Design of a Split-CLS Pipelined ADC With Full Signal Swing Using an Accurate But Fractional Signal Swing Opamp

Benjamin Hershberg, Skyler Weaver, and Un-Ku Moon, *Fellow, IEEE*

***Abstract*—Building on the technique of correlated level shifting (CLS), Split-CLS is introduced as a viable way to enable the design of high performance, high resolution A/D converters in deep submicron CMOS processes. One possible implementation of Split-CLS is presented, which achieves very high effective gain, and combines the fast, high efficiency charging of a zero-crossing based circuit (ZCBC) with the high-accuracy, low power settling of a double-cascode telescopic opamp. A dynamic zero-crossing detector (ZCD) conserves power in the ZCBC by only creating high bandwidth in the ZCD near the zero-crossing instant. Measured results are presented from a pipelined A/D converter fabricated in 0.18 μm CMOS. Using the Split-CLS structure, an opamp with 300 mV output swing is used to produce a pipeline stage output swing of 1.4 V. The proof-of-concept test chip achieves 68.3 dB SNDR (11.1b ENOB) and 76.3 dB SFDR while sampling at 20 MHz, and consumes 17.2 mW at 1.8 V supply.**

***Index Terms*—Split-CLS, scaled CMOS amplification technique, pipelined analog-to-digital converter, A/D, ADC, switched capacitor amplification, correlated level shifting, CLS, zero crossing based circuit, ZCBC, comparator based switched capacitor circuit, CBSC, dynamic zero crossing detector, ZCD.**

I. Introduction

SWITCHED capacitor amplification techniques are a nearly ubiquitous building block in the design of modern CMOS analog circuits. These methods allow designers to perform high accuracy, discrete time mathematical functions in the analog domain. The accuracy of this operation is often limited by finite loop gain, an effect commonly referred to as finite opamp gain error. Traditionally, opamp designs such as telescopic, folded-cascode, and compensated multi-stage [1] were used to provide the requisite high gains needed. However, as CMOS feature sizes continue to scale, the supply voltage and intrinsic device gain of the newer processes have decreased. In the regime of supply voltages below 1.8 V the use of traditional opamps becomes quite costly in terms of power, because the headroom requirements for proper biasing of all transistors in the opamp has not scaled as fast as supply voltage. The amplifier's transistors now demand most of the available supply voltage just to ensure high enough gain, leaving little or no voltage swing left for actual signal amplification. Less swing means that more power must be spent reducing noise by using larger capacitors and higher current in the active components. In modern supplies below 1 V, using these traditional opamps becomes unfeasible.

There are several circuit and calibration techniques that have been proposed as a solution to this challenge. An early method proposed was to employ gain-boosting amplifiers [2], [3]. While this will increase gain without sacrificing headroom, it comes at the cost of increased power and in deep submicron may still fail to free up sufficient headroom. Calibration is another technique that has been proposed in various ways [4], [5], with the general idea of using a low accuracy amplifier and calibrating out the error digitally. Calibration has proven effective, even in deep submicron, although it often comes at the price of increased power and complexity and often still doesn't solve the low output swing problem. In a circuit based approach, correlated double sampling (CDS) is a switched capacitor technique which samples the opamp finite gain error present at the input of the opamp and cancels it from the feedback path [6]. However, CDS comes with a number of drawbacks in terms of power, speed, noise, and complexity. Many of the drawbacks of CDS derive from the fact that it removes the error at the input of the opamp. This is the key insight of correlated level shifting (CLS), which introduces a method for cancelling the error in the feedback path at the output of the opamp [7]. CLS suffers from almost none of the drawbacks of CDS and also comes with the very significant advantage of extending the effective output swing of an amplifier well beyond its normal range.

A different approach to scaling challenges is to look at amplification methods other than opamps. [8] suggests a residue amplification method using a single transistor amplifier. While this approach is feasible for lower and medium resolution, it does require calibration and is not suitable for high resolution applications. Another option is to use a comparator and current source to charge towards and detect the feedback circuit's desired final output value, rather than explicitly settle to it with an opamp (Fig. 1). This is the core idea of zero-crossing based circuits (ZCBC), first proposed in [9], and expanded on in [10], [11]. This has proven its applicability over a range of design specifications in modern processes, achieving power efficiency in line with the state of the art. However, the open-loop nature of ZCBC operation makes this technique difficult to design in

Manuscript received April 13, 2010; revised June 29, 2010; accepted July 31, 2010. Date of publication October 14, 2010; date of current version December 03, 2010. This paper was approved by Guest Editor Boris Murmann. This work was supported by the Semiconductor Research Corporation and the Center for Design of Analog-Digital Integrated Circuits.

B. Hershberg and U. Moon are with the Department of Electrical Engineering and Computer Science, Oregon State University, Corvallis, OR 97331-5501 USA (e-mail: benjamin@hershberg.com).

S. Weaver was with Oregon State University, Corvallis, OR 97331-5501 USA, and is now with Intel Corporation, Hillsboro, OR 97124 USA.

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JSSC.2010.2073190

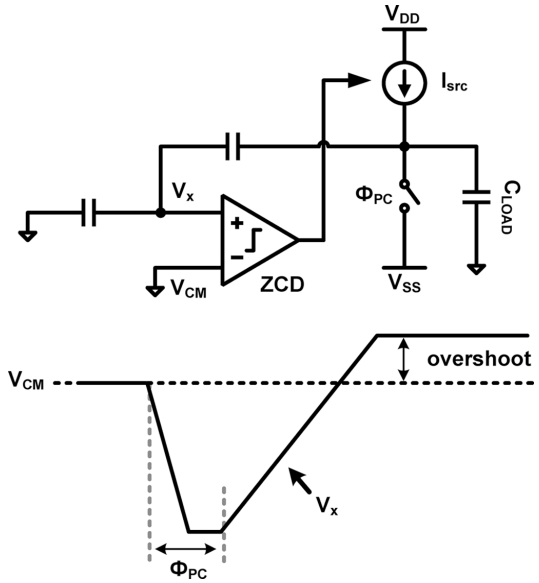

Fig. 1. Basic theory of operation for a zero-crossing based circuit (ZCBC). Rather than use an opamp to force $\mathrm{V_X}$ to $\mathrm{V_{CM}}$, we can use a zero-crossing detector (ZCD) and current source to charge the output and detect when $\mathrm{V_X} = \mathrm{V_{CM}}$.

a way that ensures sufficient linearity, particularly in a setting where significant tolerance to process, voltage, and temperature are important.

Although all these design approaches offer solutions in the regime of medium accuracy, analog techniques applicable to high accuracy remain limited. In this paper, we propose a technique for achieving efficient high accuracy amplification in modern processes. It was implemented as a pipelined ADC, which serves to further highlight the need for such a technique; although ADCs have scaled exceptionally well in medium resolution architectures (particularly low power SAR ADCs [12]–[14]), there is a notable scarcity of high resolution Nyquist converters implemented in deep submicron processes. The Split-CLS technique of this paper offers a solution to this technology gap.

