A 12b 250MS/s Pipelined ADC with Virtual Ground Reference Buffers

Hyun H. Boo, Duane S. Boning, Hae-Seung Lee

Building 39-328, 60 Vassar St., Massachusetts Institute of Technology, Cambridge, MA 02139

Tel(Hyun Boo): (617) 901-3763 Email: hyunboo@mit.edu, boning@mtl.mit.edu, hslee@mit.edu

Abstract

The virtual ground reference buffer (VGRB) technique is introduced as a means to improve the performance of switched-capacitor circuits. The technique enhances the performance by improving the feedback factor of the op-amp without affecting the signal gain. The bootstrapping action of the level-shifting buffers relaxes key op-amp performance requirements including unity-gain bandwidth, noise, open-loop gain and offset compared with conventional circuits. This reduces the design complexity and the power consumption of op-amp based circuits. Based on this technique, a 12b pipelined ADC is implemented in 65nm CMOS that achieves 67.0dB SNDR at 250MS/s and consumes 49.7mW of power from a 1.2V power supply.

I. Introduction

The op-amp is a key building block of switched-capacitor circuits such as pipelined ADCs. Unfortunately, it has become increasingly difficult to implement a high performance op-amp [1]. The scaling of CMOS results in low intrinsic gain and therefore low open-loop op-amp gain. Techniques such as gain enhancement and multi-stage configurations have been used at the cost of added power, noise and speed. To maintain the SNR at ever decreasing supply voltages, noise is reduced by increasing the sampling capacitance, which increases power consumption. Also, implementing high bandwidth op-amp circuits for high-speed operation increases power consumption and noise bandwidth. This has led to several directions in the development of pipelined ADCs.

A common approach is to employ low-gain and/or non-settling op-amps, and digitally calibrate the resulting nonlinearity [2-5]. In order to track the power supply voltage and temperature, the

calibration in general must run continuously in the background, which can consume significant power. Analog techniques have also been employed to relax the open-loop gain requirements of the op-amp [6-8]. In one technique, the finite gain and in-band noise of the main op-amp is canceled by sampling a correction voltage onto the next-stage sampling capacitance through an auxiliary path [5]. The drawback is the power consumption in the auxiliary path and the adaptive calibration to determine the optimum auxiliary gain. Correlated level shifting (CLS) is another technique that relaxes op-amp gain requirement and increases its output swing [7-8]. The drawback is the reduced speed due to the sampling and two settling periods required for its operation.

   In other alternative approaches, the op-amp is replaced completely. Zero-crossing-based circuits (ZCBC) [9-13] are based on the insight that detecting the virtual ground node is more power-efficient than enforcing it with an op-amp in closed-loop. While energy efficient, managing signal-dependent DC voltage drops across many switches requires considerable complexity [12]. In the pulsed bucket brigade (PBB) ADC [14-15], the sampled input charge is re-used to charge the capacitors in the following stage and the voltage gain is achieved by capacitor scaling. Due to the nonlinearity of the circuit, continuous background calibration is necessary. Another alternative amplifier topology is a ring amplifier (RA) [16-18]. The speed of the RA-based switched-capacitor circuit is determined by the oscillating frequency, and the maximum sampling rate demonstrated thus far is 100MS/s at a modest 9-bit ENOB [18].

   In this work, we exploit op-amp-based switched-capacitor circuits whose principles are very well understood and the design techniques are mature. The proposed technique significantly relaxes key op-amp specifications including unity-gain bandwidth, noise, open-loop gain, and offset voltage. Since the op-amp is allowed to settle fully, calibration to remove charge-transfer error in PBB and in low gain or non-settling op-amp-based circuits is unnecessary. In addition, the transient currents and corresponding voltage drops across switches and reference buffers in the

ZCBCs and in non-settling op-amp-based circuits are avoided. The circuit is also shown to achieve higher maximum operating speed than alternative methods.

An overview of the conventional pipeline circuit is provided in Section II. Section III presents the VGRB technique. Section IV describes the implementation details of the circuit blocks. Measured results are presented in Section V and Section VI concludes the paper.

II. Conventional MDAC Switched-Capacitor Circuit

The conventional switched capacitor circuit for a 1-bit flip-around multiplying D/A converter (MDAC) is shown in Fig. 1. In the sampling phase, the input voltage is sampled on C_1 and C_2. In the charge-transfer phase, C_2 flips around the op-amp, and C_1 is driven by either the positive reference voltage, V_{REFP}, or the negative reference voltage, V_{REFM}, depending on the sub-ADC bit-decision. In a multi-bit/stage implementation, C_1 represents an array of capacitors. The reference voltages are referenced to the system ground, and on-chip buffers are typically used for high speed ADCs.

The signal gain G_s is determined by the ratio between C_1 and C_2. For example, in a typical 4-bit/stage ADC, C_1 is 7 times C_2, and G_s is 8. High signal gain corresponds to resolving a large number of bits in the corresponding stage. Assuming the input voltage is sampled on all MDAC capacitors, the feedback factor β of the circuit is the inverse of the signal gain in the conventional circuit, where parasitic capacitance is initially ignored for simplicity. Therefore, for G_s of 8, β is 1/8. The closed-loop bandwidth is given by βf_u where f_u is the unity-gain bandwidth. Since the signal gain is typically much larger than unity, the unity-gain bandwidth of the op-amp must be much higher than the closed-loop bandwidth required for settling. For given sampling capacitance of the subsequent stage, correspondingly high transconductance is required in the op-amp, thus resulting in high power consumption. In addition, the signal gain is equal to the op-amp noise gain, thus the integrated input-referred noise of the switched-capacitor circuit is equivalent to the input-

referred noise of the op-amp integrated over the closed-loop bandwidth. High op-amp open-loop gain is also required for small charge-transfer error.

III. Virtual Ground Reference Buffer Technique

The proposed VGRB technique in [19] in its charge-transfer phase is shown in Fig. 2. The sampling phase is identical to that of the conventional scheme in Fig. 1(a). In the charge-transfer phase, similarly to the conventional circuit, C_2 flips around the op-amp and C_1 is driven by either the positive or negative reference voltage. Here, however, the reference voltages are referenced to the virtual ground node instead of the system ground. Reference voltages are generated by level-shifting the virtual ground potential.

A variety of circuits, such as a simple source follower, can be used as the level-shifting buffer. For example, a PMOS source follower generates V_{REFP} by level-shifting the virtual ground up, and an NMOS source follower generates V_{REFM} by level-shifting the virtual ground potential down. The output voltage of the circuit is identical to that of the conventional circuit once the signals settle.

The benefit of this technique is revealed in the AC equivalent incremental circuit model in Fig. 3. Assuming an ideal buffer for simplicity, any change in the virtual ground node voltage is reflected at the output of the buffer, effectively bootstrapping C_1 away. Therefore, C_1 is removed from the feedback network of the op-amp, resulting in a unity feedback factor independent of the signal gain. This is in contrast to the conventional circuit, where the feedback factor β is the inverse of the signal gain G_s. In the proposed technique, that constraint no longer applies and both high signal gain and unity feedback factor can be achieved.

