A 14-b 12-MS/s CMOS Pipeline ADC With Over 100-dB SFDR

Yun Chiu, *Student Member, IEEE*, Paul R. Gray, *Fellow, IEEE*, and Borivoje Nikolić, *Member, IEEE*

***Abstract*—A 1.8-V 14-b 12-MS/s pseudo-differential pipeline analog-to-digital converter (ADC) using a passive capacitor error-averaging technique and a nested CMOS gain-boosting technique is described. The converter is optimized for low-voltage low-power applications by applying an optimum stage-scaling algorithm at the architectural level and an opamp and comparator sharing technique at the circuit level. Prototyped in a 0.18-μm 6M-1P CMOS process, this converter achieves a peak signal-to-noise plus distortion ratio (SNDR) of 75.5 dB and a 103-dB spurious-free dynamic range (SFDR) without trimming, calibration, or dithering. With a 1-MHz analog input, the maximum differential nonlinearity is 0.47 LSB and the maximum integral nonlinearity is 0.54 LSB. The large analog bandwidth of the front-end sample-and-hold circuit is achieved using bootstrapped thin-oxide transistors as switches, resulting in an SFDR of 97 dB when a 40-MHz full-scale input is digitized. The ADC occupies an active area of 10 mm^2 and dissipates 98 mW.**

***Index Terms*—Analog integrated circuits, capacitor mismatch, comparator sharing, discrete-time common-mode voltage regulation, early comparison, low power, low voltage, nested CMOS gain boosting, opamp sharing, passive capacitor error-averaging, pipeline analog-to-digital converter, pseudo-differential, subsampling.**

I. Introduction

In LOW-COST highly integrated system-on-a-chip (SoC) digital communication receivers, the boundary between analog and digital functionalities is constantly redefined as a result of fabrication technology advancement. The trend toward more digital signal processing (DSP) for multistandard agility in receiver designs has recently created a great demand for low-power low-voltage analog-to-digital converters (ADCs) that can be realized in a mainstream deep-submicron CMOS technology. Intended for embedded applications, the specifications of such converters emphasize high dynamic range and low spurious spectral performance. In a highly integrated wireless receiver, regardless of whether frequency translation is accomplished with a single conversion (direct conversion and low-IF architecture) or a wideband-IF double conversion, the lack of high-Q on-chip IF channel-select filters inevitably leads to a large dynamic range imposed on the baseband circuits in the presence of in-band blockers (strong adjacent channel interfering signals). For example, the worst case blocking specs of some wireless standards, such as GSM, dictate a conversion linearity of 14–16 bits to avoid losing a weak received signal due to distortion artifacts [1]–[3]. Recent works also underline the trend toward the IF-digitizing architecture to enhance programmability and to achieve a more "digital" receiver [4]–[6]. However, advancing the digitizing interface toward the antenna exacerbates the existing dynamic range problem, as it also requires a high oversampling ratio. To achieve high linearity, high dynamic range, and high sampling speed simultaneously under low supply voltages in deep-submicron CMOS technology with low power consumption has so far been conceived of as extremely challenging.

Among various ADC architectures, the pipeline converter is widely used in Nyquist sampling applications that require a combination of high resolution and high throughput. This study describes the prototype design of a 14-b pipeline ADC fabricated in a 0.18-μm digital CMOS technology that achieves a 75.5-dB SNDR at 12 MS/s with a power consumption of 98 mW [21]. The measured SFDR of over 97 dB up to 40 MHz of input frequency demonstrates the outstanding linearity and bandwidth of the analog circuitry operating at a low supply voltage of 1.8 V. The design techniques and strategies that have enabled this performance include a power-efficient ratio-independent conversion scheme [7], a pipeline ADC stage-scaling algorithm, a nested CMOS gain-boosting technique, a $\Delta\Sigma$ common-mode (CM) voltage regulation circuit, an opamp and comparator sharing technique, and the use of minimum channel-length, thin-oxide transistors for the switches combined with clock bootstrapping [10] and in-line switch techniques [11].

Section II of this paper reviews the pipeline ADC architecture and discusses the design challenges for switched-capacitor circuits in deep-submicron CMOS technology. To alleviate the prominent issue of power efficiency under low supply voltages, a pipeline ADC stage-scaling analysis is then introduced that determines the optimum stage resolution and scaling factor simultaneously. Following this, Section III highlights the key linearity technique of this design—the passive capacitor error-averaging (PCEA) technique. In Section IV, the details of the circuit implementation issues are presented. The experimental results of the chip are summarized in Section V with a short conclusion following in Section VI.

II. Architecture Power Efficiency

A. Pipeline ADC Architecture

A pipeline ADC is inherently a multistep amplitude quantizer in which digitization is performed by a cascade of many topologically similar or identical stages of low-resolution analog-to-digital encoders. Pipelining enables high conversion

Manuscript received April 19, 2004; revised July 28, 2004. This work was supported by the California MICRO Program, the National Science Foundation, Intel Corporation, and MARCO C2S2.

The authors are with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720-1770 USA (e-mail: chiuyun@eecs.berkeley.edu).

Digital Object Identifier 10.1109/JSSC.2004.836232

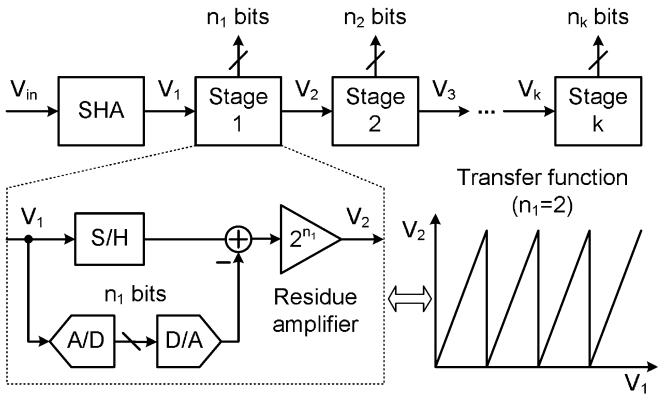

Fig. 1. Pipeline ADC block diagram.

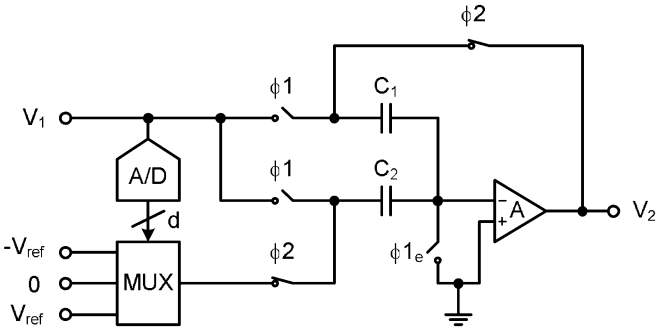

Fig. 2. Circuit diagram of the 1.5-b/s MDAC.

throughput by inserting analog registers, i.e., sample-and-hold amplifiers (SHAs), in between stages that allow a concurrent operation of all stages. This is done at the cost of an increased latency. In a typical pipeline stage as shown in Fig. 1, the signal passed to the subsequent stage is the conversion residue of the current stage created by a digital-to-analog converter (DAC) and a subtraction circuit. The maximum swing of this residue signal is often brought back to the full-scale reference level with a precision amplifier, also shown in Fig. 1. This keeps the signal level constant and allows the sharing of an identical reference throughout the pipeline stages. Breaking a high-resolution conversion into multiple steps greatly reduces the total number of comparators in contrast to a flash converter. In the limiting case, a 1-b/stage (b/s) pipeline ADC only needs N comparators to resolve an N-bit word as opposed to 2^N comparators required by a flash converter. The large accumulative interstage gain also relaxes the impact of later stage circuit nonidealities (noise, nonlinearity, and offset) on the overall conversion accuracy. For medium- to high-resolution Nyquist applications, pipeline ADCs have been demonstrated to achieve the lowest power consumption at high conversion rates [12]–[21].

In CMOS circuit technology, a typical pipeline ADC stage usually consists of a coarse comparator and a compact switched-capacitor (SC) circuit termed the multiplier DAC (MDAC), which integrates the sample-and-hold (S/H), the DAC, the subtraction, and the residue-gain functions. The circuit diagram of a single-ended 1.5-b/s MDAC is shown in Fig. 2. This architecture is also known to tolerate large comparator offsets due to the built-in decision level overlaps between successive stages, usually referred to as digital redundancy or digital error correction [14]. The conversion accuracy thus solely relies on the precision of the residue signals; the conversion speed, on the other hand, is largely determined by the settling speed of the residue amplifiers.

B. Power Efficiency Under Low Supply Voltage

While the scaling of CMOS technology offers a potential for improvement on the operating speed of mixed-signal circuits, the accompanying reduction in the supply voltage and various short-channel effects create both fundamental and practical limitations on the achievable gain, signal swing, and noise level, particularly under a low power constraint. Oxide reliability dictates the use of a low supply voltage, which limits the signal swing of analog circuits. It was derived in [10] that the power consumption of an SC circuit is inversely proportional to the supply voltage for a fixed dynamic range.

The sampling process inherent in SC circuits introduces kT/C noise at each pipeline stage when a residue voltage is captured. The sampled noise usually comprises two major contributions—the channel noise of the switches and the opamp noise. Since no direct current is conducted by the switch right before a sampling takes place (the circuit is assumed settled), the $1/f$ noise is not of concern here, and only the thermal noise contributes, which is a function of the channel resistance that is weakly affected by technology scaling [22]. On the other hand, the opamp output noise is mostly dominated by the channel noise of the input transistors, where the thermal noise and the $1/f$ noise both contribute. Because the input transistors are usually biased in saturation, velocity saturation and hot carrier effects tend to enhance the thermal noise level [23], [24]; the $1/f$ noise increases as well due to the reduced gate capacitance. It follows that, as CMOS technology scaling continues, the operational amplifier is becoming the dominant noise source for SC circuits.