This paper is organized as follows. Section II introduces the concept of Split-CLS, which is then extended to a specific implementation of Split-CLS incorporating ZCBC in Section III. Section IV addresses various implementation details of the prototype Split-CLS structure. Additionally, a dynamic ZCD designed to improve power efficiency of ZCBC systems is introduced in Section V. Section VI presents measured results taken from a fabricated test chip, and Section VII concludes.

II. Split-CLS

A. Basic CLS Structure

The basic CLS MDAC structure introduced in [7] and shown in Fig. 2 can be used to reduce finite opamp gain error and increase the opamp's useful output swing. After sampling the input (Φ_{S}), the MDAC switches into an amplification configuration (Φ_{A}). CLS occurs during this amplification phase in two steps: estimation (Φ_{EST}) and level shifting (Φ_{CLS}). In Φ_{EST} the opamp is connected directly to the output, which allows the CLS capacitor to sample an estimate of the correct output voltage with respect to the opamp output common mode (or any other analog reference). In $\Phi_{\mathrm{CLS}}, \mathrm{C_{CLS}}$ is connected between the MDAC output and the opamp output, which level shifts the opamp output back to $\mathrm{V_{CMO}}$. The opamp now only processes the error of the initial estimate (the signal has been cancelled in the feedback path), which will reduce the finite opamp gain error and requires only enough opamp output swing as is necessary to correct this error. Using the same set of assumptions and simplifications made in the derivation of (22) of [7] and adding an additional assumption that the opamp is single stage (which we will later see is true for our use), we see that the final effective gain $(\mathrm{A_{EFF}})$ of the structure is related to the opamp gain during Φ_{EST} and Φ_{CLS} as

$$A_{\mathrm{EFF}} \propto A_{\mathrm{EST}} \cdot A_{\mathrm{CLS}} \tag{1}$$

B. Split-CLS

A useful observation that we can make about CLS is that the gains $\mathrm{A_{EST}}$ and $\mathrm{A_{CLS}}$ do not necessarily need to be the same or even come from the same amplifier. Moreover, the amplifier requirements for Φ_{EST} and Φ_{CLS} differ. By splitting the amplifier in Fig. 2 into two separate amplifiers and then designing each amplifier with its specific requirements in mind, we can improve the overall performance of CLS in terms of power, speed, and accuracy. Fig. 3 shows a generalized single ended Split-CLS structure. AMP1 processes the full signal during Φ_{EST} and is directly connected to the output load. Therefore, for optimum performance it should have as large of an output swing as possible and high slewing capabilities. By contrast, AMP2 must only process the small error term remaining after Φ_{EST} and charges the output indirectly through $\mathrm{C_{CLS}}$; if $\mathrm{C_{CLS}}$ is sized appropriately, the output swing and slew requirements for AMP2 are much smaller than for AMP1.

Before looking at possible implementation choices for AMP1 and AMP2, a few more observations about Split-CLS are helpful. First, the opamp outputs are connected to different nodes, which comes with some advantages. First of all, we can move the switch at the output of AMP2 outside of the signal path. Switch $\mathrm{S_2}$ is used during Φ_{EST} to both sample $\mathrm{C_{CLS}}$ and short the opamp output. Without this shorting switch, AMP2's output would swing heavily during Φ_{EST} due to AMP1's settling of $\mathrm{V_X}$. This swinging would in turn affect the voltage at $\mathrm{V_X}$ via parasitic kickback through AMP2's input transistors, creating a second feedback loop and potential settling issues. Likewise, in the generalized form of Split-CLS, switch $\mathrm{S_1}$ must be used at the output of AMP1 during Φ_{CLS}, but as we will see later, if the device used for AMP1 can turn its output off internally, $\mathrm{S_1}$ is unnecessary.

Another difference between Split-CLS and CLS is the transient effects that occur during the transition from Φ_{EST} to Φ_{CLS}. In traditional CLS, at the beginning of Φ_{CLS} the amplifier output will quickly transition from the full-swing estimation voltage back to $\mathrm{V_{CMO}}$, causing a signal dependent glitch at node $\mathrm{V_{CLS}}$ that can degrade settling speed. The size of this glitch depends on the swing of the opamp during Φ_{EST} and the size of the opamp's internal capacitances. While it was mentioned in [7]

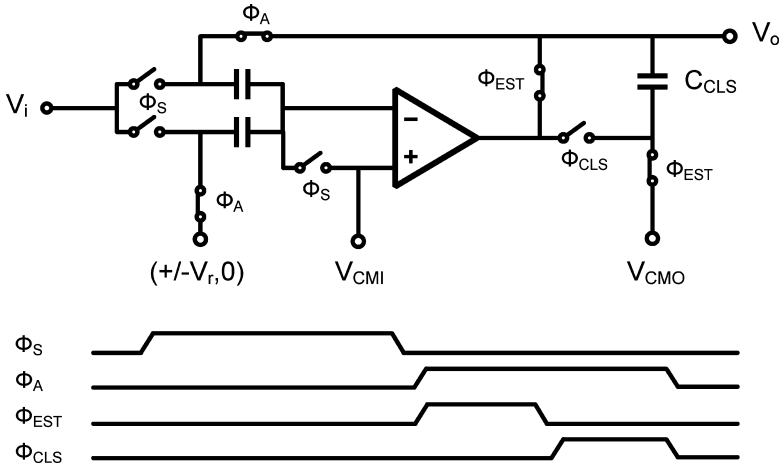

Fig. 2. Correlated level shifting (CLS).

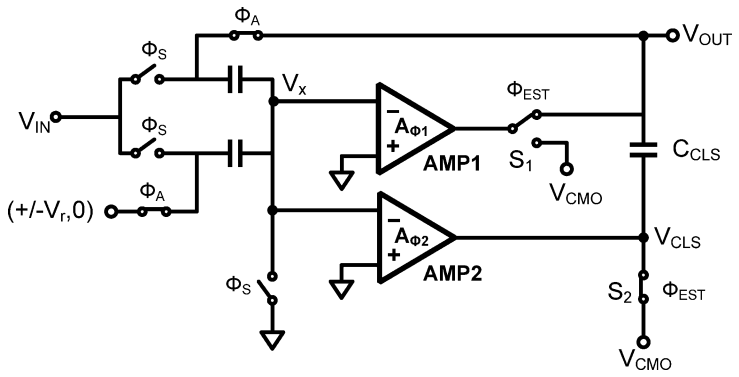

Fig. 3. Split-CLS structure.

that this glitch can be tolerated with careful sizing of the opamp transistors, this limits the opamp design space, and it is not easy to predict the actual parasitic effects of such subtle parameters. Split-CLS doesn't have this glitch problem because AMP2 is held during Φ_{EST} at the exact bias condition which we desire it to be at when beginning Φ_{CLS}; upon S_2 opening, AMP2 immediately begins settling the MDAC output where AMP1 left off.