The performance of the VGRB technique is compared with that of the conventional circuit, both without parasitic capacitance, in Table 1. The effects of the parasitic capacitance is examined in the next subsection. In the proposed circuit, the signal gain G_s is identical to that of the conventional circuit. However, the feedback factor β is unity independent of the signal gain. Thus,

the closed-loop bandwidth is higher by a factor of the signal gain compared with that of the conventional circuit. This means either lower op-amp power or faster circuit operation can be achieved. Also, the op-amp noise gain is reduced to unity and thus the op-amp noise is reduced by a factor of the signal gain G_s resulting in a factor of G_S^2 reduction of noise spectral density referred to the ADC input. This is a major advantage since reducing input referred noise requires significant power consumption. In addition, the op-amp open-loop gain and offset requirements are reduced by a factor of the signal gain. Therefore, all key performance parameters are improved by a factor of the signal gain.

For a fair comparison, however, bandwidth, noise, and power consumption of the buffers must be considered. In the presented ADC, the buffers replace existing reference buffers, thus incurring insignificant penalty; while off-chip bypass capacitors can be used in conventional circuits to ease the reference buffer requirements [20-22], on-chip high-speed buffers are strongly desired in high-speed, high-SNR ADCs to avoid reference voltage ringing [23-25]. More detailed analysis of the buffer bandwidth and noise is presented in the next subsection.

It is important to note that the benefits are obtained without affecting the signal gain. The performance is improved by breaking away from the conventional inverse relationship between the signal gain and the feedback factor. Although the VGRB technique is applied in the implementation of a pipeline ADC in this paper, the technique can be applied to a broader range of switched-capacitor circuits, for example, track-and-hold circuits and delta-sigma ADCs.

A. Effect of Parasitic Capacitance

To be more realistic, the effect of parasitic capacitance must be considered as shown in Fig. 4. The actual improvement in the feedback factor depends on the value of the parasitic capacitance C_p at the virtual ground node. There are four sources of C_p: the finite gain of the buffer, the buffer input capacitance $C_{i,b}$, the op-amp input capacitance $C_{i,oa}$, and routing capacitance C_w. In practice, the incremental gain of the source follower buffer is typically less than unity and C_1 is not

completely bootstrapped away; a portion of C_1 loads the virtual ground node. If the buffer gain is 0.9 instead of unity, the buffer bootstraps away only 90% of the C_1 capacitance, and the remaining 10% of C_1 effectively shows up at the virtual ground node. The source follower input capacitance is dominated by the gate-to-drain capacitance C_{gd} because C_{gs} is bootstrapped away and is not a major source of the input capacitance. The op-amp input capacitance is another source of parasitic capacitance, and it affects both conventional and proposed designs. In the proposed scheme, for the same op-amp noise and bandwidth requirements, the op-amp requires much smaller input transistors, and thus has much smaller $C_{i,oa}$. The routing of the virtual ground node contributes parasitic capacitance in the conventional circuit. In the proposed scheme, it can be mostly bootstrapped away by the buffers as shown later in the implementation. Considering all sources of parasitic capacitance, the prototype chip achieves an actual feedback factor of 0.3 compared with 0.1 in a conventional circuit, while a feedback factor of 0.5 would be realistically achievable by further optimization.

B. Noise Analysis

The noise contribution from the op-amp and the level-shifting buffers in the first stage is compared with that in a conventional circuit that has the same bandwidth. The additional loading capacitance from the sub-ADC sampling network is not considered and the comparison is made with simplifying assumptions.

Conventional Circuit

The sampling capacitance in the second stage, C_L, is assumed to be scaled by a factor of four compared to the sampling capacitance $C_1 + C_2$ in the first stage. Although 3.9-bits are resolved in the first stage and more aggressive capacitor scaling is possible for better power optimization, often the capacitance sizing is dictated by the capacitor matching requirement to ensure the MDAC linearity [26]. For a signal gain of $G_s = 8$, the feedback factor $\beta = \frac{1}{10}$ considering the effect of

parasitic capacitance, which is approximately twice the feedback capacitance C_2 in a typical implementation, at the virtual ground node. Thus $C_1 = 7C_2$ and $C_p = 2C_2$. The total load capacitance C_T is

$$C_T = \frac{C_2(C_1+C_p)}{C_1+C_2+C_p} + C_L = 2.9C_2. \qquad (1)$$

The op-amp noise referred to the ADC input during the charge-transfer phase is

$$\overline{V^2_{n,oa,in}} = \left(\frac{4kT\gamma(2g_{mn,oa}+2g_{mp,oa})n_f}{g^2_{mn,oa}}\right)\frac{BW}{4}\frac{1}{\beta^2}\frac{1}{G^2_s} \qquad (2)$$

where $g_{mp,oa}$ and $g_{mn,oa}$ are the transconductance of the PMOS and NMOS device in the input stage of the op-amp, respectively, and γ, typically $2/3$ in strong inversion, is a coefficient dependent on the region of operation [1]. An NMOS transistor input pair with PMOS current source load is assumed in the op-amp topology. The noise from the added circuitry such as gain boosting and multi-stage amplifier is considered by the noise multiplication factor, n_f. The closed-loop bandwidth $BW = \frac{g_{mn,oa}}{C_T}\beta = \frac{g_{mn,oa}}{29C_2}$. For simple PMOS and NMOS source follower reference buffers, noise power referred to the ADC input is shown to be

$$\overline{V^2_{n,b,in}} = \left(\frac{4kT\gamma}{g_{mp,b}} + \frac{4kT\gamma}{g_{mn,b}}\right)\frac{N^2}{12}\frac{BW}{4}\frac{1}{G^2_s} \qquad (3)$$

where MDAC code-dependent partial cancellation of reference buffer noise in fully-differential topology is considered. Here $g_{mp,b}$ and $g_{mn,b}$ are the transconductance of the PMOS and NMOS device in the buffers and N is the maximum sub-ADC code, which is 14 in the prototype, and uniform-code distribution is assumed.

Assuming $g_m = g_{mp,oa} = g_{mn,oa}$, and $g_m = g_{mp,b} = g_{mn,b}$ for simplicity, the total input-referred noise power in the charge-transfer phase is

$$\overline{V^2_{n,in}} = \overline{V^2_{n,oa,in}} + \overline{V^2_{n,b,in}} = \frac{kT\gamma}{C_2}\left(0.216n_f + 0.035\right). \qquad (4)$$

The noise contribution of the reference buffers is insignificant because their noise is partially cancelled by the differential architecture and further filtered by the op-amp closed-loop bandwidth, and because their noise gain is lower than the signal gain.