C. Stage-Scaling Analysis of Pipeline ADC

Exploiting the fact that later stages contribute a diminishing input-referred noise because of the accumulative interstage gain, one architectural approach to maximize the signal-to-noise ratio (SNR) with a given power budget is to determine an optimum way of distributing the biasing current to each pipeline stage [25], [26], [15]. This involves choosing an optimum per-stage resolution and reducing the sampling capacitor sizes along the pipeline. Intuitively, a too low per-stage resolution (hence more stages) increases the number of residue resampling events. Coupled with a low interstage gain, this leads to larger capacitors and more biasing current. Conversely, although a high per-stage resolution (hence fewer stages) minimizes the number of resampling events and allows a rapid tapering of the capacitor size, the per-stage power increases exponentially due to the reduction of the feedback factor of residue amplifiers. This is clearly not power-efficient as well. It follows that the optimum stage resolution has to be somewhere in between, but the exact answer depends on the conversion speed, the technology used, the circuit topology, and a specific layout. Due to the complexity of this problem, simple hand analyses, albeit lacking in numerical accuracy, are commonly used to reveal qualitative parametric tradeoffs and to offer insight to circuit designers.

The first such analysis was introduced in [25], where the most important conclusion was that an optimum stage-scaling factor exists and is approximately equal to the interstage gain. In the Appendix, a new analysis is presented to optimize the per-stage resolution and the scaling factor concurrently, taking into account both the amplifier and the switch noises. One prominent feature of this analysis is the inclusion of the parasitic-loading effect on the residue amplifiers, manifested by two parasitic capacitors, C_o and C_g, at the output node and the summing node of the residue amplifier, respectively (Fig. 19). An accurate account of the parasitics requires a post-layout extraction that is design specific. At the architectural level, two observations are noteworthy in modeling these parasitics. First, to keep up the conversion speed, the transconductance of the opamp (hence its size and biasing current) must increase exponentially as a function of the stage resolution due to the drop of the feedback factor. Meanwhile, the number of comparators and sampling switches and the wiring complexity of layout exhibit the same dependence. Second, the parasitic-loading effect worsens as the conversion rate increases. In a high-speed converter, the parasitic capacitance can be comparable to the total value of the sampling capacitors. The model $C_o \approx C_g \approx \eta\,(2^n\gamma_i C_u)$ captures these dependences, where a technology-independent "speed factor" η is introduced to model the loading effect as a function of the conversion speed. The value of this "speed factor" varies between 0 and 1 in this analysis by observing the following facts.

1) When conversion speed is low, the loading due to parasitics is insignificant and the residue amplifier is mainly loaded by the sampling capacitors. This indicates a small "speed factor," hence η is set to 0 for this scenario.
2) When conversion speed is high, the parasitic-loading effect is severe. If a specific technology is given, an exact relationship between η and the sampling rate may be derived. To make this analysis more general and technology-independent, the maximum parasitic capacitance is set to be equal to the total sampling capacitors, i.e., $C_o \approx C_g \approx 2^n\gamma_i C_u$ or $\eta = 1$.

For a uniform per-stage resolution n and a uniform scaling factor γ, the total power consumption can be derived as (see the Appendix)

$$P \propto \text{SNR} \cdot kT \cdot f_s \cdot \left(\frac{V_{gs} - V_{th}}{V_{dd}}\right) \cdot g(n,\gamma,\eta) \qquad (1)$$

where kT is the thermal energy, f_s is the sampling rate, $(V_{gs} - V_{th})$ is the overdrive voltage of the opamp input transistors, and V_{dd} is the supply voltage.

Equation (1) supports the result of [10] with an addition of the new $g(.)$ function, which captures the dependence of the power consumption on the per-stage resolution n, the scaling factor γ, and the speed factor η. For a given speed factor, minimizing this function yields the optimum stage resolution and scaling factor. Important observations drawn from the study of this function are listed as follows.

1) For a small η or a low conversion speed, the analysis degenerates to that of [25]. A higher per-stage resolution always results in a lower total power consumption for a fixed resolution.
2) When η approaches one, i.e., in high-speed cases, the amplifier is increasingly loaded by parasitics. The total power consumption quickly rises when the stage resolution exceeds 3 b/s. The optimum resolution is 2 or 3 b/s. This is the typical scenario encountered in a practical design.
3) The optimum scaling factor is a weak function of the speed factor, i.e., for a given stage resolution, the optimum scaling factor is relatively constant for different conversion speeds. This is probably justified by the fact that the speed factor influences all stages uniformly and the interstage gain stays constant.

These observations serve as the design guidelines for pipeline ADCs when the tradeoff between SNR and power consumption is critical. Note that the derivation in the Appendix does not include the SHA noise if one is used. The uniform scaling factor throughout the pipeline may not be practical due to an increased layout effort. In addition, scaling of the last few stages may be difficult when they become too small. Lastly, although a uniform per-stage resolution helps to keep a design modular, increasing the resolution of the backend stages (a nonuniform scaling essentially) may result in more power savings.

III. Capacitor Error-Averaging

In a fully optimized pipeline ADC, random capacitor mismatch is the most important error source of nonlinearity. In the absence of a post-fabrication component trimming or calibration, the conversion accuracy is usually limited to about 10–12 bits.

A. Active Capacitor Error-Averaging

One approach to treat mismatch error is capacitor error-averaging (CEA). For a 1-b/s architecture, the error correction is performed by interchanging the roles of two sampling capacitors during the amplification phases [8], [9]. Two residue voltages that contain complementary errors are generated consecutively. A second opamp and two extra capacitors are used to obtain the average of the residue voltage pair (hence the name "active"). It has been shown that the first-order gain error resulted from capacitor mismatch is removed by averaging [8].

The active CEA (ACEA) is capable of realizing an excellent linearity with poorly matched capacitors. However, this is achieved at the cost of an added circuit complexity and more power consumption due to the residue resampling process. Assuming that the averaging amplifier generates an equal amount of noise while draining the same current as the residue amplifier, the ACEA leads to a power and area increase of four times for the same SNR compared to the conventional architecture (Table I). Moreover, the averaging amplifier and the residue resampling process impose additional nonlinearities on the signal path, compromising the overall linearity of the converter.

B. Passive Capacitor Error-Averaging

In this paper, a PCEA technique is used. Calculation and Monte Carlo simulation demonstrate a better averaging result achievable with this technique [7]. Fig. 3 shows the PCEA circuit diagram. In each clock cycle T, two residue voltages are

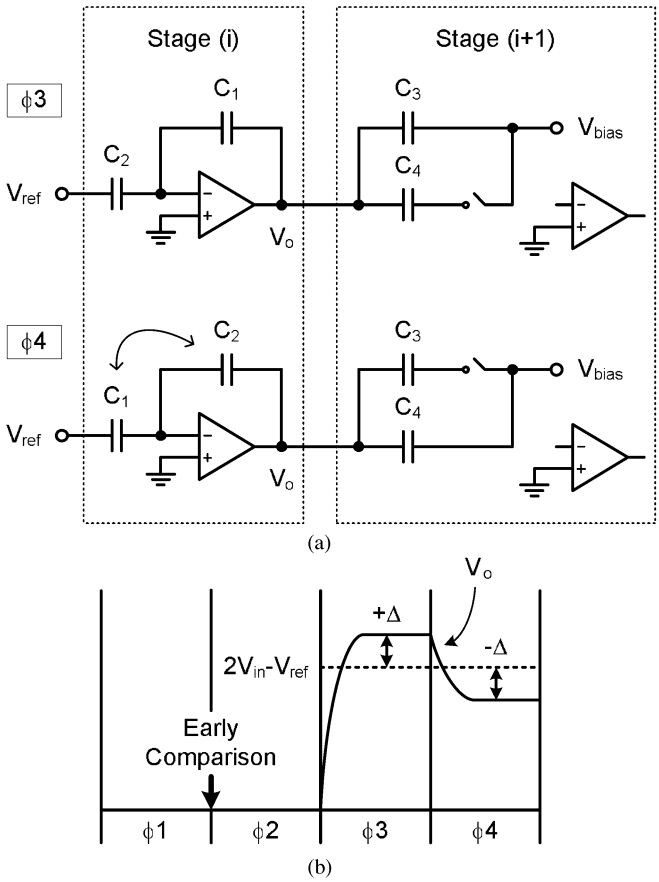

Fig. 3. Passive capacitor error-averaging technique. (a) Stage (i) in $\phi 3$ and $\phi 4$ phases. (b) Timing diagram for stage (i) with the early comparison marked at the end of $\phi 1$.

TABLE I
ARCHITECTURE POWER EFFICIENCY

Architecture (1.5-b/s)	Conventional	ACEA	PCEA
Averaging	No	Yes	Yes
Power	1	4	1/2
Speed	1	2/3	1/2
Power/speed	1	6	1

produced by interchanging the roles of the sampling capacitors C_1 and C_2. Two capacitors C_3 and C_4 from the next pipeline stage directly sample the residue voltage pair sequentially. The averaging amplifier is eliminated; averaging takes place when C_3 and C_4 merge charge in the subsequent phases. Avoiding the residue resampling process reduces the total conversion power by a factor of four. In addition, the uncorrelated noise in two separately acquired residue samples further increases the SNR by 3 dB. Furthermore, the loading of the residue amplifier is reduced because C_3 and C_4 never switch on at the same time.

A limitation of the PCEA is that it takes four clock phases to complete the sampling-comparison-amplification process, in contrast to two for a conventional implementation with no averaging or three for the ACEA (Table I). However, the extra phases allow a long settling time for the comparators if an early comparison is performed against the first residue as indicated in Fig. 3

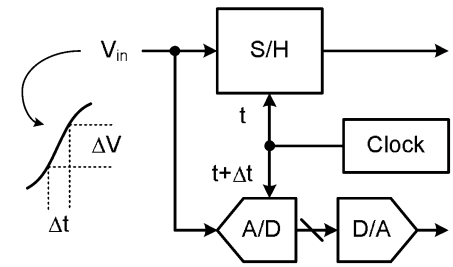

Fig. 4. Clock skew in the front-end sampling network.