III. Opamp/ZCBC Split-CLS

There are many possible amplifier topologies that could be promising choices for AMP1 and AMP2 in Split-CLS. For the remainder of this paper we will explore the design of one possible configuration, the Opamp/ZCBC Split-CLS structure.

A. AMP1

As mentioned in Section I, designing a high gain amplifier to have large output swing in deep sub-micron technologies using traditional opamp topologies is difficult. Ultimately, it will probably require a multi-stage approach so that the output stage can be a single PMOS/NMOS pair. The compensation required to make such an amplifier stable will place an upper bound on the best case power-to-speed ratio. On the other hand, if we consider some alternative amplification techniques, we may get better results. A ZCBC is a particularly attractive AMP1 choice. Unlike traditional opamps, the power requirement of the sense circuitry (the ZCD) is ideally unrelated to the size of the output load, and in practice only weakly related. For this reason, ZCBCs are capable of excellent slewing efficiencies. Another key benefit of using ZCBC amplification is that the current source and ZCD can be shut off after Φ_{EST}, which will save considerable power compared to using an opamp for AMP1. Since the current sources shut off internally, S_1 of Fig. 3 is no longer needed.

A major challenge in ZCBC designs is achieving high-accuracy amplification at high speed. Any signal-dependent effects in the ZCD decision time-delay, output current source, or switches will directly impact the overall accuracy. Slower operation reduces the influence of all of these problems, but ideally we would like AMP1 to charge the output very quickly. Deterministic errors are not the only challenge, however. The open-loop nature of ZCBC makes it susceptible to internal and external transient variations and integrated noise. Whereas an opamp in feedback can suppress these memory effects, a ZCBC is not always able to.

Ultimately, the strength of pure ZCBC appears to lie more in its high power efficiency at medium to high speeds with medium resolution than in the realm of high accuracy. For a conventional single-shot differential ZCBC implementation such as [11], at low speeds, high accuracy is less attractive because of the higher ratio of ZCD static current to current source dynamic current. At higher speeds, high accuracy is hard to achieve due to the linearity challenges mentioned previously. We can relax many of the accuracy challenges of ZCBC by treating it as a coarse,

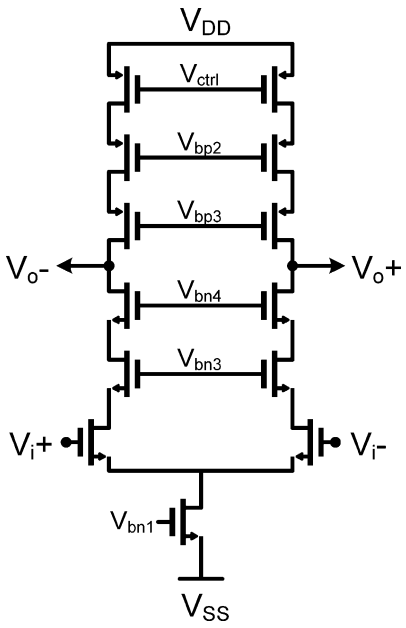

Fig. 4. Double cascode opamp used for AMP2.

fast, high efficiency charging device, and leave the high accuracy settling to an opamp (AMP2) which is better suited to such a task. From the traditional amplification perspective, this approach can be seen as a way to improve the slewing efficiency of an opamp. From the perspective of ZCBC, this can be seen as a method for extending the use of ZCBC into higher resolution applications by alleviating many of its critical issues through the application of linear feedback.

This idea of using a coarse charging device paired with an opamp was previously explored in [15] by connecting both coarse and fine charging devices directly to the output. Although this could potentially improve power efficiency, errors from low opamp gain and low output swing are not reduced like they are with CLS. In addition, the opamp must be designed so that it will not charge the output while the coarse charging device is operating. By contrast, AMP2 in Split-CLS is held in a fixed state during Φ_{EST} by simply shorting its output.

B. AMP2

The device chosen for AMP2 in the implemented Split-CLS structure is a double-cascode telescopic opamp, shown in Fig. 4. This structure's merits lie in its high speed and high gain relative to power consumption. It is much easier to guarantee stability for a single stage opamp than for a two-stage opamp. By trading output swing for gain with additional cascode transistors, high gain is obtained without the use of gain-boosting amplifiers or other techniques that would require additional power. Typically, the drawback of a telescopic opamp is that it has a small output swing due to the cascoding transistors and a low slew rate due to the lack of a high current output stage. However, as discussed earlier, AMP2 requires neither high output swing nor high slew rate. The opamp employs the switched capacitor common-mode feedback of [16].

The amount of error remaining at the end of Φ_{EST} and the size of $\mathrm{C_{CLS}}$ determine AMP2's output swing requirement. $\mathrm{C_{CLS}}$ also affects the effective gain of the overall structure, and for a single stage amplifier, increasing $\mathrm{C_{CLS}}$ will increase effective gain (at the cost of higher output load) [7]. Therefore, after sizing $\mathrm{C_{CLS}}$ to meet the output swing requirement, only if more effective gain is needed in the structure should the value of $\mathrm{C_{CLS}}$ be increased further.

C. Full Structure

The Opamp/ZCBC Split-CLS MDAC amplifier structure is presented in Fig. 5 with corresponding timing and waveform diagrams in Fig. 6. In addition to the opamp and ZCBC circuitry, the structure contains a set of capacitive DACs which are used to cancel ZCBC overshoot prior to the start of opamp settling (explained in Section IV).

During Φ_{S} the input capacitors sample the input while the ZCD and current sources are shut off and the opamp idles. At the beginning of Φ_{A} the amplification operation begins with a short pre-charge phase Φ_{PC}, during which several things occur simultaneously: the pre-charge switches $(\mathrm{S_{PC}})$ charge the output nodes to $\mathrm{V_{DD}}$ and $\mathrm{V_{SS}}$, the current sources $(\mathrm{I_N}, \mathrm{I_P})$ are turned on, the ZCD internally switches into a startup configuration (discussed in Section V), the opamp output and bottom plate of $\mathrm{C_{CLS}}$ is connected to a low-impedance reference through $\mathrm{S_{OP}}$, and the overshoot cancellation DACs $(\mathrm{C_{DAC}})$ are connected to the output. When the pre-charge switches open at the end of Φ_{PC} the current sources freely charge the output load with a linear ramp until the zero crossing of $\mathrm{V_{X+}}$ and $\mathrm{V_{X-}}$ is detected by the ZCD and the current sources are switched off. When the digital output of the ZCD flips, an asynchronous timing block is activated. This block facilitates the cancellation of the ZCBC overshoot by flipping the $\mathrm{S_{DAC}}$ switches to the alternate supply voltage. After the overshoot cancellation is settled, the timing block disconnects $\mathrm{C_{DAC}}$ from the output (in order to minimize the capacitive load that the opamp must drive) and opens the shorting switches $\mathrm{S_{OP}}$, allowing the opamp to begin freely settling to the final output voltage until the end of Φ_{A}.

