both reduced by 43%.

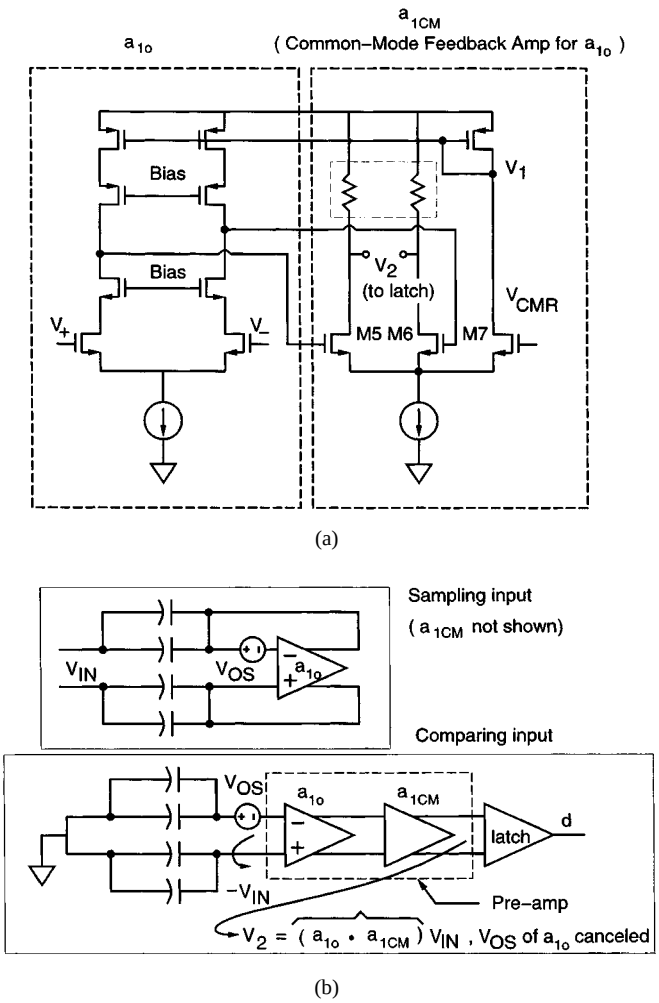

Fig. 9. Reusing the first stage of an op-amp as a comparator pre-amp: (a) circuit schematics and (b) block diagram.

For larger number of bits per stage, since the contribution from the following stages is negligible, the power saving can approach 50%. The timing necessary to implement this technique requires that the digital decision be carried out after the sampling phase, while the timing for the pipeline requires the digital decision to be ready at the onset of the amplifying phase. Hence, a pseudophase is necessary between the sampling and the amplifying phases. This added phase will be referred to as the comparison phase. For a conversion rate of 5 Msample/s, each phase is 100 ns. For a comparison time of 10 ns, the additional penalty of 10% is a small price to pay for the benefit of power saving. The pseudophase scheme has the added advantage that the follower circuit used to prevent the kickback of the latch [7] can be eliminated.

VI. Exploiting Parasitic Capacitors

The common-mode feedback circuit of a_2, the second stage of the op-amp, is implemented by exploiting existing parasitic capacitors which would otherwise load the amplifier. This technique is described in this section.

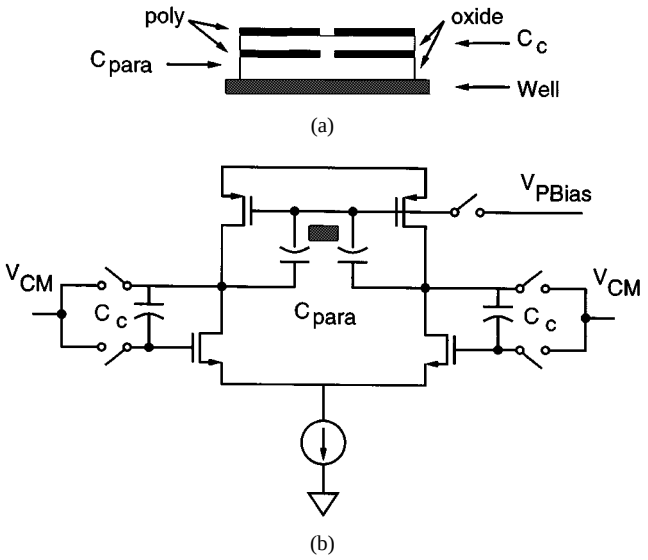

Fig. 10. Exploiting parasitic capacitance for the common-mode feedback of the second stage of the op-amp.