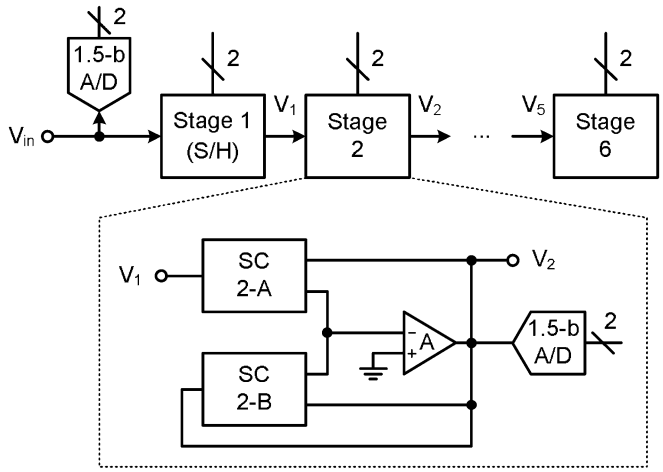

Fig. 5. Block diagram of the prototype 14-b pipeline ADC.

[7], [9]. In a typical scenario, this is as long as $T/4$ as opposed to just the nonoverlapping time in the conventional architecture with a two-phase clock.

Table I summarizes the power-saving results of the new averaging approach. For the same target SNR, the power efficiency is defined by normalizing the total conversion power to the sampling rate. It follows that the PCEA is as efficient as the conventional technique and six times more efficient than the ACEA.

IV. CIRCUIT IMPLEMENTATION

In order to demonstrate the effectiveness of the stage-scaling algorithm and the PCEA technique, a prototype 14-b pipeline ADC was designed and fabricated. This section describes the implementation details of this prototype and the next section presents the measurement results.

A. Sampling Clock Skew

In a typical pipeline ADC implementation, a dedicated SHA at the front-end is often used to enhance the dynamic performance of the converter. It mitigates the effect of the timing skew between the passive switched-capacitor sampler and the sub-ADC of the first stage. As indicated in Fig. 4, this aperture error effectively creates a dynamic offset between the two paths when the input signal exhibits a large slew rate. The offset ultimately results in a hard clipping error when the digital error-correction range of the subsequent stages is exceeded.

Nonetheless, the dedicated SHA is accompanied by a substantial power penalty; as it usually acts as a unity-gain buffer, the SHA provides no attenuation to the ADC noise referred to

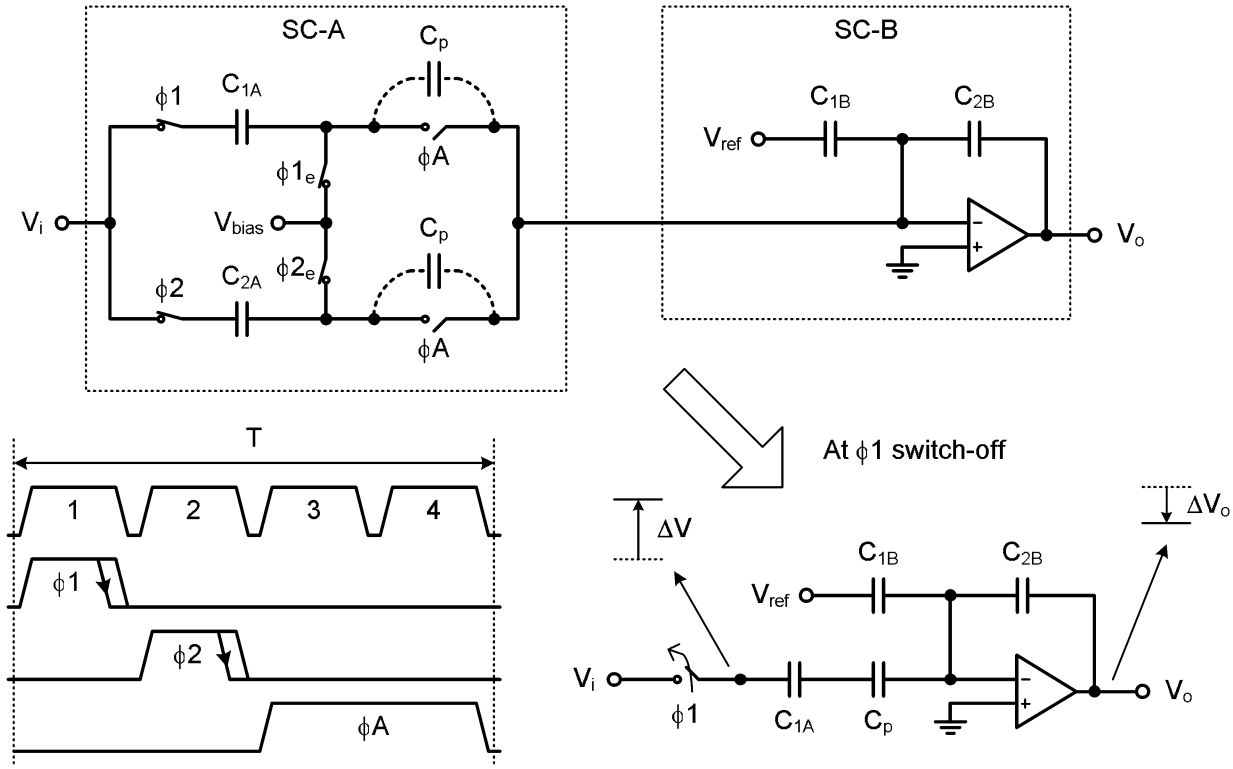

Fig. 6. Summing-node crosstalk.

the input, meanwhile adding its own contribution. Assuming that the SHA and the ADC contribute an equal amount of noise and consume an equal amount of power, then for the same target SNR, having a dedicated SHA translates into a fourfold increase in the total conversion power. This is too high a price for low-voltage designs that are SNR-limited. Notwithstanding, alternative solutions exist to remedy this problem. One approach is to use digital techniques to monitor the skew and adaptively compensate for it with a variable delay circuit. In this prototype, a simpler approach is adopted that exploits the large built-in digital redundancy of a 1.5-b/s architecture [26]. For example, it is straightforward to calculate that, for a 40-MHz full-scale sinusoidal input, the maximum clock skew tolerable to this architecture is one nanosecond. It is believed that the choice of a 1.5-b/s topology achieves a judicious architectural tradeoff and the best power efficiency in this prototype when the stage resolution and scaling, the SHA power penalty, the clock skew, and the averaging overhead are all taken into account. With this choice, the optimum stage-scaling factor is determined to be 0.5.

B. Opamp and Sub-ADC Sharing

Fig. 5 shows the single-ended equivalent block diagram of the prototype ADC employing the PCEA technique. The actual implementation is fully differential. The opamps and sub-ADCs are interleaved between the successive SC stages to save more power (e.g., SC 2-*A* and 2-*B* in Fig. 5). The 14-b ADC is partitioned into six pipeline stages with a total of six opamps and 14 comparators. Three bits are resolved by each stage effectively.

The opamp sharing technique was previously used in pipeline ADCs with an 8–10-b resolution [27], [20]. The challenge of this technique is to maintain the charge fidelity at the summing node, which is particularly difficult at the accuracy level of 14 bits. A potential crosstalk path between SC-*A* and SC-*B* (Fig. 6) arises due to a drain-to-source stray capacitor C_p of

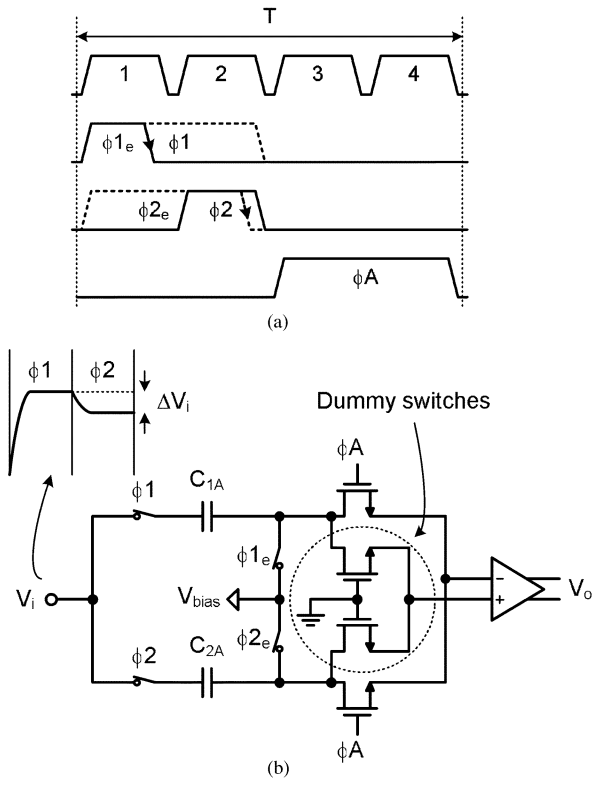

Fig. 7. Crosstalk eliminated by (a) modified timing and (b) dummy switch.

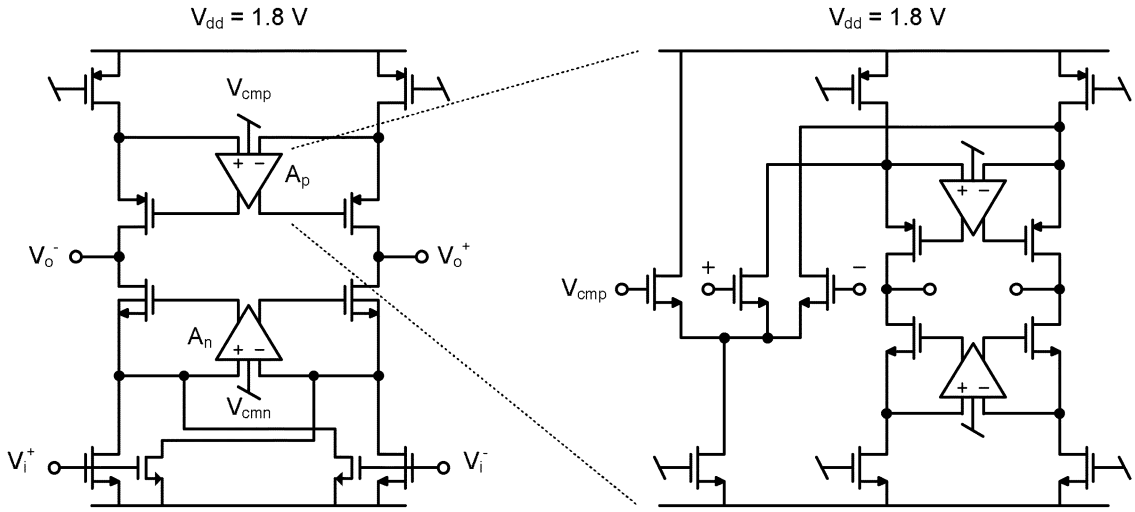

Fig. 8. Nested CMOS gain-boosting technique.

the off-switch ϕA (the same problem exists for ϕB, which is not shown). At the moment $\phi 1$ switches off, a signal-dependent charge injection induces an error voltage ΔV on the top plate of the capacitor C_{1A}; through the series connection of C_{1A}, C_p, and C_{2B}, it produces a small error voltage ΔV_o in the output residue, destroying the accuracy of SC-B. The effect of C_p, albeit small, can be significant at the 14-b level. The effect of this crosstalk can be gauged by the voltage gain through the coupling path, which is essentially the capacitor ratio C_p/C_{2B}.