IV. Implementation Details

A. Current Sources

One of the current sources used in the ZCBC implementation $(\mathrm{I_N})$ is shown in Fig. 7. The design of $\mathrm{I_P}$ is an analogous pMOS current source. Transistor $\mathrm{S_{ISRC}}$ is designed to operate more like a switch than a cascode device in order to maximize the ZCBC output swing. As a fast, coarse charging device, a single transistor current source provides sufficient linearity. By switching the current source at its drain, there will be some initial nonlinearity as the inversion layer of $\mathrm{S_{ISRC}}$ is created and the drain of $\mathrm{M_1}$ rises from $\mathrm{V_{SS}}$ to a voltage slightly below $\mathrm{V_O-}$. It is for this reason that digital signal $\mathrm{D_C}$ (generated by the ZCD) transitions low at the beginning of Φ_{PC} while the output is being shorted to

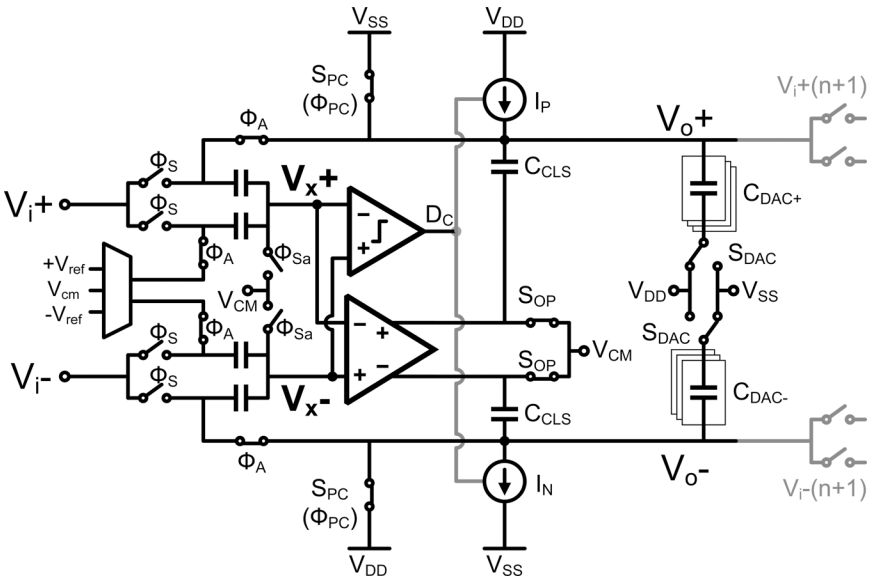

Fig. 5. Opamp/ZCBC Split-CLS structure.

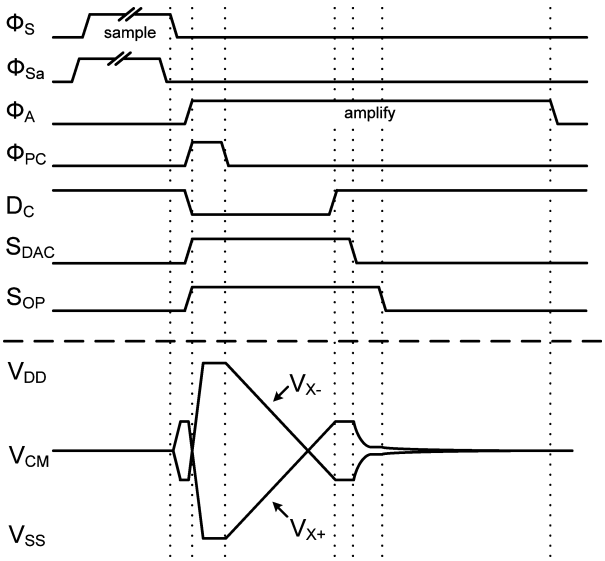

Fig. 6. Timing and waveforms of the Opamp/ZCBC Split-CLS structure.

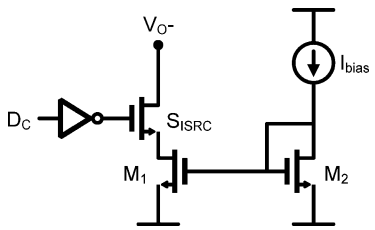

Fig. 7. ZCBC current source $\mathrm{I_N}$.

a supply, allowing the initial switching transients to resolve and ensuring a linear ramp on the output right after Φ_{PC} goes low.

B. Switches

Unlike opamp settling, where the amount of current (and any resulting IR drop) flowing through the switches in the MDAC will converge to zero as the output settles, in ZCBC the current through the switches remains high. Even as a coarse charging device, switch linearity during the ZCBC operation is something that must be addressed.

The bootstrapping switch of [17] was used to create a small signal-independent on-resistance for switches in the feedback path. All of the switches that are closed during Φ_{S} are bootstrapped. Even though the bottom plate sampling switches controlled by Φ_{Sa} have a fixed gate-to-source voltage, that voltage is only $\mathrm{V_{DD}}$-$\mathrm{V_{CM}}$. By bootstrapping the switch, $\mathrm{V_{GS}}$ is made twice as large, yielding an overall reduction in switching power compared to using a single transistor switch. Of the switches that are asserted in Φ_{A}, only the ones in the MDAC feedback loop are bootstrapped. The MDAC reference voltage switches used in Φ_{A} already have constant and sufficiently large $\mathrm{V_{GS}}$. $\mathrm{S_{OP}}$ is also outside of the feedback path and with fixed $\mathrm{V_{GS}}$, so it is implemented as a single-transistor switch.

C. Overshoot Cancellation

The primary challenge of pairing a ZCBC with an opamp is the differential output offset (or "overshoot") introduced by the finite time delay between when the inputs of the ZCD cross to when the current sources turn off. In a pure ZCBC pipelined ADC this overshoot is passed down the pipeline and the signal-independent portion of it will become a DC offset in the final digitized signal. This static offset can be canceled on a global level, and one such approach to this is shown in [11]. Depending on the ZCBC accuracy requirements, a global correction like this is not necessarily possible in the case of Split-CLS; the overshoot must be cancelled completely before being passed on to the next stage. This is because the opamp will attempt to settle such that $\mathrm{V_X}- = \mathrm{V_X}+$. Any differential DC offset at these nodes will be processed by the opamp as if it were error. In practice, this overshoot can be several hundred millivolts, and if the opamp is left to cancel it directly, we must sacrifice an inordinate amount of either opamp output swing or increased size of $\mathrm{C_{CLS}}$ in order to keep the opamp transistors in saturation—negating

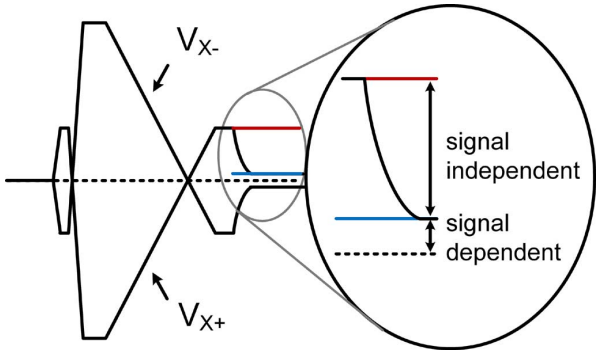

Fig. 8. All components of ZCBC overshoot will be processed by the opamp as error. To minimize this error, we must cancel the signal-independent portion of the overshoot as best as possible prior to level-shifting.