A. Motivation

In a typical pipelined ADC, the total op-amp load capacitance consists of three components: the sampling capacitance, the common-mode feedback capacitance, and the parasitics from the various devices connected at the output. In a two-stage design, a pole-splitting compensation capacitor (C_c) is used as shown in Fig. 10. The bottom-plate of C_c is typically connected to the output of the op-amp and therefore adds additional loading to the output.

B. Principle of Operation

During the comparison phase when the second stage of the op-amp is not in use, the switches shown in Fig. 10 are closed. In this case, the inputs and outputs are shorted to the common-mode voltage V_{CM}, while the gates of the PMOS load are shorted to the voltage $V_{P\,Bias}$. A PMOS diode-connected transistor is used to generate a $V_{P\,Bias}$ so that each of the PMOS devices is biased at half of the tail current. In this manner, the desired voltage difference between V_{CM} and $V_{P\,Bias}$ is stored on the common-mode feedback capacitors.

By inserting a well underneath the capacitor and connecting this well to the gates of the PMOS load as shown, the parasitic capacitance from the bottom plate of the compensation capacitor to this well can be exploited as part of the common-mode feedback capacitor for the second stage of the op-amp.

The switches shown in Fig. 10 are open during the normal mode of operation. When the common-mode output is too high, the common-mode feedback capacitors act as a dc level shifter, pulling up the gates of the PMOS devices. This, in turn, decreases the common-mode output voltage back to the desired V_{CM}.

In this manner, the bottom-plate parasitics from C_c are exploited as common-mode feedback capacitors. Power saving is achieved by reducing the total load capacitance and hence power consumption by approximately 20% while maintaining the same settling time.

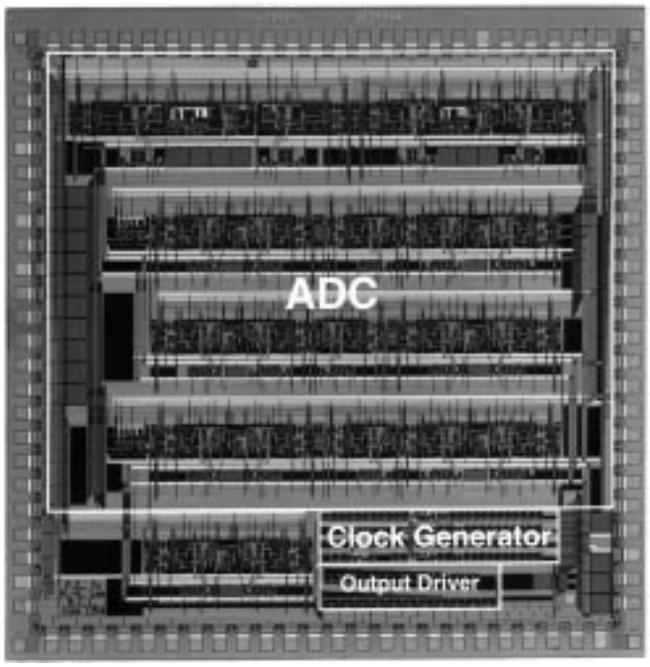

Fig. 11. Chip micrograph.

C. Common-Mode Settling Consideration

The sheet resistance of the well is typically on the order of 2 k$\Omega/\square$. Since the well is inserted underneath a square capacitor, the well presents approximately 2 kΩ of series resistance when considered as a lumped resistance. Together with a worst-case capacitance of 5 pF, the pole associated with the RC circuit is at 16 MHz. If the unity-gain frequency of the common-mode circuit is also at 16 MHz, the resulting phase margin of 45° is more than sufficient for a 5-MHz ADC. When modeled as a distributed RC circuit, the phase shift is less than the lumped RC circuit until a frequency that is a few times 16 MHz. So again, for 5-MHz operation, the effect of well resistance on the common-mode settling is not critical.