A simple remedy to this problem is to tie the bottom plates of C_{1A} and C_{2A} (C_{1B} and C_{2B} as well) to an AC ground at all time, such that the coupling through C_p only results in a fixed offset error. This is accomplished for C_{2A} by simply advancing the rising edge of $\phi 2_e$ (the early phase of $\phi 2$) to the rising edge of $\phi 1$, as shown in Fig. 7(a). However, the same operation for C_{1A} is not possible because it has to take a sample at the end of $\phi 1$ and its bottom plate becomes floating afterwards. An alternative solution is to postpone the trailing edge of $\phi 1$ to the end of $\phi 2$. Since the switch-off of $\phi 1$ is delayed, no charge is injected onto the top plate of C_{1A} at the first place. Thus, the crosstalk is eliminated. Nonetheless, a careful examination of this solution still reveals a problem—it relies on the fact that V_i remains constant between $\phi 1$ and $\phi 2$, which is hardly guaranteed. In practice, V_i can be different between $\phi 1$ and $\phi 2$ due to a capacitor mismatch error in the previous pipeline stage. The difference between the dual residues, ΔV_i, albeit considerably smaller than the charge injection error, can still generate a significant error at the output. To mitigate this problem, gate-grounded dummy switches are introduced at the summing nodes [Fig. 7(b)]. Cross-coupled between the p- and n-sides of the virtual ground, they convert the residual crosstalk into a common-mode signal, which is rejected by the differential architecture. Combining all of these techniques, the magnitude of the resultant crosstalk can be estimated by the following equation:

$$\frac{\Delta V_o}{V_{FS}} \approx \frac{\Delta C}{C} \cdot \frac{C_p}{C_{2B}} \cdot \frac{\Delta C_p}{C_p} \approx \frac{1}{2^7} \cdot \frac{1}{2^8} \cdot 10\% \ll \frac{1}{2^{14}} \qquad (2)$$

where $\Delta C/C$ represents the capacitor mismatch error of the previous stage, $\Delta C_p/C_p$ corresponds to the cancellation accuracy of the cross-coupled dummy switches, and V_{FS} is the full-scale reference voltage. Even with the conservative estimates of (2), the aggregated attenuation of the crosstalk is large enough to ensure a 14-b accuracy at the summing node.

C. Nested CMOS Gain Boosting

Delivering sufficient dc gain at a high sampling rate with low power dissipation is a difficult challenge for opamp design at a low supply voltage. Although a multistage architecture offers high open-loop gain, the necessity of frequency compensation makes it power-inefficient. Single-stage architectures, on the other hand, offer large gain–bandwidth products with limited dc gain due to the low output resistance of short-channel devices. The CMOS gain-boosting technique was previously introduced to enhance the output resistance of a single-stage operational transconductance amplifier (OTA) [28]. In this design, a nested gain-boosting technique is used. As illustrated in Fig. 8, a two-level recursive boosting with devices of 0.2-μm gate length results in a minimum open-loop gain of 130 dB across process corners in simulation, exceeding the accuracy requirement of a 14-b converter. A manual analysis of the dynamics of the nested feedback loop is difficult; the simulation method proposed in [29] was used instead to verify the stability.

To maintain high current efficiency and large output swing simultaneously, the main amplifier uses a pseudo-differential (PD) architecture, as shown in Fig. 8. With four transistors in a stack, the peak-to-peak output swing of the opamp exceeds 2 V with a supply voltage of 1.8 V. The boosting amplifiers all use a folded-cascode structure with a p- or n-type input differential pair to allow a flexible input CM range. The nested boosters are the scaled version of the main boosters. The current ratios among the main amplifier, the gain boosters, and the nested boosters are 64:8:1. The latter two also share a common bias circuit.

D. Discrete-Time CM Regulation

The amplifier and switch-induced offset voltages are problematic in PD pipeline architectures. Without compensation, the

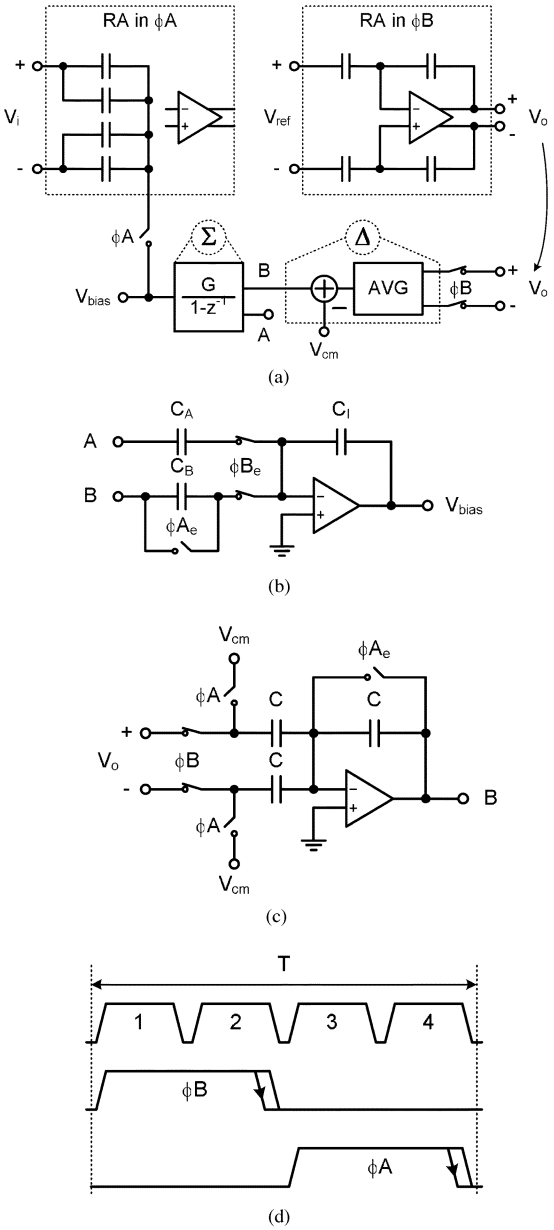

Fig. 9. (a) Diagram of the $\Delta\Sigma$ common-mode voltage regulation scheme. (b) Discrete-time integrator with the look-ahead capacitor C_A. (c) Averaging and differencing circuit. (d) Timing diagram.

offset will quickly accumulate and saturate the usable signal swing due to the large interstage gain of the pipeline. Reverting to a fully differential topology with instantaneous CM feedback (CMFB) every a few stages was suggested to break the offset propagation [19], resulting in a compromised hybrid design. In this paper, an analog $\Delta\Sigma$ loop that facilitates the CM control is introduced. Shown in Fig. 9, an averaging circuit derives the output CM voltage and compares it to the desired reference (Δ); the resultant error voltage is then accumulated with a discrete-time integrator (Σ) and fed back to be the bottom-plate bias for the S/H circuit. When the loop settles, the long-term average of the integrator input has to be zero, which forces the output CM to equal the reference voltage. The SC circuits realizing the Δ and Σ blocks are also shown in Fig. 9.

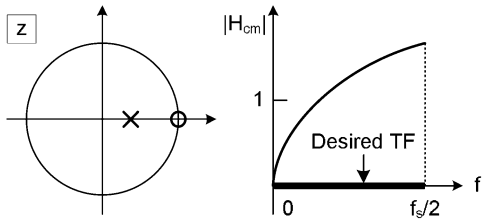

Fig. 10. Pole-zero plot and frequency response of the CMFB loop.

One key design aspect of the CM regulation loop is to ensure its stability. Assuming an ideal integrator and ignoring parasitic capacitance, the z-domain closed-loop CM voltage transfer function (TF) can be derived as

$$H_{cm}(z) = \frac{V_{o,cm}}{V_{i,cm}} = \frac{1}{\beta} \cdot \frac{1 - z^{-1}}{1 - \left(1 - \frac{2G}{\beta}\right) \cdot z^{-1}} \tag{3}$$

where β is the feedback factor of the main amplifier and $G = C_B/C_I$ is the integrator gain. The pole-zero location and the frequency response of this function are plotted in Fig. 10. The high-pass nature resembles the characteristic noise-shaping function of a sigma-delta modulator. To ensure stability, the condition $0 \leq G/\beta \leq 1$ must be satisfied. Because G and β are determined by capacitor ratios, stability can be guaranteed over process corners. However, stability is not sufficient to make this scheme fully functional. For a small integrator gain G, the magnitude of the closed-loop TF near the Nyquist frequency is approximately $1/\beta$, which is typically greater than one. It follows that a high-frequency CM variation will still be amplified. The desired closed-loop gain should be as close to zero as possible at all frequencies. This is accomplished by introducing a look-ahead path in the integrator as shown in Fig. 9. The modified TF becomes

$$H_{cm}(z) = \frac{1}{\beta} \cdot \frac{\left(1 - 2\beta \cdot \frac{C_A}{C_I}\right) \cdot (1 - z^{-1})}{1 - \left(1 - \frac{2}{\beta} \cdot \frac{C_B}{C_I}\right) \cdot z^{-1}}. \tag{4}$$

The condition $2\beta \cdot C_A/C_I = 1$ sets $H_{cm}(z) = 0$ identically for all frequencies. Again, this is determined by capacitor ratios and insensitive to process variations.