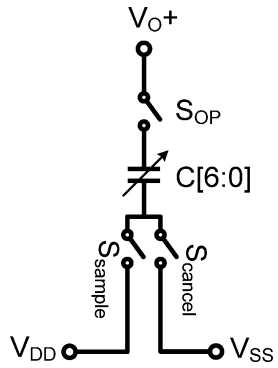

Fig. 9. Programmable overshoot cancellation DAC.

our previous assumptions and benefits related to the low-swing and low-slew requirements for the opamp.

As shown in Fig. 8, the overshoot consists of a signal-independent portion and a smaller signal-dependent portion. The signal-dependent portion is not easily predictable or cancellable, but the signal-independent portion is. Fig. 8 illustrates the cancellation method implemented in this paper—the explicit cancellation of the overshoot voltage. Cancellation with the digitally programmable capacitor DAC of Fig. 9 has the benefit of flexibility in a test environment, and was chosen over a single capacitor for this reason. Explicit cancellation also makes the Split-CLS structure compatible with integrators, where the accumulation of ZCBC overshoot can quickly saturate the integrator output. The DAC illustrated in Fig. 9 is $\mathrm{C_{DAC+}}$ of Fig. 5. $\mathrm{C_{DAC-}}$ is similar, but with $\mathrm{V_{DD}}$ and $\mathrm{V_{SS}}$ swapped on the lower reference nodes.

Rather than cancelling the overshoot explicitly, it is also possible to prevent it from occurring in the first place by designing the ZCD with an input referred offset so that it will trip earlier. Likewise, we could let the overshoot occur, but design the opamp with an input referred offset to match the magnitude of the overshoot at its inputs. Yet another option, suggested in [18], is to charge both $\mathrm{V_O+}$ and $\mathrm{V_O-}$ from the same direction during ZCBC operation. The resulting overshoots would ideally have the same magnitude and direction, with zero differential overshoot.

D. Asynchronous Timing Block

There are several signals that must be generated between the time that the ZCD transitions and the opamp begins settling the output. These signals are generated by an asynchronous timing block, shown in Fig. 10. The switch $\mathrm{S_{DAC}}$ in Fig. 5 is actually two switches: $\mathrm{S_{sample}}$ and $\mathrm{S_{cancel}}$ of Fig. 9. These switches are either nMOS or pMOS depending on the supply that they're connected to (which are opposite for $\mathrm{C_{DAC+}}$ and $\mathrm{C_{DAC-}}$), so all polarities must be generated. The asynchronous timing block allows the level shifting phase to be started automatically by the ZCD's output $(\mathrm{D_C})$, maximizing the time that the opamp has to settle and greatly simplifying the clocking scheme compared to [7].

V. Dynamic Zero Crossing Detector

The ZCD in a differential ZCBC system is typically statically biased while the current sources are ramping the output [9], [11]. A useful observation that we can make in this regard is that the ZCD's bandwidth requirements depend on the relative position of its inputs. When the inputs of the ZCD are close to crossing, we know that the ZCD must trip soon, so we would like to have a high bandwidth in order to minimize both the ZCD's absolute time-delay and variation in time-delay. By contrast, when the inputs are still far apart we can infer that the ZCD is not yet close to tripping, so there is very little benefit to having a high bandwidth in the detection circuit. Therefore, if we redistribute the power consumption of the ZCD to concentrate the use of current around the detection instant itself, we can optimize power efficiency and minimize static power dissipation.

A. Basic Structure and Operation

A simple way to dynamically scale the power of a ZCD is to adjust the tail current ($\mathrm{M_1}$ in the ZCD of Fig. 11). A more challenging question is how to detect when this current should be scaled. One possibility lies in the output nodes themselves; because the ZCD's gain $(\mathrm{A_{ZCD}})$ is finite, the outputs will begin to transition before the actual zero-crossing. The output itself provides us with a transitioning signal that always precedes the zero-crossing by a fixed amount.

Fig. 11 presents the basic structure and theory of operation of the proposed dynamic ZCD. Initially, during Φ_{PC}, tail transistor $\mathrm{M_1}$'s gate node $(\mathrm{V_B})$ is charged via switch $\mathrm{M_3}$ to a static bias. On the falling edge of Φ_{PC}, $\mathrm{V_B}$ is disconnected from this reference and left at a fixed but floating voltage. The inputs $(\mathrm{V_i+}, \mathrm{V_i-})$ are still relatively far apart, and if $\mathrm{A_{ZCD}}$ is sufficiently large, the outputs $\mathrm{V_O+}$ and $\mathrm{V_O-}$ are saturated at the maximum value of the ZCD's output swing. Node $\mathrm{V_O-}$ will remain close to $\mathrm{V_{SS}}$ in this saturated state until

$$A_{\mathrm{ZCD}}(V_{i+} - V_{i-}) \approx (V_{o+} - V_{o-}) \tag{2}$$

At this point the voltage at $\mathrm{V_O-}$ will begin to rise and, via the feedback capacitor $\mathrm{C_{FB}}$, the floating node $\mathrm{V_B}$ will also rise by some proportional amount, increasing the tail current $(\mathrm{I_{tail}})$. Around the detection instant, the $\mathrm{g_m}$ of the ZCD is at a maximum. Once the ZCD decision is registered by the dynamic latch

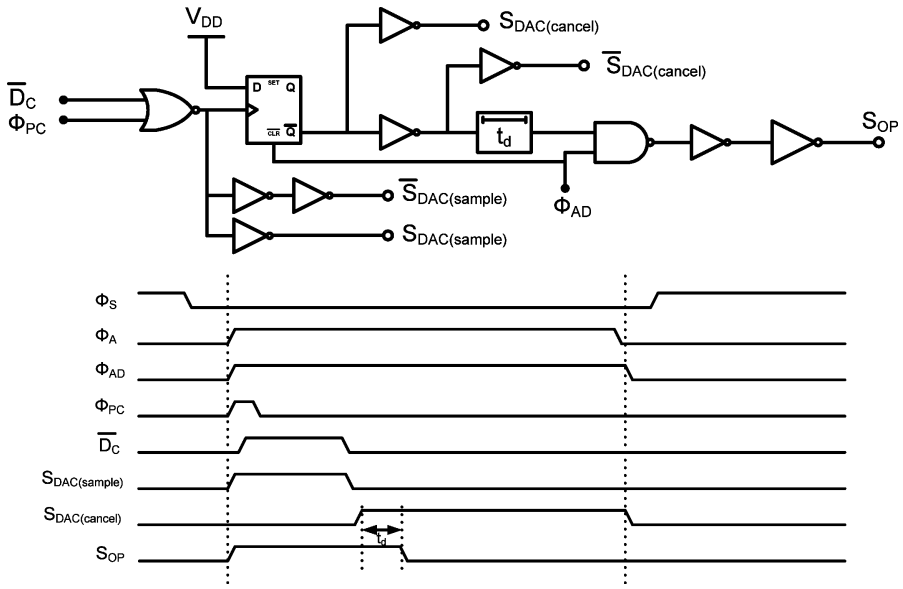

Fig. 10. Asynchronous timing block used to facilitate the transition from ZCBC to opamp operation.