VII. Experimental Results

The power minimization techniques described in Sections III–VI are incorporated in an experimental chip fabricated in a 1.2-μm, double-poly, double-metal CMOS process. Fig. 11 shows the chip micrograph. All of the experimental results presented in this section were obtained when the chip was powered by a 2.5-V analog supply and dissipates 33 mW including the on-chip bias generators. For the digital supply, although on-chip charge-pumping techniques could be used to turn on the MOS switches [8], [9], for ease of design and testing, a 4.2-V supply was used. When a charge pumping circuit is used, the digital supply is about one V_T below two times the analog supply. The 4.2-V digital supply is roughly equal to this pumped supply when the analog supply is 2.5 V.

A. Histogram Code Density Test

A 9.875 60-kHz sine wave is applied to the ADC operating at a 5.0-Msample/s conversion rate from a 2.5-V analog

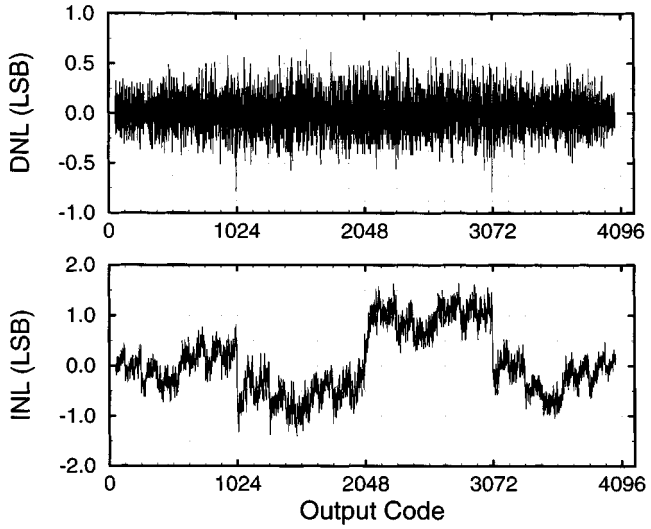

Fig. 12. Measured DNL and INL at $f_s = 5.0$ MHz and $f_{in} = 9.875\,60$ kHz.

supply. A total of 128 000 points are collected to compute the DNL and INL which are shown in Fig. 12 at 12-b level. The INL jumps near the major carries are larger than expected because CFCS is supposed to remove them.

One possible source of this is a parasitic coupling capacitance in the second MSB stage. This parasitic capacitance can be from the top plate of a sampling capacitor to a digital node such as the gate of a bottom-plate switch. During the sampling phase, the parasitic capacitance C_p couples to the digital node, which is reset to a high voltage level at V_H. During the amplifying phase, C_p couples to either V_H or V_L depending on the digital code d. The effect of C_p is to introduce a code-dependent offset to the residue voltage. As a result, the residue drop is decreased from the ideal V_{REF} by this offset [10]. For $V_H = 4.2$ V, $V_L = 0$, $V_{\mathrm{REF}} = 1$ V, and $C_2 = 0.5$ pF, a C_p of 0.12 fF is sufficient to cause a -1 LSB DNL error at the second MSB transition. Note that the effect of this parasitic capacitance is not a problem unique to the CFCS technique. For a conventional switching technique, the same parasitic capacitance still gives a code-dependent offset.

In future design, the top plates can be either shielded or sandwiched between two layers of bottom plates to eliminate this type of parasitics capacitance. In addition, if C_p is better matched between the two sides of a differential implementation, the parasitic effect is common-mode and can be rejected by the op-amp.

B. The FFT Test

In the fast Fourier transform (FFT) test, a 2.2-MHz sine wave is digitized by the ADC under test. A total of 16 384 samples of data are collected. The FFT plot is shown in Fig. 13.