It is well known that the finite open-loop gain leads to a leakage problem in discrete-time integrators. However, this is not of concern in this design as it effectively only introduces a small offset in setting the output CM. As a result, single-transistor amplifiers are used in the integrator and the averaging circuit. Computer simulation reveals a systematic offset of approximately 10 mV resulted from the finite-gain effect, which is small compared to the full-scale output swing.

E. Dynamic Comparator

The 1.5-b/s pipeline architecture greatly relaxes the offset tolerance of the comparators. In addition, an early comparison is performed that exploits the dual residue feature of the CEA technique [7], [9]. This allows the comparators a complete quarter clock cycle to resolve the digital code. As a result, the comparator design is quite relaxed; dynamic comparators with minimum size devices are used. The schematic diagram of the

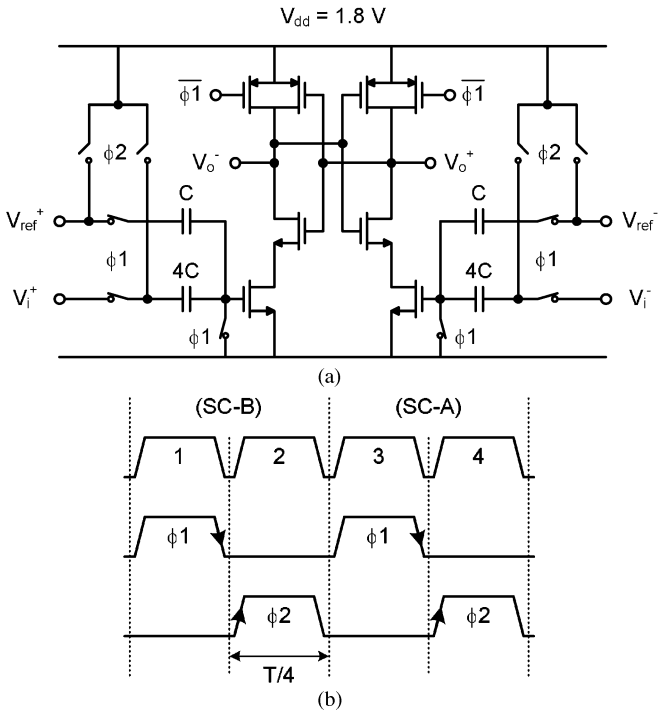

Fig. 11. (a) Dynamic comparator. (b) Comparator timing diagram.

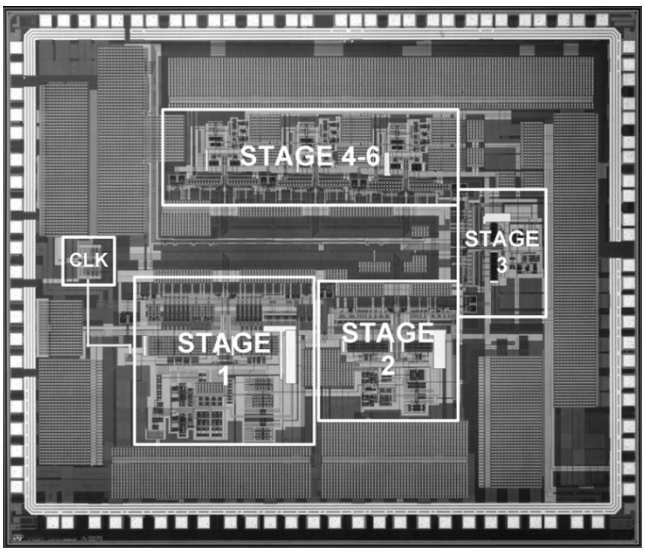

Fig. 12. Die photo.

comparator is shown in Fig. 11. The comparison threshold is determined by the size ratio of the sampling capacitors, which is 4 : 1 in this design.

F. Sampling Switch

The signal-dependent charge injection and on-resistance variation of switches pose fundamental limits to the achievable distortion levels of SC circuits. The increased f_T through technology scaling improves the switch performance. In this study, the use of minimum channel-length devices combined with clock bootstrapping [10] and in-line switch [11] techniques resulted in an outstanding SFDR of 97 dB with a full-scale 40-MHz input [21].

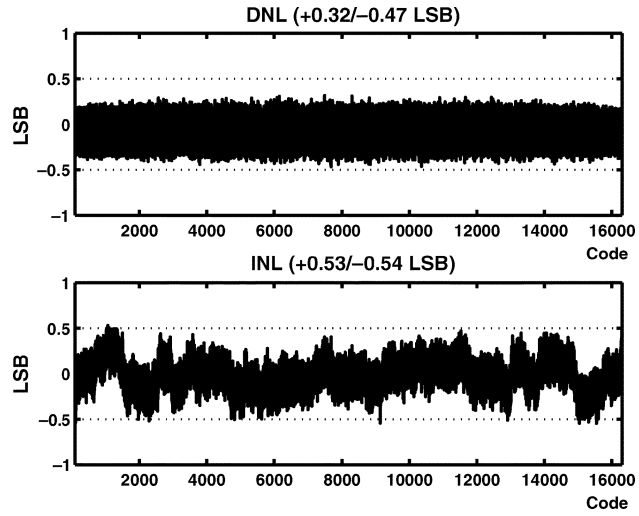

Fig. 13. Measured DNL and INL ($f_s = 12$ MS/s, $f_{\rm in} = 1$ MHz).

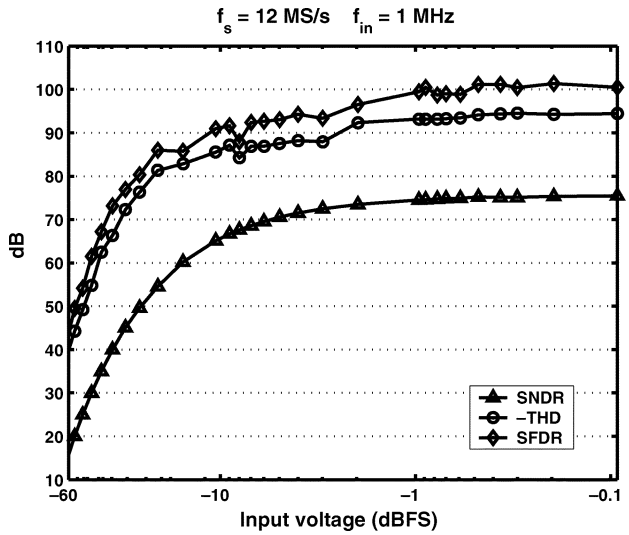

Fig. 14. Measured ADC performance versus input signal level.

V. Experimental Results

The prototype ADC was fabricated in a 1.8-V 0.18-μm six-metal one-poly digital CMOS process. Capacitors were implemented using metal–insulator–metal (MIM) structures with no special attention paid in layout to match them. A die photo is shown in Fig. 12. The size of the chip measures 4.3×3.5 mm^2 with the ADC occupying approximately 10 mm^2.

The differential nonlinearity (DNL) and the integral nonlinearity (INL) were measured using a code-density test with a 1-MHz full-scale sinusoidal input. At 12 MS/s, eight million samples were collected. The measured DNL and INL profiles are shown in Fig. 13. The maximum DNL is 0.47 LSB and the maximum INL is 0.54 LSB.

The dynamic linearity of the ADC was characterized by analyzing a fast Fourier transform (FFT) of the output codes with a single-tone input. Shown in Fig. 14, the measured peak SNDR reaches 75.5 dB with a 1-MHz input, equivalent to 12.25 effective number of bits (ENOB). Under the same condition, the peak total harmonic distortion (THD) and the peak SFDR are −94.5 dB and 101 dB, respectively (the THD figure corresponds

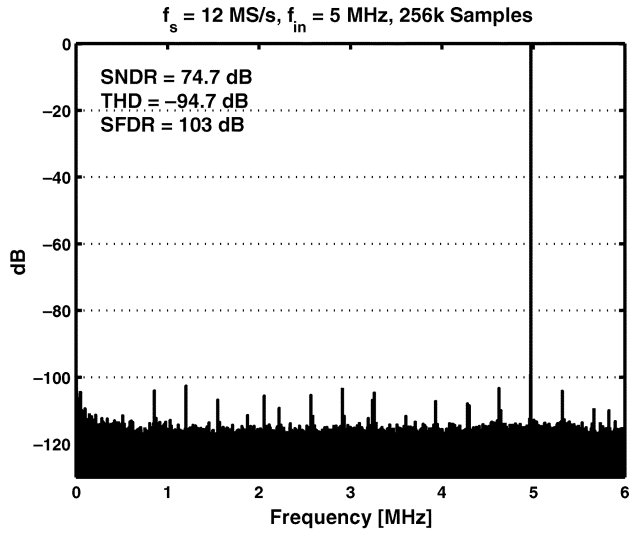

(a)

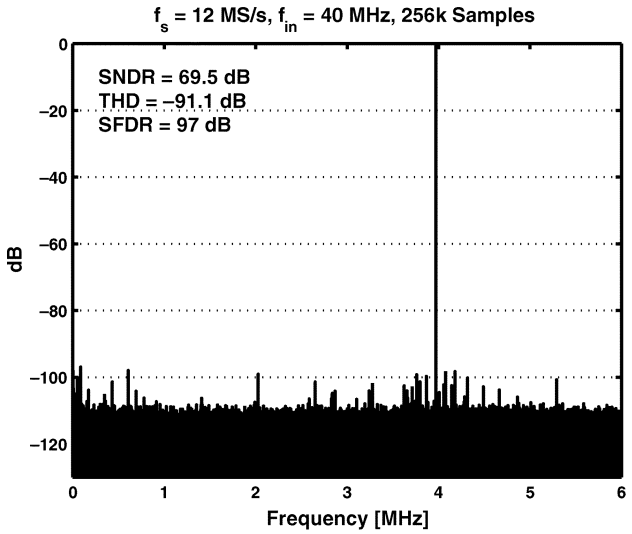

(b)

Fig. 15. Measured FFT spectrum at (a) $f_{\rm in} = 5$ MHz and (b) $f_{\rm in} = 40$ MHz.