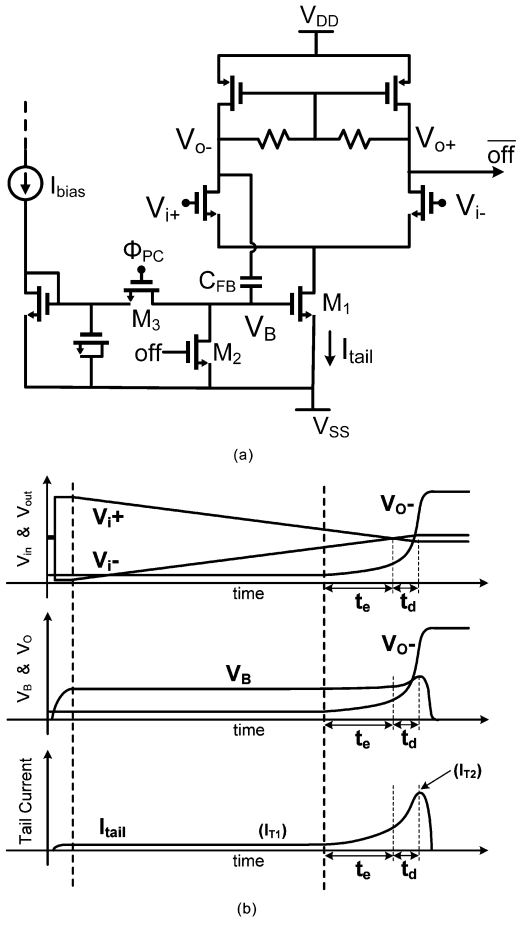

Fig. 11. Differential dynamic zero-crossing detector. (a) Simplified schematic of proposed ZCD. (b) Example ZCD waveforms demonstrating dynamic rationing of power.

at the output of $\mathrm{V_O}+$, switch $\mathrm{M_2}$ connects $\mathrm{V_B}$ to $\mathrm{V_{SS}}$ and shuts off the ZCD.

B. Implementation Details

The complete schematic of the dynamic ZCD is given in Fig. 12. Specific implementation considerations are discussed in the following sub-section.

i) Startup Behavior: For the ZCD to be properly biased at the end of Φ_{PC}, both $\mathrm{V_B}$ and $\mathrm{V_O}-$ must be at a constant, settled voltage. At the beginning of Φ_{PC} the ZCD is brought out of shutoff, and $\mathrm{V_O}- = \mathrm{V_{DD}}$. $\mathrm{V_O}-$ must then settle during the short Φ_{PC} phase to its initial steady-state voltage, $\mathrm{V_{O-(sat)}}$ (which will be a few hundred millivolts above $\mathrm{V_{SS}}$). Incomplete settling of $\mathrm{V_O}-$ is problematic: $\mathrm{V_O}-$ is coupled to $\mathrm{V_B}$ through $\mathrm{C_{FB}}$ and if $\mathrm{V_O}-$ continues to fall after the end of Φ_{PC}, $\mathrm{V_B}$ and $\mathrm{I_{tail}}$ will drop as well. This scenario is very likely, because the initial bandwidth of the ZCD during Φ_{PC} is intentionally low. Due to these concerns, $\mathrm{V_O}-$ is pre-charged to $\mathrm{V_{pco-}}$. Ideally, $\mathrm{V_{pco-}}$ would be equal to $\mathrm{V_{O-(sat)}}$, but for simplicity we can set it to any convenient voltage smaller than $\mathrm{V_{O-(sat)}}$, such as $\mathrm{V_{SS}}$. This ensures fast and signal-independent settling and guarantees that $\mathrm{I_{tail}}$ will be greater than zero. The resulting transient ZCD waveforms are shown in Fig. 13.

The initial voltage of the other output node, $\mathrm{V_O}+$, is also critical. During Φ_{PC} $\mathrm{V_O}+$ is set to $\mathrm{V_{DD}}$ for two reasons. First, $\mathrm{V_O}+$ must remain above the pMOS threshold voltage $(\mathrm{V_{TP}})$ of the dynamic output latch. Second, with $\mathrm{V_O}-$ simultaneously held at $\mathrm{V_{pco-}}$, this will provide a good approximation of the common mode feedback (CMFB) bias condition that will exist right after Φ_{PC} ends. It is important to match this CMFB condition as best as possible in order to minimize any settling ripple which could occur. Any transient ripple larger than $\mathrm{V_{TP}}$ will also cause the dynamic latch to trip in error. Another such scenario where this could occur is in a design where the CMFB bandwidth is not high enough to track the dynamic increase in tail current.

ii) Detection and Shutoff: If not for the presence of the shutoff switches controlled by $\mathrm{S_{off}}$, the ZCD's internal voltages would continue to change even after the ZCD trips because the

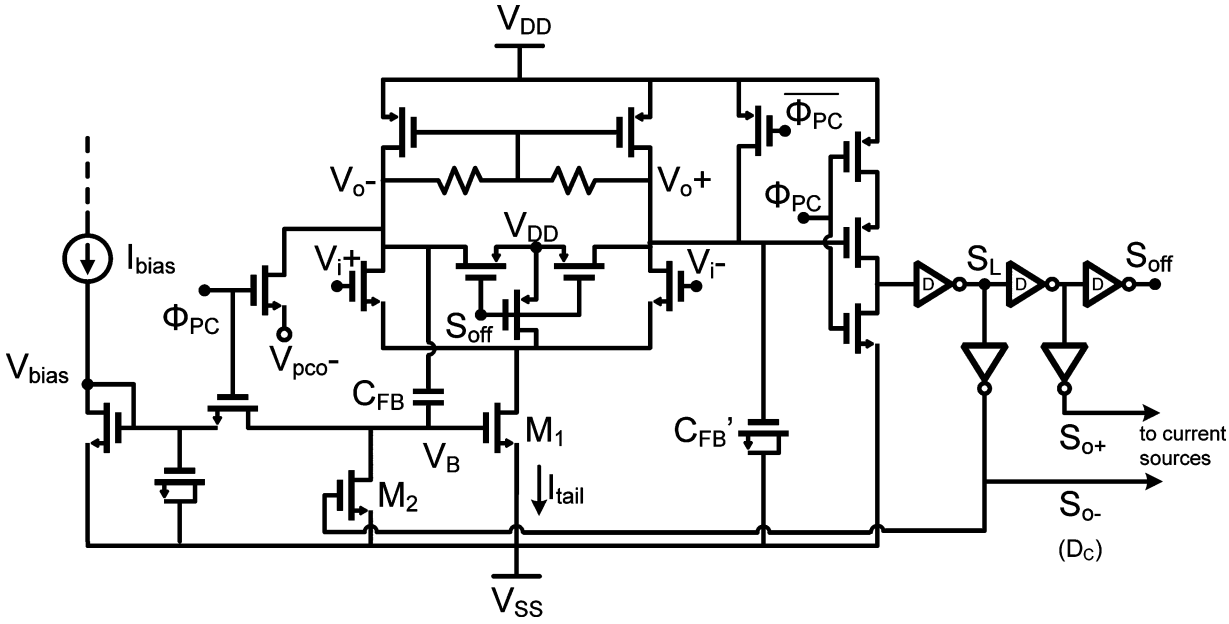

Fig. 12. Detailed view of the dynamic ZCD, including pre-charge switches, output load matching, dynamic output latches, and shutoff switches.