The -66.5 dB second harmonic is consistent with a 10-b capacitor matching [10]. The remaining harmonics are due to the discontinuities in the INL.

The signal power, the higher harmonics power, and the noise power can be calculated from the same data. Fig. 14

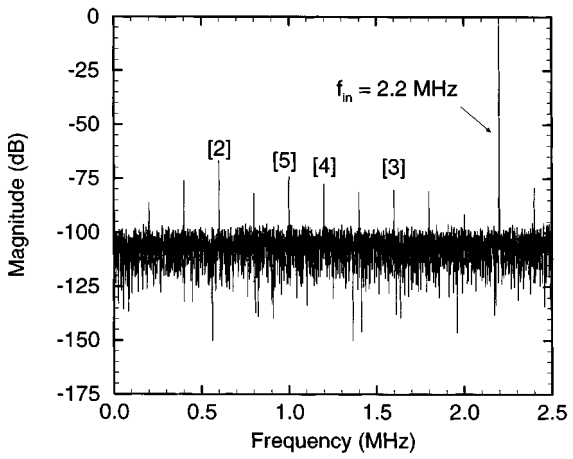

Fig. 13. Measured FFT plot at $f_s = 5.0$ MHz.

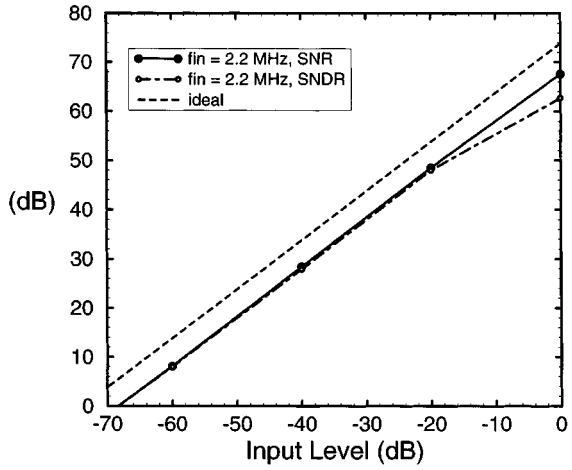

Fig. 14. Measured SNR and SNDR versus input signal level at $f_s = 5.0$ MHz.

shows the SNR with a 2.2-MHz input frequency. The peak SNR is 67.6 dB. This is about 5.4 dB lower than the expected SNR of 73 dB when the rms kT/C noise is about a factor of two smaller than the rms quantization noise. For reference, the theoretical SNR is shown. The experimental SNDR at the same input frequency of 2.2 MHz is also shown with a peak value of 62.7 dB. At low input amplitudes, the SNDR is nearly equal to the SNR. The performance parameters are summarized in Table II.

VIII. Conclusion

A set of power minimization techniques for pipelined ADC's is described. The first is a commutated feedback capacitor switching scheme that saves power by relaxing the capacitor matching requirement. When the relaxed matching requirement is relatively easy to satisfy, as is the case in most modern process technologies, the CFCS technique can be used to scale capacitors down to the kT/C limit. The second technique is an op-amp sharing method that reduces power consumption by as much as 40%. The third technique saves power by reusing the first stage of an op-amp as a pre-amp for the comparator. The fourth technique exploits

TABLE II
EXPERIMENTAL PERFORMANCE SUMMARY

Technology	1.2-μm CMOS
Resolution	12 b
Conversion Rate	5.0 Ms/s
Minimum Supply Voltage	2.5 V (Analog)
Power	33 mW (Analog)
Input Capacitance	1.0 pF (single-ended)
DNL	+0.63/-0.78 LSB
Peak SNR	67.6 dB at f_{in} = 2.2 MHz
Active Die Area	4.1 x4.2 mm^2

parasitic capacitors as the common-mode feedback capacitors, resulting in a reduced op-amp power consumption for the same settling speed. The set of techniques proposed here for pipelined ADC's can be investigated for applications to other analog subsystems such as switched-capacitor filters and sigma-delta converters.