VGRB Circuit

We assume that the VGRB technique improves the feedback factor to 0.3 as in the prototype, although it can be further improved in future designs. The transconductance $g_{mp,b}$ and $g_{mn,b}$ are assumed to be identical to those in the reference buffer in the conventional circuit for a straightforward comparison. The load capacitance C_L is assumed to be identical as well. Given the feedback factor $\beta = \frac{C_2}{C_2 + C_p} = 0.3$, the effective parasitic capacitance at the virtual ground node is $C_p = \frac{7}{3} C_2$ giving the total load capacitance of $C_T = 2.7 C_2$. Due to the larger feedback factor and reduced load capacitance, the input transconductance of the op-amp can be reduced to $g'_m = 0.31 g_m$ for the same closed-loop bandwidth as that of the conventional circuit, and the op-amp power consumption can be reduced by the same factor. The op-amp noise power referred to the ADC input is reduced to $\overline{V_{n,oa,in}^2} = 0.077 \frac{kT\gamma n_f}{C_2}$. Note that this is much lower than that of the conventional circuit due to the increased feedback factor, even at much lower power consumption. Unlike the conventional circuit, noise from the level-shifting buffers does not partially cancel [27]. Again, assuming a simple PMOS and NMOS source follower, for the implemented MDAC scheme in IV.B, the buffer noise referred to the ADC input is

$$\overline{V_{n,b,in}^2} = \left(\frac{8kT\gamma}{g_{mp,b}} + \frac{8kT\gamma}{g_{mn,b}} \right) \frac{N^2}{12} \frac{BW}{4} \frac{1}{G_s^2}. \qquad (5)$$

The maximum sub-ADC code, N, is identical to the maximum sub-ADC code in Eqn. (3) for comparison and the signal gain is maintained. The buffer noise is still limited by the op-amp bandwidth since the op-amp transconductance is reduced to g'_m. Finally, again assuming $g_m = g_{mn,b} = g_{mn,b}$ for simplicity, the total input-referred noise power is

$$\overline{V_{n,in}^2} = \frac{kT\gamma}{C_2}\left(0.077n_f + 0.070\right). \qquad (6)$$

Compared to the conventional circuit the op-amp power is reduced by a factor of 3.2 and the noise contribution from the op-amp is reduced by a factor of 2.8. The reference buffer noise contribution increases by a factor of two due to the lack of the partial noise cancellation, but the overall noise is still much lower in the proposed circuit. Even in the conservative case of $n_f = 1$ where only the input stage transistors contribute noise, the total input-referred noise power of the op-amp and the reference buffers is reduced by a factor of 1.6. It is interesting to note that with $\gamma = 2/3$ in strong inversion, the total noise power during the charge-transfer phase is lower than the sampling noise. In reality, the op-amp in the conventional circuit will likely have higher n_f since more devices are required to achieve the necessary higher op-amp open-loop gain, and the input-referred noise power as well as op-amp power consumption in the VGRB technique become even more favorable. Moreover, the feedback factor in the proposed circuit can be improved further, which will increase the improvement factors in both the power consumption and noise.

IV. Circuit Implementation

A. Overall Structure

The top-level block diagram of the pipelined ADC implemented based on the VGRB technique is shown in Fig. 5. The ADC consists of five stages, the first stage resolving 3.9-bits with 14 comparators and a signal gain of 8. The second stage to the fourth stage are identical, each stage resolving 2.6-bits with 6 comparators and a signal gain of 4. These stages resolve 2-bits excluding the over-range correction bits. The final stage is a 4.5-bit flash. The digital bits in each stage are time-aligned, then added to convert to a binary code from a thermometer code, and the low-voltage differential signal (LVDS) buffers drive the digital codes off-chip where they are truncated to 12

bits resulting in 11.9 bit effective quantization. The chip also includes circuits for clock buffering and biasing.

B. Multiplying Digital-to-Analog Converter

A total of four level-shifting buffers are employed in each stage of the fully-differential prototype as shown in Fig. 6 since each virtual ground node requires two level-shifting buffers to generate the positive and negative reference voltages. Therefore, the implementation in this prototype chip consumes twice as much power in the buffers compared with conventional circuits that have two reference buffers. Although not implemented in the prototype, since each stage has a separate set of reference buffers that are engaged only in the charge-transfer phase, the reference buffers can be duty-cycled to save approximately 50% of power. Therefore, the total buffer power would be comparable to those in conventional circuits. Moreover, by breaking up the buffers into multiple parallel segments and selectively enabling them based on the number of capacitors connected to them, buffer power can be reduced by an additional 50%.

C. Stage Residue and Over-Range Correction

The first stage has a peak-to-peak differential input signal range of $1.5V$ and resolves 3.9-bits. The comparator bit-decision threshold levels in the first stage sub-ADC are $100mV$ apart, resulting in a peak-to-peak differential output swing of $800mV$ in the MDAC for a signal gain of 8. The over-range correction extends the peak-to-peak differential output swing to $1V$. In the second to fourth stages, comparator bit-decision threshold levels in the sub-ADC are $200mV$ apart and the over-range can be applied to a differential input swing range as large as $1.4V$. However, the residue of the preceding stage is limited to a differential swing of $1V$ since the op-amps suffer from nonlinearity beyond that range.

D. Gain-Boosted Op-Amp

Due to the reduced open-loop gain requirement, the proposed technique allows a single-stage gain-boosted telescopic op-amp to be used in the pipelined stages for 12-bit accuracy. In Fig. 7 the

NMOS and PMOS cascode transistors are gain-boosted by a folded cascode boosting amplifier with a PMOS and NMOS differential input pair, respectively. The gain-boosting amplifier uses a single transistor for common-mode feedback for simplicity [28]. A conventional switched-capacitor common-mode feedback circuit is applied to the op-amp.

Since the parasitic capacitance of the input pairs degrades the feedback factor of the closed-loop op-amp switched-capacitor circuit as described in the preceding section, the input transistors are kept at minimum length to minimize parasitic contributions.

E. Level-Shifting Buffers

The flipped voltage follower (FVF) [29-30] in Fig. 8 is used as the level-shifting buffer. Compared with the conventional source follower, the PMOS and NMOS FVFs have better current sourcing and sinking ability, respectively. In the charge-transfer phase, the PMOS buffers only need to source current to C_1 and the NMOS buffers only need to sink current to C_1 for an input voltage within the full-scale. Therefore, the current sourcing from the PMOS FVFs and the current sinking from the NMOS FVFs give much improved slewing.

In addition to using low-V_T input transistors, a switched-capacitor circuit sets the body-to-source voltage, V_{BS}, of the input transistor to further reduce its V_T. The bulk biasing voltages V_{BIASP} and V_{BIASM} reduce the $|V_{GS}|$ in the NMOS and PMOS FVFs, respectively, to achieve the targeted reference voltage $V_{BIASP} - V_{BIASM} \approx 800mV$. In the sampling phase, Φ_1, the bulk biasing voltage is sampled on C_{s1}. In the charge-transfer phase, Φ_2, the charge on C_{s1} is shared with C_{s2} applying the voltage across the source and bulk terminals of the input transistor. Throughout the ADC, all PMOS and NMOS level-shifting buffers share the same V_{BIASP} and V_{BIASM}, respectively.