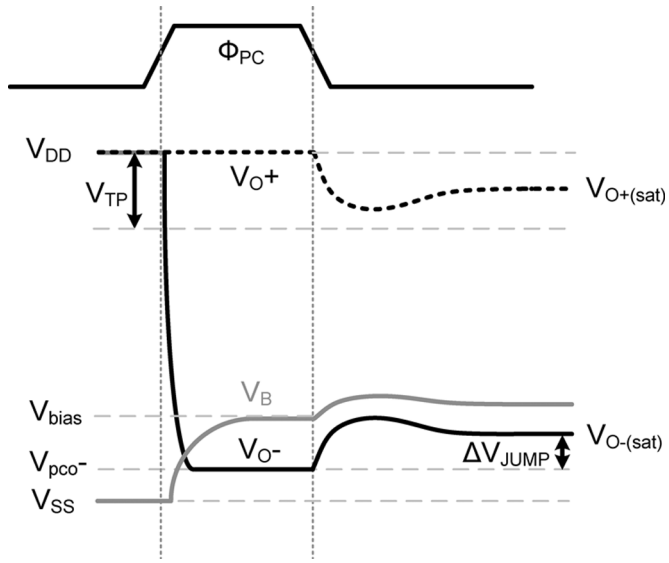

Fig. 13. Transient behavior of $\mathrm{V_O}+, \mathrm{V_O}-$, and $\mathrm{V_B}$ during and immediately after pre-charge. To ensure proper operation, $\mathrm{V_O}+$ must not dip below $\mathrm{V_{TP}}$ and $\mathrm{V_B}$ must stay above $\mathrm{V_{bias}}$ after Φ_{PC} goes low.

CMFB will begin to pull $\mathrm{V_O}+$ and $\mathrm{V_O}-$ toward $\mathrm{V_{DD}}$ after M1 is turned off. To ensure that the ZCD doesn't affect the charge at nodes $\mathrm{V_X}+/\mathrm{V_X}-$ (in Fig. 5) except during its own operation, the internal nodes of the input amplifier are brought to $\mathrm{V_{DD}}$ immediately after the zero-crossing is detected, and held there until the next pre-charge phase.

iii) Dynamic Latches: After the zero-crossing, the ZCD decision must be latched before $\mathrm{V_O}+$ is pulled back to $\mathrm{V_{DD}}$. Dynamic latches were used to accomplish this. The dynamic latch connected to $\mathrm{V_O}+$ is shown explicitly, and the rest are denoted with a "D". Alternating pull-high and pull-low latch styles are used in the output chain. During Φ_{PC} the latches are reset.

iv) Symmetric Loading: $\mathrm{V_O}+$ and $\mathrm{V_O}-$ as shown in Fig. 11 have unbalanced capacitive loads, which affect the transient behavior of the CMFB and time-delay of the ZCD. We can better balance the outputs by adding a small capacitance $\mathrm{C_{FB}}$' to $\mathrm{V_O}+$. Even more importantly, the dip in $\mathrm{V_O}+$ directly following pre-charge due to settling can be reduced by the additional capacitance of $\mathrm{C_{FB}}$', ensuring that $\mathrm{V_O}+$ will remain above $\mathrm{V_{TP}}$.

C. Design Considerations

The ramp rate of the inputs, the values of $\mathrm{A_{ZCD}}, \mathrm{C_{FB}}$ and $\mathrm{V_{bias}}$, and the pre-charge value of $\mathrm{V_O}-$ are inter-dependent design variables which all factor into how much power savings can ultimately be achieved with the dynamic ZCD versus a more typical ZCD. We will begin by looking at the relation between ramp rate and $\mathrm{A_{ZCD}}$. As shown in Fig. 11(b), the time when $\mathrm{V_O}+$ begins rising to when the zero-crossing occurs is denoted $\mathrm{t_e}$, and the time delay between the zero-crossing moment and the output of the ZCD latching is denoted $\mathrm{t_d}$. Considering the condition defined in (2) when the outputs begin transitioning, the value of $\mathrm{t_e}$ can be defined as

$$t_e = \frac{V_{o+(\mathrm{sat})} - V_{o-(\mathrm{sat})}}{A_{\mathrm{ZCD}} \cdot \frac{d(V_{i+} - V_{i-})}{dt}} \tag{3}$$

where $\mathrm{V_{O(sat)}}$ is the saturated full-swing value of $\mathrm{V_O}$. If $\mathrm{t_e}$ is larger than some optimal value, we are increasing the bandwidth earlier than necessary, because the value of $\mathrm{t_d}$ depends mainly on the bandwidth of the ZCD at the zero-crossing instant. If $\mathrm{t_e}$ is smaller than this optimal value, we will not have peak bandwidth at the zero-crossing instant due to practical speed limitations of the ZCD itself. In other words, the lower bound of $\mathrm{t_e}$ depends mostly on internal design constraints rather than external factors such as ramp rate. Because of this, we can approximate the optimal value of $\mathrm{t_e}$ for a given design to be constant with respect to

different ramp rates. From (3), the relation between A_{ZCD} and the ramp rate now becomes

$$\frac{1}{A_{\mathrm{ZCD}}} = c_1 \cdot \frac{dV_i}{dt} \tag{4}$$

where c_1 is a constant. The conclusion that we can draw from (4) is that for an optimum design, A_{ZCD} and ramp rate are tightly coupled variables, and the slower the ramp rate, the larger A_{ZCD} should be made. If A_{ZCD} is not made large enough to satisfy (3) and (4), then power is not being optimally utilized.