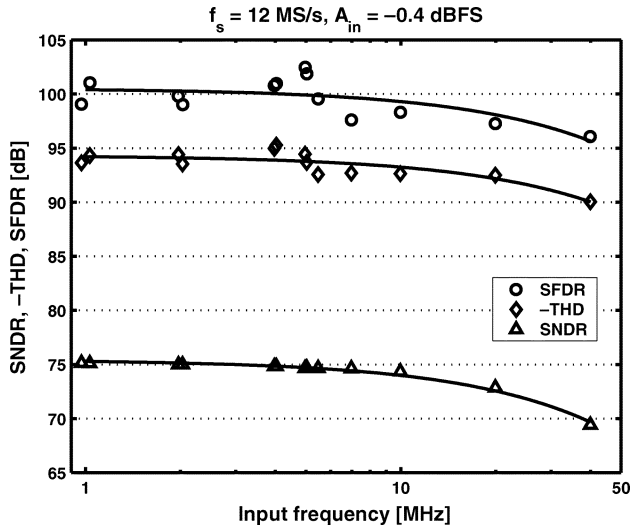

Fig. 16. Measured ADC dynamic performance.

TABLE II
MEASURED ADC PERFORMANCE (1.8 V, 25 °C)

Resolution	14 bits	
Reference voltage	0.4 V and 1.4 V	
Packaging	QFP100	COB
Sampling rate	10 MS/s	12 MS/s
DNL @ 1 MHz	-0.31/0.31 LSB	-0.47/0.32 LSB
INL @ 1 MHz	-0.58/0.53 LSB	-0.54/0.53 LSB
Peak SNDR	73.6 dB	75.5 dB
Peak SFDR	99 dB	103 dB
SFDR @ 40 MHz	84 dB	97 dB
Power	112 mW	98 mW
Technology	0.18-μm 6M-1P CMOS	

to the power sum of the first 15 harmonics). Fig. 15(a) shows the measured FFT spectrum with a −0.4-dBFS 5-MHz input. The SFDR in this case reaches the value of 103 dB, and the SNDR is 74.7 dB. To measure the input analog bandwidth and to verify the architecture choice without a dedicated SHA, the ADC was also tested with a −0.4-dBFS 40-MHz sine-wave subsampled at 12 MS/s. Fig. 15(b) shows the digital spectrum of the ADC output. In this case, the measured SNDR and SFDR are 69.5 and 97 dB, respectively. Fig. 16 summarizes the measured dynamic performance of this 14-b ADC with an input frequency span from 1 to 40 MHz. The random jitter accumulated during the generation and distribution of the clock signal limits the SNR performance at high frequencies. A locked histogram test revealed a 1.5-ps rms jitter in the system including the clock generator, the synthesizer, the ADC chip, and the board, which translates to a 70-dB SNR at approximately 40 MHz [30]. This confirms the observation that the performance of this converter is limited by the clock jitter at high input frequencies.

The performance sensitivity against the supply and CM voltages were also verified. The measurement results are summarized in Fig. 17. The minimum supply voltage at which this ADC still works without noticeable performance degradation is 1.65 V. The total power consumption of the chip is 97.7 mW, excluding the LVDS digital output drivers. Out of this, 95.4 mW is consumed by the analog circuits, 1.4 mW is consumed by the digital circuits, and 0.9 mW goes to the clock buffer.

All measurements were performed with a 1.8-V supply at room temperature (25 °C). Table II summarizes the measurement results of the prototype ADC.

VI. CONCLUSION

A 14-b pipeline ADC in 0.18-μm CMOS technology using exclusively thin oxide transistors and a low supply voltage of 1.8 V is demonstrated. The prototype design achieves a true 14-b linearity and a 12.25 ENOB in experiments without trimming, calibration, or dithering. This study demonstrates the feasibility of sustained scaling of high-resolution CMOS ADCs in the deep-submicron regime and the potential improvement on performance harvestable through technology scaling.

As a reference, Fig. 18 shows the comparison of this design (marked by a square) and the previously reported high-resolution ADCs (marked by diamonds) with a 12-b and higher res-

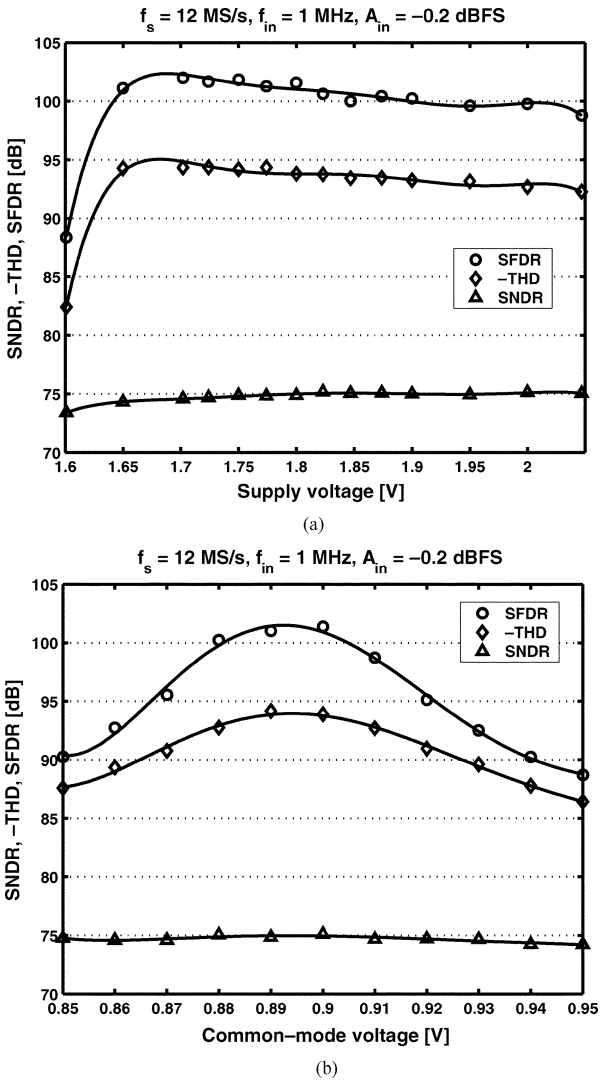

Fig. 17. Measured ADC performance versus (a) supply voltage and (b) CM voltage.

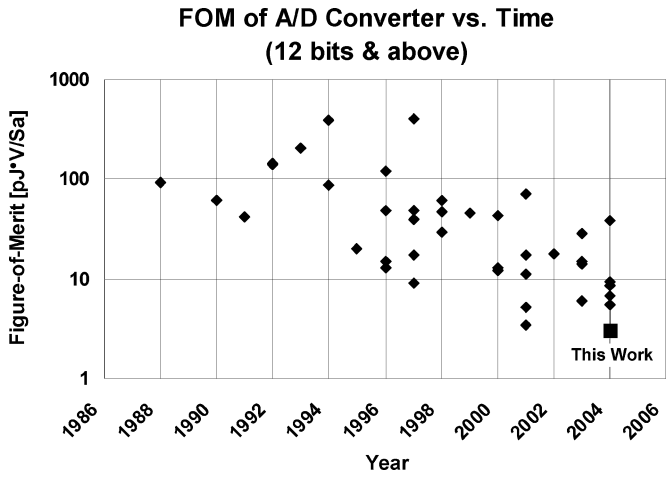

Fig. 18. Comparison of this prototype (square) and previously published ADCs (diamonds).

olution dated from 1988 to 2004. For SNR-limited designs, a more appropriate figure of merit (FOM) used in this comparison is defined as

$$\text{FOM} = \frac{\text{Power}}{2^{\text{ENOB}} \cdot f_s} \cdot V_{dd}. \tag{5}$$

This is in accordance with (1) that the conversion power is inversely proportional to the supply voltage, as manifested by the normalization to the supply voltage in (5). This 14-b pipeline ADC has achieved the lowest FOM in this category of Nyquist converters.

APPENDIX
STAGE-SCALING ANALYSIS OF PIPELINE ADC

The circuit diagram of two consecutive MDAC stages of a typical pipeline ADC is shown in Fig. 19, with the ith stage sampling the output of the $(i-1)$th stage. The per-stage resolution n is assumed constant. The unit sampling capacitor of the ith stage is $\gamma_i C_u$, where C_u is the unit capacitor of the first stage and γ_i is the scaling factor of the ith stage. The assumption is that all circuit components, including the capacitor size, the amplifier g_m, and the biasing current, are scaled. The noise sampled on C_L comprises the opamp noise (g_m block) and the switch noises (r_1 through r_4). The parasitic-loading effect is modeled by two capacitors, C_o and C_g, at the output node and the summing node of the residue amplifier, respectively. The model $C_o \approx C_g \approx \eta\,(2^n \gamma_i C_u)$ captures the exponential dependence of the loading effect on the stage resolution n, where a "speed factor" η is introduced varying between 0 and 1 in this analysis. The assumption $C_o \approx C_g$ is approximate to make the analysis tractable. A feedforward factor can be defined as

$$\alpha = \frac{C_s}{C_s + C_f + C_g} = \frac{1}{1+\eta}\left(1 - \frac{1}{2^n}\right) \tag{6}$$

and a feedback factor as

$$\beta = \frac{C_f}{C_s + C_f + C_g} = \frac{1}{1+\eta} \cdot \frac{1}{2^n}. \tag{7}$$

The noise sampled on C_L is usually taken as kT/C_L without a specific account of the noises originating from the amplifier and the switches [25]. Assuming a first-order frequency response, the g_m noise floor at the output of the $(i-1)$th stage is $(1/\beta)^2 4kT(N_{\text{op}}/g_m)$, where N_{op} is the noise factor of the amplifier. If the input transistor is the only noise contributor in a single-ended amplifier, $N_{\text{op}} = 2/3$ holds for long-channel devices. Nonetheless, N_{op} can be substantially larger than 2/3 if the noise model for short-channel devices is used [23], [24]. In addition, the current source and the cascode devices also contribute noise. In this analysis, $N_{\text{op}} = 3$ is assumed for single-ended opamps.