Since individual stages have separate reference buffers, mismatches in the effective reference voltage between stages results in linearity errors. For this reason, reference calibration may be necessary as explained in Section V. In the prototype, both the positive and the negative reference

buffers have gain of 0.91 according to simulation. However, the gain mismatch between the positive and negative reference buffers can degrade the common-mode rejection ratio (CMRR). Gain matching within 0.02 is sufficient to provide an average 40dB CMRR over the input range. It should also be noted that the worst-case feedback factor is realized when all MDAC capacitors are switched to the lower gain reference buffer.

F. Sub-ADC

A flash ADC is used as the sub-ADC in the pipelined stage. The 3.9-bit flash ADC in the first stage with 14 decision levels is shown in Fig. 9. The sampling process is identical to that of the MDAC. To minimize the dynamic offsets generated from the timing mismatch in the sampling network between the MDAC and the sub-ADC, the sub-ADC in the first stage uses unit capacitance and switch sizes identical to those used in the MDAC. The dynamic offset in the second stage and beyond is relatively less serious since both the MDAC and sub-ADC sample a DC signal from the previous stage. Therefore, the sub-ADC unit capacitor is scaled down to 25fF to reduce the capacitive loading to the op-amp in the previous stage. Comparator offsets are calibrated using calibration voltages V_{CALP} and V_{CALM} to ensure that the amplified residue in the MDAC stays within the over- range correction. In order to measure comparator offsets, for each sub-ADC code in the first stage, the sub-ADC codes from the second stage to the last stage are combined to construct the back-end code. The back-end code corresponds to the digital representation of the first stage residue. If the comparator offsets in the first stage sub-ADC are not calibrated, then the back-end codes do not fit in the nominal range. V_{CALP} and V_{CALM} are adjusted to bring the back-end codes within the nominal range.

The StrongArm comparator with an extra input pair for calibration in the first stage sub-ADC is shown in Fig. 10(a). The calibration input pair is scaled down in size by a factor of 5 compared with the main input pair. A dedicated resistor ladder in Fig. 10(b) generates an array of calibration

voltages that are multiplexed into the calibration input pair. These differential calibration voltages control the comparator offsets by a step size of $\pm 8mV$ up to $\pm 40mV$.

The sub-ADCs in the second to the fourth stage are identical and have the same structure as the first stage sub-ADC. Each sub-ADC resolves 2-bits and the input transistors in the StrongArm comparators are scaled up for smaller offsets to avoid offset calibration.

G. Bootstrapping the Virtual Ground Node Parasitics

In conventional circuits, the routing capacitance between the virtual ground and the substrate degrades the feedback factor. The input-referred noise and charge-transfer error increases, and the bandwidth decreases as a result. In the VGRB technique, the routing capacitance can be bootstrapped away by the level-shifting buffers as shown in Fig. 11. The level-shifting buffer must drive a slightly larger capacitive load since it now sees the additional parasitic capacitance $C_{p,rout}$ and $C_{p,substrate}$. However, $C_{p,rout}$ and $C_{p,substrate}$ are typically small compared with the C_1 the level-shifting buffers are designed to drive, thus the additional capacitance is not significant. For the prototype ADC, the NMOS level-shifting buffer is used to bootstrap the routing capacitance since it has a higher transconductance.

H. Feedback Factor in Post-Simulation

An ideal unity gain is desired in the level-shifting buffers to completely bootstrap away C_1, which is 840fF in the prototype. In simulation, both level-shifting buffers were found to have gain of 0.91 instead. Therefore, 840fF×0.09 = 75.6fF of C_1 effectively appears at the virtual ground node. The PMOS and NMOS level-shifting buffers combined contribute 87.7fF, mostly from C_{gd} of the input transistors. The op-amp input capacitance adds another 78fF. Extra parasitic capacitance comes from the routing as the routing capacitance was not bootstrapped aggressively enough. The final achieved feedback is approximately β=0.3, which can be realistically improved to β=0.5 by increasing the gain of the level-shifting buffer via cascoding the current sources and bootstrapping the drain of the source-follower buffers, which also bootstraps away the C_{gd}.

V. Measurement Results

The prototype chip was fabricated in a 65nm LP technology. The die, whose core occupies 0.59mm^2, is shown in Fig. 12. The full-scale input signal range is approximately 1.5V_{pp} differentially. One-time foreground capacitor mismatch calibration for the first stage is applied off-chip and the calibration coefficients are frozen throughout the measurements. The subsequent stage capacitors are sized to provide sufficient matching without calibration. During the normal conversion, the capacitor mismatch calibration requires coefficient additions only, and is estimated to consume $100\mu\text{W}$ if implemented on-chip.

Fig. 13 shows the measured linearity of the ADC before the reference voltage calibration. The INL has 14 vertical voltage drops if reference voltage calibration is not applied. This indicates mismatch in the reference voltage between the first stage and the second stage. In the prototype, the reference voltage is systematically lower in the first stage than in the second stage due to the DC voltage drop along the power supply and ground routing. To resolve the difference, the differential reference voltage in the first stage is increased by increasing the current for the first stage NMOS level-shifting buffers. Also, the differential reference voltage in the second stage is decreased by decreasing the current for the second stage NMOS level-shifting buffers. Within each 14 segments in the INL plot, additional vertical drops exist that arise from the reference voltage mismatch between the second stage and third stage, as well as the third stage and the fourth stage. The mismatch is calibrated by adjusting the current level of the NMOS level-shifting buffers in the corresponding stages. Fig. 14 shows the DNL and INL after the reference calibration. The DNL and INL are within -0.86/+0.52 LSB and -0.90/+1.08 LSB, respectively, as shown in Fig. 14(b). Although automatic foreground or background calibration of references is relatively straightforward, it is not implemented in the prototype for simplicity. Capacitor mismatch calibration can automatically remove reference mismatches. However, reference mismatches can vary with temperature and power supply voltages necessitating background calibration. Thus,

simple analog background calibration may be more power efficient. For example, the output of each reference buffer is periodically compared with a desired reference voltage, and the reference current is incremented or decremented in a feedback control loop.

The measured noise floor of the ADC is -68dB, the sampling and quantization (11.9b) noise contributing -75dB and -73dB respectively, and the op-amp and buffers together contributing -71dB. Fig. 15 shows the output spectra with input frequencies of 12.1MHz and 121.8 MHz at the sampling rate of 250MS/s. At 12.1MHz, SNDR of 67.0dB (10.84-b ENOB) and SFDR of 84.6dB are achieved. At 121.8MHz, the SNDR degrades by 1.3dB to 65.7dB, which is believed to be due to the sampling clock jitter. Fig. 16 shows the SNDR and SFDR performance of three randomly selected dies at a sampling rate of 250MS/s as the input frequency is swept. The SNDR is consistent across the three chips. Sampling frequency is swept in Fig. 17 while the input signal is set at a low frequency. The ENOB is 10.3-bits even at 280MS/s sampling rate.