ACKNOWLEDGMENT

Dr. A. Karanicolas provided the data acquisition hardware and software as well as technical discussions. During the testing phase, Analog Devices provided advice and assistance.

REFERENCES

[1] S. H. Lewis and P. R. Gray, "A pipelined 5 Msample/s 9-b analog-to-digital converter," *IEEE J. Solid-State Circuits*, vol. SC-22, pp. 954–961, Dec. 1987.
[2] 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.
[3] B.-S. Song, M. F. Tompsett, and K. R. Lakshmikumar, "A 12-b 1-Msamples/s capacitor error-averaging pipelined A/D converter," *IEEE J. Solid-State Circuits*, vol. SC-23, pp. 1324–1333, Dec. 1988.
[4] P. C. Yu and H.-S. Lee, "A pipelined A/D conversion technique with near-inherent monotonicity," *IEEE Trans. Circuits Syst. II: Analog and Digital Sig. Process.*, vol. 42, pp. 500–502, July 1995.
[5] J. Yang and H.-S. Lee, "A CMOS 12-bit 4 MHz pipelined A/D converter with commutative feedback capacitor," presented at *IEEE Custom Integrated Circuits Conf.*, May 1996.
[6] K. Nagaraj, "Efficient circuit configuration for algorithmic analog to digital converters," *IEEE Trans. Circuits Syst. II: Analog and Digital Sig. Process.*, vol. 40, pp. 777–785, Dec. 1993.
[7] A. N. Karanicolas, "Digitally self-calibration techniques for high-accuracy, high-speed analog-to-digital converters," Ph.D. dissertation, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Feb. 1994.
[8] Y. Nakagome, H. Tanaka, K. Takeuchi, E. Kume, Y. Watanabe, T. Kaga, Y. Kawamoto, F. Murai, R. Izawa, D. Hisamoto, T. Kisu, T. Nishida, E. Takeda, and K. Itoh, "An experimental 1.5-V 64-Mb DRAM," *IEEE J. Solid-State Circuits*, vol. 26, pp. 465–472, Apr. 1991.
[9] T. B. Cho and P. R. Gray, "A 10 b, 20 Msample/s 35 mW pipeline A/D converter," *IEEE J. Solid-State Circuits*, vol. 30, pp. 166–172, Mar. 1995.
[10] P. C. Yu, "Low-power design techniques for pipelined analog-to-digital converters," Ph.D. dissertation, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, June 1996.

Paul C. Yu (S'92–M'96) was born in Taipei, Taiwan, in 1967. He received the B.S. and M.S. degrees in 1990 and the Ph.D. degree in 1996, all in electrical engineering, from the Massachusetts Institute of Technology (MIT), Cambridge, MA.

Between 1990 and 1996, he was with the MIT Microsystems Technology Laboratories. Since September, 1996, he has been with Texas Instruments Inc., Dallas, TX, where his research is focused on analog circuit design.

Hae-Seung Lee (M'85–SM'92–F'96) was born in Seoul, South Korea, in 1955. He received the B.S. and the M.S. degrees in electronic engineering from Seoul National University, Seoul, Korea, in 1978 and 1980, respectively. He received the Ph.D. degree in electrical engineering from the University of California, Berkeley, in 1984, where he developed self-calibration techniques for A/D converters.

In 1980, he was a Member of Technical Staff in the Department of Mechanical Engineering at the Korean Institute of Science and Technology, Seoul, Korea, where he was involved in the development of alternative energy sources. Since 1984, he has been with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, where he is now Professor. Since 1985, he has acted as Consultant to Analog Devices, Inc., Wilmington, MA. His research interests are in the areas of analog integrated circuits, early vision circuits, fabrication technologies, and solid-state sensors.

Dr. Lee is a recipient of the 1988 Presidential Young Investigators' Award. He has served on a number of technical program committees for various IEEE conferences, including the International Electron Devices Meeting, the International Solid-State Circuits Conference, the Custom Integrated Circuits Conference, and the IEEE Symposium on VLSI circuits. From 1992 to 1994, he was an Associate Editor for the IEEE JOURNAL OF SOLID-STATE CIRCUITS.