The noise transfer functions of the switches r_1 through r_4 can be calculated individually. As the switch noises are band-limited by the amplifier, their contribution is reduced compared to the scenario where the circuit bandwidth is only determined by the low-pass filter formed by the switch and the sampling capacitors. The exact solution, taking into account the effect of

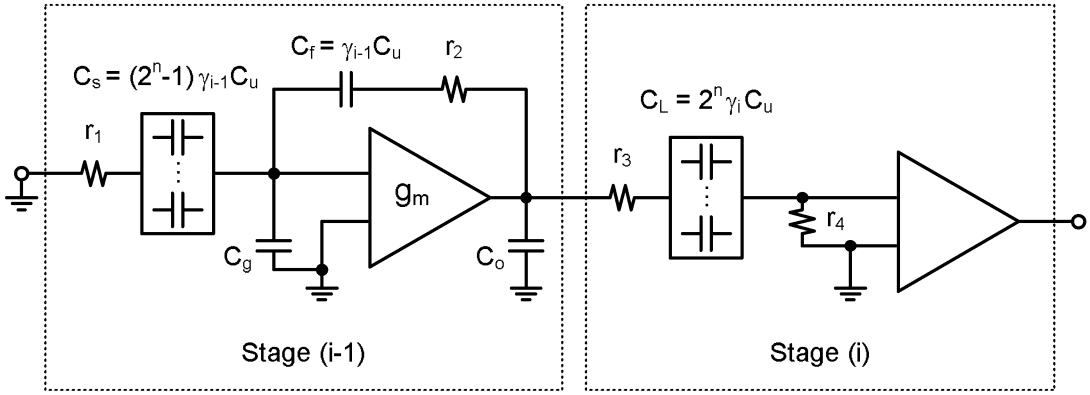

Fig. 19. Residue amplifier model used in the stage-scaling analysis.

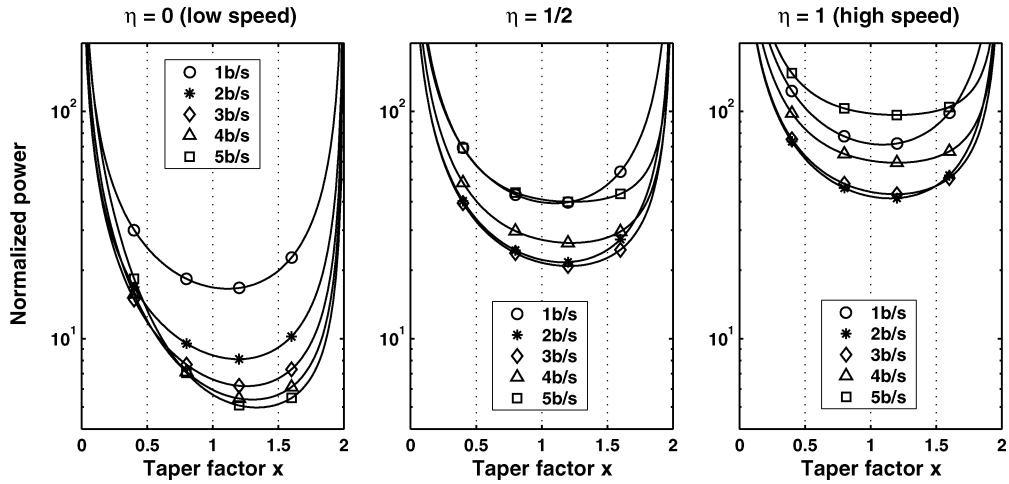

Fig. 20. Evaluation of $g(n,\gamma,\eta)$ versus the taper factor x.

the switch resistance on the frequency response, is quite involved. Instead of an accurate account of these second-order effects, the total noise floor at the output of the $(i-1)$th stage is approximated as

$$N_i(f) = 4kT \cdot \left[\left(\frac{1}{\beta}\right)^2 \frac{N_{\text{op}}}{g_m} + \left(\frac{\alpha}{\beta}\right)^2 r_1 + r_2 + (r_3 + r_4)\right] \quad (8)$$

where the first term stems from the opamp noise and the rest from the switches. The frequency response is assumed first-order, band-limited by the closed-loop bandwidth ($\omega_{-3\text{ dB}}$) of the residue amplifier. In addition, the switches are sized such that they will not limit the settling speed of the amplifier. For this reason, assume

$$r_1 C_s, r_2 C_f, (r_3 + r_4) C_L \le \frac{1}{5}\frac{1}{\omega_{-3\text{ dB}}} = \frac{1}{5}\frac{C_T}{\beta \cdot g_m} \quad (9)$$

where C_T is the total output load capacitance of the residue amplifier, given by

$$C_T = \left[1 - \frac{1}{(1+\eta)\cdot 2^n} + \eta \cdot 2^n\right] \gamma_{i-1} C_u + 2^n \gamma_i C_u. \quad (10)$$

Combining (8) and (9), we have

$$N_i(f) = \frac{4kT}{5g_m} \cdot \left[\frac{5N_{\text{op}}}{\beta^2} + \left(\frac{\alpha}{\beta}\right)^2 \frac{C_T}{\beta C_s} + \frac{C_T}{\beta C_f} + \frac{C_T}{\beta C_L}\right]. \quad (11)$$

The input-referred integrated noise sampled by the ith stage is

$$N_i = \frac{1}{(4^n)^{i-1}}\left[\frac{\pi}{2} N_i(f)\text{BW}\right] = \frac{1}{(4^n)^{i-1}} \frac{kT}{5C_u} \cdot \left[\frac{(1+\eta)2^n 5N_{\text{op}}}{\frac{C_T}{C_u}} + \frac{2^n}{\gamma_{i-1}} + \frac{1}{2^n \gamma_i}\right]. \quad (12)$$

If a uniform scaling is assumed, i.e., $\gamma_i = \gamma^{i-1}$ for $i = 1, \ldots, \infty$, $\sum_i N_i$ is summable for $\gamma > 1/4^n$ and is given by

$$\sum_{i=1}^{\infty} N_i = \frac{kT}{2^n C_u} + \frac{kT}{5C_u}\left(\frac{1}{4^n\gamma - 1}\right) \times \left[\frac{(1+\eta)\cdot 2^n \gamma \cdot 5N_{\text{op}}}{1 - \frac{1}{(1+\eta)\cdot 2^n} + \eta \cdot 2^n + 2^n\gamma} + 2^n\gamma + \frac{1}{2^n}\right] \quad (13)$$

where the noise from the first stage is $kT/2^nC_u$ because no front-end SHA is assumed.

In addition, if slewing is ignored, the settling speed is determined by the small-signal closed-loop bandwidth ($\omega_{-3\ \mathrm{dB}}$) of the residue amplifier. The total conversion power can be derived as

$$P \propto SNR \cdot kT \cdot f_s \cdot \left(\frac{V_{gs}-V_{th}}{V_{dd}}\right) \cdot g(n,\gamma,\eta) \qquad (14)$$

where f_s is the sampling rate, $V_{gs}-V_{th}$ is the overdrive voltage of the amplifier input transistor, and V_{dd} is the supply voltage. Function $g(.)$ is given by

$$\begin{aligned} &g(n,\gamma,\eta) \\ &= \left(\frac{1+\eta}{1-\gamma}\right)\left(2^n - \frac{1}{1+\eta} + \eta\cdot 4^n + 4^n\gamma\right) \\ &\quad \times \left\{\frac{1}{2^n} + \frac{\frac{1}{5}}{4^n\gamma - 1}\left[\frac{(1+\eta)\cdot 4^n\gamma \cdot 5N_{\mathrm{op}}}{2^n - \frac{1}{1+\eta} + \eta\cdot 4^n + 4^n\gamma} + 2^n\gamma + \frac{1}{2^n}\right]\right\}. \end{aligned} \qquad (15)$$

A "taper factor" x is defined in [25], which relates to the scaling factor by the equation $\gamma = 1/2^{nx}$. In Fig. 20, $g(.)$ is plotted against x for η =0, 1/2, 1 (corresponding to different conversion speeds) and $n = 1 \ldots 5$. The plot indicates that the speed factor has a significant impact on the optimum stage resolution. A 2–3-b/s architecture is in general preferred for high-speed applications.

ACKNOWLEDGMENT

The authors would like to thank ST Microelectronics for providing the chip fabrication, and the Data Conversion Systems Group of National Semiconductor Corporation for help with testing. The authors are also grateful to Prof. B. Brodersen for providing support on experimental setup for the later evaluation of the prototype chip at the Berkeley Wireless Research Center (BWRC), UC Berkeley. The contributions of J. Vanderhaegen and S. Wang for improving the manuscript are also acknowledged.

REFERENCES

[1] T. L. Brooks, D. H. Robertson, D. F. Kelly, A. Del Muro, and S. W. Harston, "A cascaded sigma-delta pipeline A/D converter with 1.25-MHz signal bandwidth and 89-dB SNR," *IEEE J. Solid-State Circuits*, vol. 32, pp. 1896–1906, Dec. 1997.

[2] A. R. Feldman, B. E. Boser, and P. R. Gray, "A 13-bit, 1.4-MS/s sigma-delta modulator for RF baseband channel applications," *IEEE J. Solid-State Circuits*, vol. 33, pp. 1462–1469, Oct. 1998.

[3] J. C. Rudell, J. A. Weldon, J. J. Ou, L. Lin, and P. R. Gray, "An integrated GSM/DECT receiver: Design specifications," Univ. California Berkeley Electronics Research Laboratory Memorandum, Memo UCB/ERL M97/82.

[4] T. Gratzek, B. Brannon, J. Camp, and F. Murden, "A new paradigm for base station receivers: High IF sampling+digital filtering," in *Proc. IEEE Radio Frequency Integrated Circuits (RFIC) Symp.*, June 1997, pp. 143–146.

[5] R. Schreier, J. Lloyd, L. Singer, D. Paterson, M. Timko, M. Hensley, G. Patterson, K. Behel, and J. Zhou, "A 10–300-MHz IF-digitizing IC with 90–105-dB dynamic range and 15–333-kHz bandwidth," *IEEE J. Solid-State Circuits*, vol. 37, pp. 1636–1644, Dec. 2002.

[6] W. H. W. Tuttlebee, "Advances in software-defined radio," *Electron. Syst. Software*, vol. 1, pp. 26–31, 2003.