The chip operates from a 1.2V power supply, and consumes 49.7mW at 250MS/s. The current consumption for each circuit block is: 23.3mA in the op-amps; 6.9mA in the buffers; 6.0mA in the clock, digital and flash ADC; 3.3mA in the flash reference ladders; and 2.0mA in the biasing cell. Even though the level-shifting buffers in the current implementation consume twice as much power as the buffers in the conventional circuits, they are far from being the dominant power consumption block, and can be further reduced by duty-cycling and multiple unit buffers as explained in the previous section. The op-amp takes more than half of the total power consumption. In this proof of concept chip, standard power saving measures such as op-amp scaling and duty cycling were not implemented for simplicity. Therefore, there is substantial room for further optimization and large power savings in future designs.

The performance of this work compared to other single-channel ADCs with a resolution of 12 bits and sampling rate of 200MS/s or greater is shown in Table 3. In [13] and [15], the op-amps are replaced, and in [3], an op-amp-based digital calibration scheme is presented.

VI. Conclusion

This paper presents the VGRB technique that relaxes the op-amp performance specifications by improving its feedback factor in switched-capacitor circuits. In conventional designs, the signal gain of a switched-capacitor circuit is largely determined by the inverse of the feedback factor and this determines the required op-amp specifications. In this work, the feedback factor is improved without altering the signal gain. Unlike alternative approaches, it retains the conventional design style as much as possible while improving the performance by presenting a simple yet power-efficient way of operating an op-amp-based switched-capacitor circuit. The proposed technique has advantages in speed, noise, and settling accuracy as well as power consumption compared with the conventional op-amp based circuit. The prototype pipelined ADC achieves 67dB SNDR at 250MS/s with 49.7mW and is comparable to the state of the art.

REFERENCES

[1] B. Razavi, "Operational Amplifiers," in *Design of Analog CMOS Integrated Circuits*. New York, NY, USA: McGraw-Hill, 2001.

[2] J.-R. Kim and B. Murmann, "A 12-b, 30-MS/s, 2.95-mW Pipelined ADC Using Single-Stage Class-AB Amplifiers and Deterministic Background Calibration," IEEE Journal of Solid-State Circuits, vol. 47, no. 9, pp. 2141–2151, Sep. 2012.

[3] B. Sahoo and B. Razavi, "A 12-Bit 200-MHz CMOS ADC," IEEE Journal of Solid-State Circuits, vol. 44, no. 9, pp. 2366–2380, Sep. 2009.

[4] A. Panigada and I. Galton, "A 130 mW 100 MS/s Pipelined ADC With 69 dB SNDR Enabled by Digital Harmonic Distortion Correction," IEEE Journal of Solid-State Circuits, vol. 44, no. 12, pp. 3314–3328, Dec. 2009.

[5] A. Verma and B. Razavi, "A 10-Bit 500-MS/s 55-mW CMOS ADC," IEEE Journal of Solid-State Circuits, vol. 44, no. 11, pp. 3039–3050, Nov. 2009.

[6] Y. Miyahara, M. Sano, K. Koyama, T. Suzuki, K. Hamashita, and B.-S. Song, "Adaptive cancellation of gain and nonlinearity errors in pipelined ADCs," in IEEE Solid-State Circuits Conference Digest of Technical Papers, Feb. 2013, pp. 282–283.

[7] B. Gregoire and U.-K. Moon, "An Over-60dB True Rail-to-Rail Performance Using Correlated Level Shifting and an Opamp with 30dB Loop Gain," in IEEE Solid-State Circuits Conference Digest of Technical Papers, Feb. 2008, pp. 540–634.

[8] B. Hershberg, S. Weaver, and U.-K. Moon, "Design of a Split-CLS Pipelined ADC With Full Signal Swing Using an Accurate But Fractional Signal Swing Opamp," IEEE Journal of Solid-State Circuits, vol. 45, no. 12, pp. 2623–2633, Dec. 2010.

[9] L. Brooks and H.-S. Lee, "A 12b 50MS/s fully differential zero-crossing-based ADC without CMFB," in IEEE Solid-State Circuits Conference Digest of Technical Papers, Feb. 2009, pp. 166–167.

[10] J. Chu, L. Brooks, and H.-S. Lee, "A zero-crossing based 12b 100MS/s pipelined ADC with decision boundary gap estimation calibration," in IEEE Symposium on VLSI Circuits, Jun. 2010, pp. 237–238.

[11] S. Lee, A. Chandrakasan, and H.-S. Lee, "A 12 b 5-to-50MS/s 0.5-to-1 V Voltage Scalable Zero-Crossing Based Pipelined ADC," IEEE Journal of Solid-State Circuits, vol. 47, no. 7, pp. 1603–1614, Jul. 2012.

[12] D.-Y. Chang, C. Munoz, D. Daly, S.-K. Shin, K. Guay, T. Thurston, H.-S. Lee, K. Gulati, and M. Straayer, "A 21mW 15b 48MS/s zero-crossing pipeline ADC in 0.13μm CMOS with 74dB SNDR," in IEEE Solid-State Circuits Conference Digest of Technical Papers, Feb. 2014, pp. 204–205.

[13] S.-K. Shin, J. Rudell, D. Daly, C. Munoz, D.-Y. Chang, K. Gulati, H.-S. Lee, and M. Straayer, "A 12 bit 200MS/s Zero-Crossing-Based Pipelined ADC With Early Sub-ADC Decision and Output Residue Background Calibration," IEEE Journal of Solid-State Circuits, vol. 49, no. 6, pp. 1366–1382, Jun. 2014.

[14] M. Anthony, E. Kohler, J. Kurtze, L. Kushner, and G. Sollner, "A process-scalable low-power charge-domain 13-bit pipeline ADC," in IEEE Symposium on VLSI Circuits, Jun. 2008, pp. 222–223.

[15] N. Dolev, M. Kramer, and B. Murmann, "A 12-bit, 200-MS/s, 11.5-mW pipeline ADC using a pulsed bucket brigade front-end," in IEEE Symposium on VLSI Circuits, Jun. 2013, pp. C98–C99.

[16] B. Hershberg, S.Weaver, K. Sobue, S. Takeuchi, K. Hamashita, and U.-K. Moon, "Ring Amplifiers for Switched Capacitor Circuits," IEEE Journal of Solid-State Circuits, vol. 47, no. 12, pp. 2928–2942, Dec. 2012.

[17] B. Hershberg and U.-K. Moon, "A 75.9dB-SNDR 2.96mW 29fJ/conv-step ringamp-only pipelined ADC," in IEEE Symposium on VLSI Circuits, Jun. 2013, pp. C94–C95.

[18] Y. Lim and M. P. Flynn, "A 100MS/s 10.5b 2.46mW comparator-less pipeline ADC using self-biased ring amplifiers," in IEEE Solid-State Circuits Conference Digest of Technical Papers, Feb. 2014, pp. 202–203.

[19] H. H. Boo, D. S. Boning, and H.-S. Lee, "12b 250MS/s pipelined ADC with virtual ground reference buffers," in IEEE Solid-State Circuits Conference Digest of Technical Papers, Feb. 2015, pp. 282–283.

[20] L. Singer, S. Ho, M. Timko, and D. Kelly, "A 12b 65MSample/s CMOS ADC with 82dB SFDR at 120MHz," in IEEE Solid-State Circuits Conference Digest of Technical Papers, Feb. 2000, pp. 38–39.

[21] Y.-J. Cho and S.-H. Lee, "An 11b 70 MHz 1.2 mm^2 49 mW 0.18 μm CMOS ADC with On-Chip Current/Voltage References," IEEE Transactions on Circuits and Systems I: Regular Papers, vol. 52, no. 10, pp. 1989–1995, Oct. 2005.

[22] K. Gulati, M. Peng, A. Pulincherry, C. Munoz, M. Lugin, A. Bugeja, J. Li, and A. Chandrakasan, "A Highly Integrated CMOS Analog Baseband Transceiver With 180 MSPS 13-bit Pipelined CMOS ADC and Dual 12-bit DACs," IEEE Journal of Solid-State Circuits, vol. 41, no. 8, pp. 1856–1866, Aug. 2006.

[23] L. Sumanen, M. Waltari, and K. Halonen, "A 10-bit 200-MS/s CMOS parallel pipeline A/D converter," IEEE Journal of Solid-State Circuits, vol. 36, no. 7, pp. 1048–1055, Jul. 2001.

[24] Z. Cao, S. Yan, and Y. Li, "A 32 mW 1.25 GS/s 6b 2b/Step SAR ADC in 0.13 m CMOS," IEEE Journal of Solid-State Circuits, vol. 44, no. 3, pp. 862–873, Mar. 2009.

[25] Y.-D. Jeon, Y.-K. Cho, J.-W. Nam, K.-D. Kim, W.-Y. Lee, K.-T. Hong, and J.-K. Kwon, "A 9.15mW 0.22mm^2 10b 204MS/s pipelined SAR ADC in 65nm CMOS," in IEEE Custom Integrated Circuits Conference, Sep. 2010, pp. 1–4.

[26] I. Ahmed and D. Johns, "An 11-bit 45 MS/s pipelined ADC with rapid calibration of DAC errors in a multibit pipeline stage," IEEE Journal of Solid-State Circuits, vol. 43, no. 7, pp. 1626–1637, Jul. 2008.

[27] H. H. Boo, "Virtual Ground Reference Buffer Technique in Switched-Capacitor Circuits," Ph.D. dissertation, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, 2015.

[28] J. Lloyd and H.-S. Lee, "A CMOS op amp with fully-differential gain-enhancement," IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing, vol. 41, no. 3, pp. 241–243, Mar. 1994.

[29] R. Carvajal, J. Ramirez-Angulo, A. Lopez-Martin, A. Torralba, J. Galan, A. Carlosena, and F. Chavero, "The flipped voltage follower: a useful cell for low-voltage low-power circuit design," IEEE Transactions on Circuits and Systems I: Regular Papers, vol. 52, no. 7, pp. 1276–1291, Jul. 2005.

[30] T. Man, K. N. Leung, C. Y. Leung, P. Mok, and M. Chan, "Development of Single-Transistor-Control LDO Based on Flipped Voltage Follower for SoC," IEEE Transactions on Circuits and Systems I: Regular Papers, vol. 55, no. 5, pp. 1392–1401, Jun. 2008.

(a)

(b)

Fig. 1. (a) Sample and (b) charge-transfer operation of conventional switched-capacitor circuits.
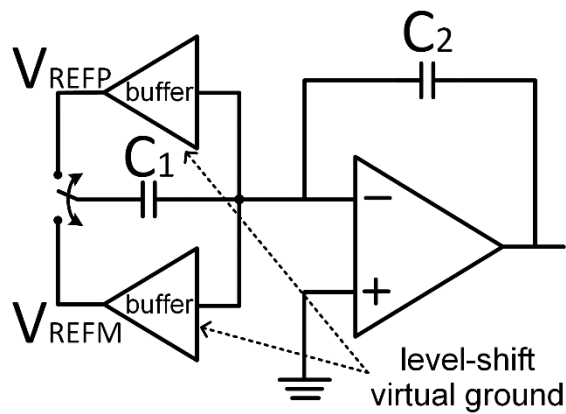
Fig. 2. Proposed virtual ground reference buffer technique.