[7] Y. Chiu, "Inherently linear capacitor error-averaging techniques for pipelined A/D conversion," *IEEE Trans. Circuits Syst. II*, vol. 47, pp. 229–232, Mar. 2000.

[8] 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.

[9] H.-S. Chen, B.-S. Song, and K. Bacrania, "A 14-b 20-Msamples/s CMOS pipelined ADC," *IEEE J. Solid-State Circuits*, vol. 36, pp. 997–1001, June 2001.

[10] A. M. Abo and P. R. Gray, "A 1.5-V, 10-bit, 14.3-MS/s CMOS pipeline analog-to-digital converter," *IEEE J. Solid-State Circuits*, vol. 34, pp. 599–606, May 1999.

[11] H. Pan, M. Segami, M. Choi, L. Cao, and A. A. Abidi, "A 3.3-V 12-b 50-MS/s A/D converter in 0.6-μm CMOS with over 80-dB SFDR," *IEEE J. Solid-State Circuits*, vol. 35, pp. 1769–1780, Dec. 2000.

[12] S. Sutardja 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.

[13] 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.

[14] S. H. Lewis, H. S. Fetterman, G. F. Gross Jr., R. Ramachandran, and T. R. Viswanathan, "A 10-b 20-Msample/s analog-to-digital converter," *IEEE J. Solid-State Circuits*, vol. 27, pp. 351–358, Mar. 1992.

[15] T. 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.

[16] S.-U. Kwak, B.-S. Song, and K. Bacrania, "A 15-b, 5-Msample/s low-spurious CMOS ADC," *IEEE J. Solid-State Circuits*, vol. 32, pp. 1866–1875, Dec. 1997.

[17] Y.-I. Park, S. Karthikeyan, F. Tsay, and E. Bartolome, "A 10-b 100-Msample/s CMOS pipelined ADC with 1.8-V power supply," in *Proc. IEEE Int. Solid-State Circuits Conf.*, Feb. 2001, pp. 130–131.

[18] W. Yang, D. Kelly, L. Mehr, M. T. Sayuk, and L. Singer, "A 3-V 340-mW 14-b 75-Msample/s CMOS ADC with 85-dB SFDR at Nyquist input," *IEEE J. Solid-State Circuits*, vol. 36, pp. 1931–1936, Dec. 2001.

[19] D. Miyazaki, M. Furuta, and S. Kawahito, "A 16-mW 30-Msample/s 10-b pipelined A/D converter using a pseudo-differential architecture," in *Proc. IEEE Int. Solid-State Circuits Conf.*, Feb. 2002, pp. 174–175.

[20] B.-M. Min, P. Kim, F. W. Bowman, D. M. Boisvert, and A. J. Aude, "A 69-mW 10-bit 80-Msample/s pipelined CMOS ADC," *IEEE J. Solid-State Circuits*, vol. 38, pp. 2031–2039, Dec. 2003.

[21] Y. Chiu, P. R. Gray, and B. Nikolic, "A 1.8-V, 14-b, 10-MS/s pipelined ADC in 0.18-μm CMOS with 99-dB SFDR," in *Proc. IEEE Int. Solid-State Circuits Conf.*, vol. 47, Feb. 2004, pp. 458–459.

[22] J. M. Rabaey, A. Chandrakasan, and B. Nikolic, *Digital Integrated Circuits: A Design Perspective*. Upper Saddle River, NJ: Pearson Education, 2003.

[23] C. Enz and Y. Cheng, "MOS transistor modeling for RF IC design," *IEEE J. Solid-State Circuits*, vol. 35, pp. 186–201, Feb. 2000.

[24] A. A. Abidi, "High-frequency noise measurements on FET's with small dimensions," *IEEE Trans. Electron Devices*, vol. ED-33, pp. 1801–1805, Nov. 1986.

[25] D. W. Cline and P. R. Gray, "A power optimized 13-b 5-Msamples/s pipelined analog-to-digital converter in 1.2-μm CMOS," *IEEE J. Solid-State Circuits*, vol. 31, pp. 294–303, Mar. 1996.

[26] T. Cho, "Low-power low-voltage analog-to-digital conversion techniques using pipelined architectures," Ph.D. dissertation, Univ. California Berkeley, 1995.

[27] K. Nagaraj, H. S. Fetterman, J. Anidjar, S. H. Lewis, and R. G. Renninger, "A 250-mW, 8-b, 52-Msamples/s parallel-pipelined A/D converter with reduced number of amplifiers," *IEEE J. Solid-State Circuits*, vol. 32, pp. 312–320, Mar. 1997.

[28] K. Bult and G. J. G. M. Geelen, "A fast-settling CMOS op amp for SC circuits with 90-dB DC gain," *IEEE J. Solid-State Circuits*, vol. 25, pp. 1379–1384, June 1990.

[29] P. J. Hurst and S. H. Lewis, "Determination of stability using return ratios in balanced fully differential feedback circuits," *IEEE Trans. Circuits Syst.*, vol. 42, pp. 805–817, Dec. 1995.

[30] M. Shinagawa, Y. Akazawa, and T. Wakimoto, "Jitter analysis of high-speed sampling systems," *IEEE J. Solid-State Circuits*, vol. 25, pp. 220–224, Jan. 1990.

Yun Chiu (S'97) received the B.S. degree in physics from the University of Science and Technology of China, Hefei, in 1993 and the M.S. degree in electrical engineering from the University of California, Los Angeles (UCLA), in 1997. He is currently working toward the Ph.D. degree at the University of California, Berkeley.

From 1995 to 1997, he was a Graduate Student Researcher with the Optoelectronic Circuits and Systems Laboratory, UCLA, where he worked on the broadband linearization of CMOS driver ICs for Mach–Zehnder optical modulator using square-law predistortion techniques. In 1997, he joined CondorVision Technology, Inc. (later Pixart Technology, Inc.), Fremont, CA, where he was a Senior IC Design Engineer in charge of developing A/D converters for CMOS digital imaging products. Since 1999, he has been with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley. From 2000 to 2003, he was also employed by the *Cal View* program (the Televised Instruction Program of the College of Engineering at Berkeley) as the Course Consultant for several graduate-level IC design classes. His current research interests include high-performance, low-power, and low-voltage A/D converters for wireless and wire-line communications, and adaptive digital-signal-processing techniques for improving performance of analog and mixed-signal circuits. He holds one U.S. patent and is the author or coauthor of several technical articles.

Mr. Chiu was the recipient of the Foreign Scholar Award at UCLA in 1994, the Regents' Fellowship at Berkeley in 1999, the Intel Ph.D. Fellowship in 2001, and the *Cal View* Teaching Fellow Award in 2003.

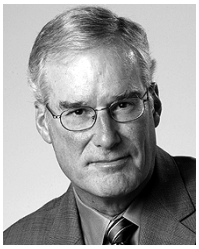

Paul R. Gray (M'69–SM'76–F'80) was born in Jonesboro, AR, on December 8, 1942. He received the B.S., M.S., and Ph.D. degrees from the University of Arizona, Tucson, in 1963, 1965, and 1969, respectively.

In 1969, he joined the Research and Development Laboratory, Fairchild Semiconductor, Palo Alto, CA, where he was involved in the application of new technologies for analog integrated circuits, including power integrated circuits and data conversion circuits. In 1971, he joined the faculty of the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley (UC Berkeley). His research interests have included bipolar and MOS circuit design, electro-thermal interactions in integrated circuits, device modeling, telecommunications circuits, and analog–digital interfaces in VLSI systems. He is the coauthor of a widely used college textbook on analog integrated circuits. During yearlong industrial leaves of absence from UC Berkeley, he served as Project Manager for Telecommunications Filters at Intel Corporation, Santa Clara, CA, in 1977–1978, and as Director of CMOS Design Engineering at Microlinear Corporation, San Jose, CA, in 1984–1985. He has held several administrative posts at UC Berkeley, including Director of the Electronics Research Laboratory (1985–1986), Vice Chairman of the Electrical Engineering and Computer Science Department for Computer Resources (1988–1990), Chairman of the Department of Electrical Engineering and Computer Sciences (1990–1993), and Dean of the College of Engineering (1996–2000). He currently serves as Executive Vice Chancellor and Provost and holds the Andrew S. Grove Chair in Electrical Engineering.

Prof. Gray was the corecipient of Best Paper Awards at the International Solid State Circuits Conference and the European Solid-State Circuits Conference and was the corecipient of the IEEE R.W.G. Baker Prize in 1980, the IEEE Morris K. Liebman Award in 1983, and the IEEE Circuits and Systems Society Achievement Award in 1987. In 1994, he was the recipient of the IEEE Solid-State Circuits Award. He served as Editor of the IEEE JOURNAL OF SOLID-STATE CIRCUITS from 1977 through 1979 and as Program Chairman of the 1982 International Solid State Circuits Conference. He served as President of the IEEE Solid-State Circuits Council from 1988 to 1990 and is a member of the National Academy of Engineering.

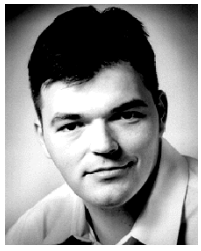

Borivoje Nikolić (S'93–M'99) received the Dipl.Ing. and M.Sc. degrees in electrical engineering from the University of Belgrade, Belgrade, Yugoslavia, in 1992 and 1994, respectively, and the Ph.D. degree from the University of California, Davis, in 1999.

He was on the faculty of the University of Belgrade from 1992 to 1996. He spent two years with Silicon Systems, Inc., Texas Instruments Storage Products Group, San Jose, CA, working on disk-drive signal processing electronics. In 1999, he joined the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, where he is now an Associate Professor. His research activities include high-speed and low-power digital integrated circuits and VLSI implementation of communications and signal processing algorithms. He is the coauthor of the book *Digital Integrated Circuits: A Design Perspective* (Prentice-Hall, 2003, 2nd ed).

Dr. Nikolić was the recipient of the National Science Foundation CAREER Award in 2003, the College of Engineering Best Doctoral Dissertation Prize, and the Anil K. Jain Prize for the Best Doctoral Dissertation in Electrical and Computer Engineering at the University of California at Davis in 1999, as well as the City of Belgrade Award for the Best Diploma Thesis in 1992.