Fig. 3. The AC equivalent incremental circuit model of the virtual ground reference buffer technique circuit in the charge-transfer phase.
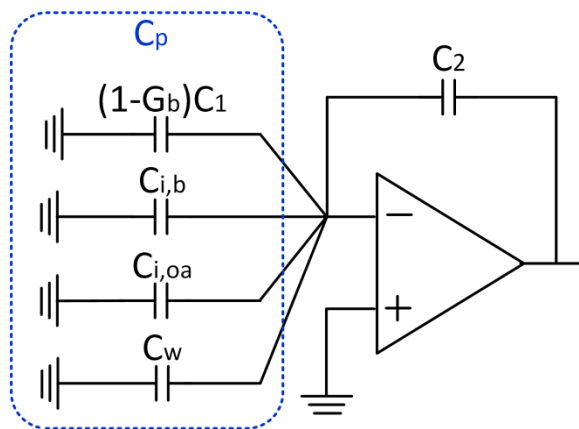
Fig. 4. Sources of parasitic capacitance that loads the virtual ground node and degrades op-amp feedback factor.
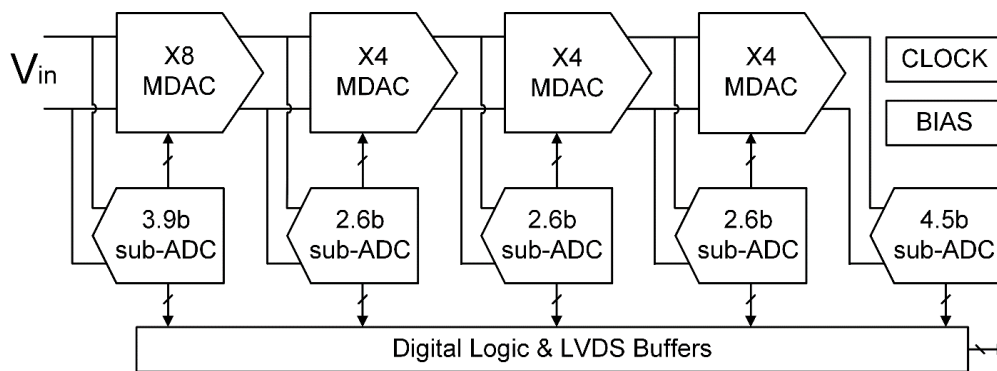
Fig. 5. Block diagram of the 12-bit pipelined ADC.
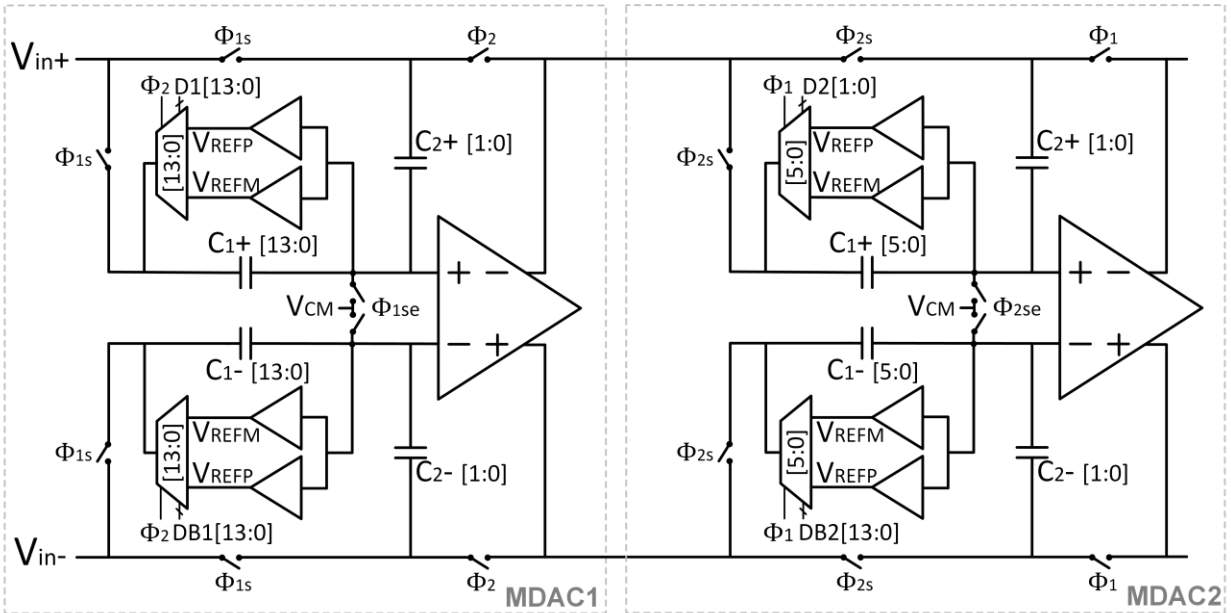
(a)

(b)

Fig. 6. Schematic of first and second MDAC stages and the corresponding timing of the operation.

Fig. 7. Gain-boosted telescopic op-amp.

(a)   (b)
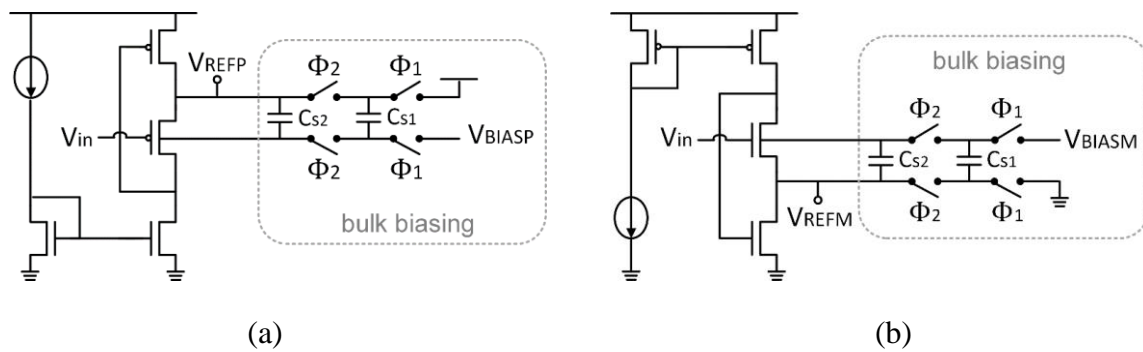
Fig. 8. (a) PMOS and (b) NMOS flipped voltage follower for generation of V_{REFP} and V_{REFM}, respectively.

(a)

(b)

Fig. 9. Sub-ADC of the first pipelined stage and (b) corresponding operation timing.
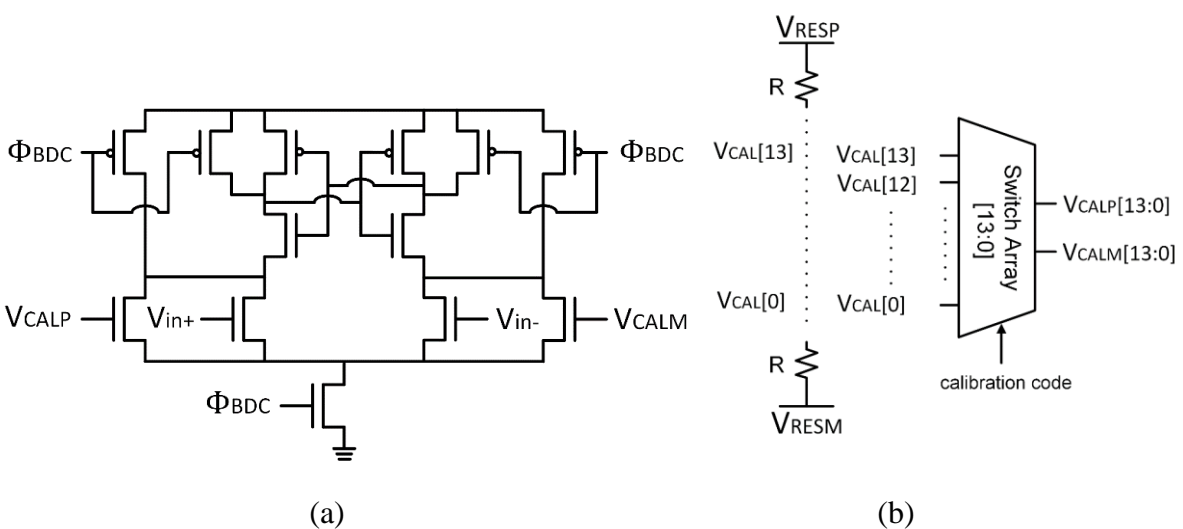
(a)

(b)

Fig. 10. (a) First stage sub-ADC StrongArm comparator with offset calibration input pair and (b) resistor ladder and multiplexor for offset calibration.

Fig. 11. The routing of the NMOS level-shifting buffer output underneath the virtual ground node.
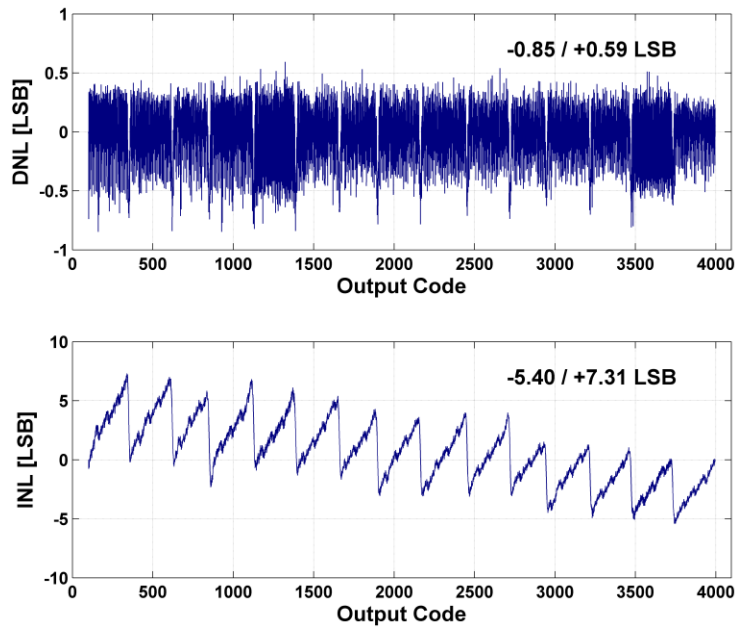

Fig. 12. Die photo.

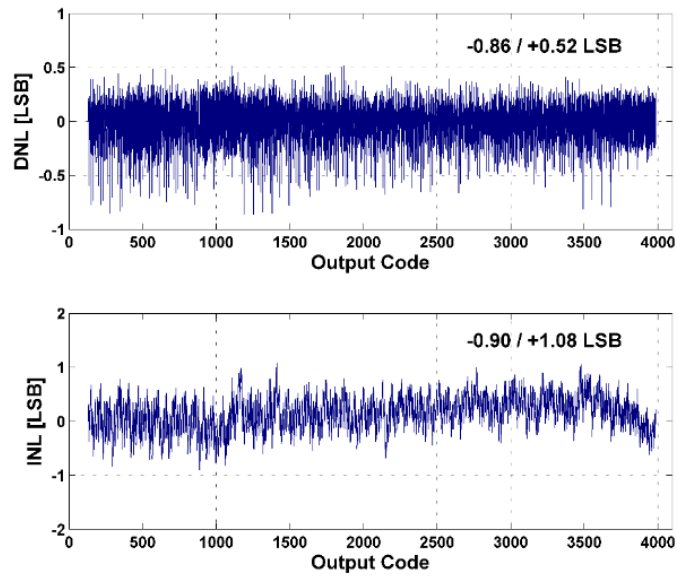Fig. 13. Measured DNL/INL of the ADC before reference voltage calibration.

Fig. 14. Measured DNL/INL of the ADC after reference voltage calibration.

(a)   (b)

Fig. 15. Measured spectrum for (a) f_{in} = 12.1MHz and (b) f_{in}=121.8MHz.
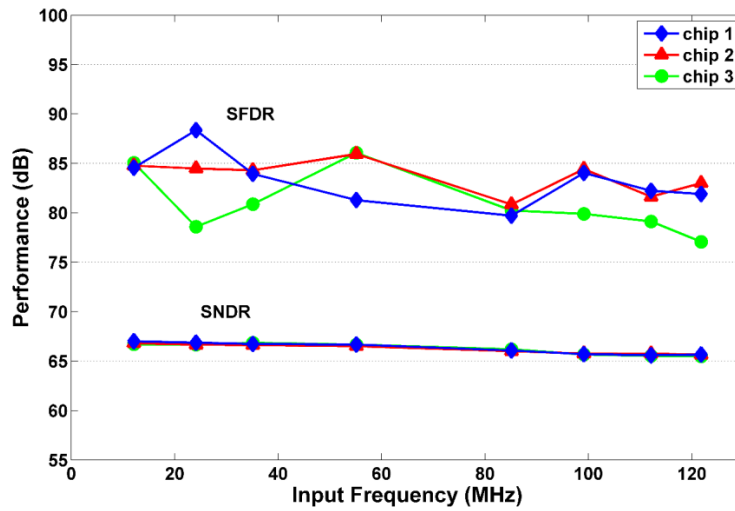

Fig. 16. Measured dynamic performance vs. input frequency for three randomly selected chips.

Fig. 17. Measured dynamic performance vs. sampling frequency.

	Conventional	Proposed
G_S signal gain	$1 + \dfrac{C_1}{C_2}$	$1 + \dfrac{C_1}{C_2}$
β feedback factor	$\dfrac{C_2}{C_1 + C_2}$	1
Closed-loop bandwidth	$\dfrac{f_u}{G_S}$	f_u
Input-referred noise density	$S_{n,oa}(f)^{**}$	$\dfrac{S_{n,oa}(f)^{***}}{G_S{}^2}$
Charge-transfer error	$\dfrac{G_S}{A^*}$	$\dfrac{1}{A}$

*   op-amp open-loop gain
**   op-amp input-referred spectral noise density, noise BW = f_u/G_S
*** op-amp input-referred spectral noise density, noise BW = f_u

Table 1. Performance comparison of a switched-capacitor circuit based on the conventional approach and the VGRB technique assuming no parasitic capacitance.

	This Work	Shin [13] JSSC 2014	Dolev [15] VLSI 2013	Sahoo [3] JSSC 2009
Technology	65nm	55nm	65nm	90nm
Type	Op-amp-based virtual ground reference buffer	Zero-crossing based	Pulsed bucket brigade digital calibration	Op-amp-based digital calibration
Sampling Rate	250MS/s	200MS/s	200MS/s	200MS/s
Resolution	12b	12b	12b	12b
Power Supply	1.2V	1.1V	1V	1.2V
Input Signal Range	1.5Vp-p	2Vp-p	-	1.2Vp-p
SNDR	67.0dB @ 12.1MHz 65.7dB @ 121.8MHz	64.6dB @ 10.1MHz 63.2dB @ 100MHz	65dB @ 1MHz 57.6dB @ 99MHz	64dB @ 3.5MHz 61.6dB @ 91MHz
SFDR	84.6dB @12.1MHz	82.9dB @ 10.1MHz	82dB @ 1MHz	70dB
Power	49.7mW	30.7mW	11.5mW*	348mW
FoM	108.5fJ/step	111.0fJ/step	39.6fJ/step*	1.34pJ/step

* Excludes power consumption in reference buffers and digital calibration circuits

Table 2. Performance summary and comparison with other single-channel ADCs with similar performance specifications.