Not only does efficiency depend on how optimally we distribute the power consumption in time, it also depends on the ratio between the initial ZCD current shortly after pre-charge (I_{T1} of Fig. 11(b)) to the final ZCD current when it latches (I_{T2} of Fig. 11(b)). Using the simplified square-law model for a MOSFET in saturation, this ratio is

$$\frac{I_{T1}}{I_{T2}} = \left(\frac{(C_{\mathrm{FB}}+C_P)(V_{\mathrm{bias}}-V_{\mathrm{TN}})+C_{\mathrm{FB}} \cdot \Delta V_{\mathrm{jump}}}{(C_{\mathrm{FB}}+C_P)(V_{\mathrm{bias}}-V_{\mathrm{TN}})+C_{\mathrm{FB}}(V_{\mathrm{trip}}-V_{\mathrm{pco-}})} \right)^2 \tag{5}$$

C_{FB} is the sum of all explicit and parasitic capacitances between V_B and V_O-, C_P is a capacitor from V_B to V_{SS} that represents all other parasitic loading on V_B, V_{pco-} is the pre-charge value of V_O-, $\Delta V_{jump} = V_{O-(sat)} - V_{pco-}$, and V_{trip} is the value of V_O- when the ZCD latches. To maximize efficiency, we want to minimize I_{T1}/I_{T2}. From (5) we see that this can be done foremost by minimizing $V_{bias} - V_{TN}$ and ΔV_{jump} so that they are much smaller than $V_{trip} - V_{pco-}$ and then by increasing the size of C_{FB} relative to C_P.

In reality, there are many other considerations which play into the overall efficiency of the dynamic ZCD. For example, the assumption that A_{ZCD} is constant is an approximation, because current is being dynamically changed. While this discussion has explored general design concepts, the inter-dependency of the design variables and time-varying nature of important circuit parameters means that no simple analytical model is available and transient simulations of the dynamic ZCD are a vital part of the design process.

D. Dynamic versus Static Comparison

Much of the dynamic ZCD structure exists to deal with issues related to the dynamic biasing. If we wish to compare the dynamically biased ZCD to an equivalent statically biased ZCD, simply removing C_{FB} will not provide a functional or fair comparison. As we modify the structure to operate with static biasing, we will eventually arrive at a design nearly identical to the ZCD of [11]. In simulation, these two ZCDs were compared within the context of this work: for an input ramp duration of 7ns, total period of 50 ns, and $t_d = 410$ ps, the average power dissipation is 51 μW for the dynamic ZCD and 175 μW for the static ZCD. The tail current of the dynamic ZCD scales from $I_{T1} = 102$ μA when the inputs are far apart to a peak current of $I_{T2} = 425$ μA at the detection instant.

This comparison is by no means a generalized conclusion, because it was done within the context of a design that does not require high accuracy, high supply rejection, or low noise. In designs where these factors are more critical, the relative efficiencies may be different. In particular, the floating gate of the tail source transistor M_1 makes the Dynamic ZCD sensitive to external transient variations, and may be unsuitable for applications which demand high power supply rejection.

TABLE I
PERFORMANCE SUMMARY

Technology	0.18µm CMOS			
Supply Voltage	1.8V			
Input Voltage Range	1.4V			
Sampling Frequency	f_S = 10 MHz		f_S = 20 MHz	
ENOB	11.3b		11.1b	
SNR	69.6dB		68.3dB	
SNDR	69.5dB		68.3dB	
SFDR	78.8dB		76.3dB	
Power (analog/digital)	7.2mW	1.2mW	15.0mW	2.2mW
FoM	343.5 fJ/step		405.5 fJ/step	

VI. EXPERIMENTAL RESULTS

A pipelined ADC incorporating the Split-CLS structure (with dynamic ZCD) was fabricated in a 0.18 μm CMOS technology. Designed for testability and proof-of-concept, the prototype ADC is composed of 10 identical 1.5b/stage pipeline stages followed by a 1.5b flash backend. The telescopic opamp was designed to maintain better than 70 dB open-loop gain for a 300 mV differential output swing. Including the contribution of the ZCBC amplification, the total effective open-loop gain of the Split-CLS structure was designed to be more than 110 dB. Measurement results, previously given in [19], are presented in Table I and Fig. 15. The SFDR is limited by even harmonics that were not present on a previous version of the test board, which suggests that the source of error may originate from outside of the chip. The ADC maintains better than 66 dB SNDR for f_{in} up to $f_{Nyquist}$, and the roll-off above this frequency seen in Fig. 15(b) is limited primarily by sampling jitter, which can be improved by more careful design of the input clock buffer, clock generation circuitry, and sampling network.

The overshoot cancellation DACs were implemented with independent digital controls for each stage as well as independent control of the two DACs within each stage. The high bandwidth of the dynamic ZCD helped to create a very consistent amount of overshoot across all stages of the pipeline. All measured results presented in this paper were taken with a single static DAC code applied to all DACs in the pipeline. This consistency across all stages indicates that the overshoot cancellation can be implemented as a single global control in the future.

The test chip did not support disabling of the opamps, so a characterization of the dynamic ZCD without the influence of the opamp was not possible. However, the high tolerance to certain internal and external variation in t_d provided by the high bandwidth at the critical zero-crossing instant was measurable. In test, the ZCD bias current I_{bias} was varied between 5 μA and 50 μA with no observable change in performance besides power, with all other controls held constant.

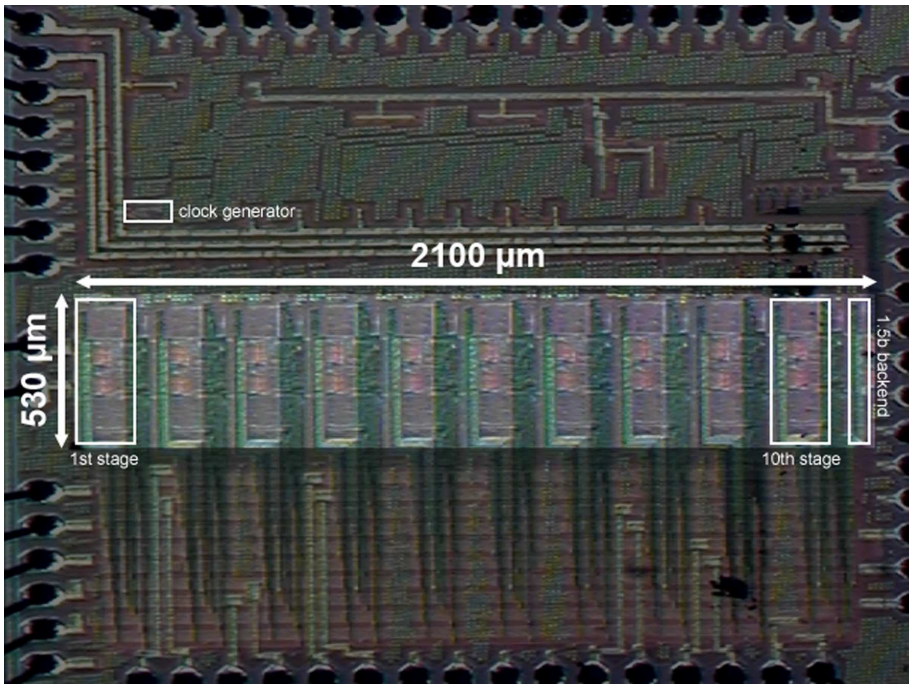

Fig. 14. Die micrograph.

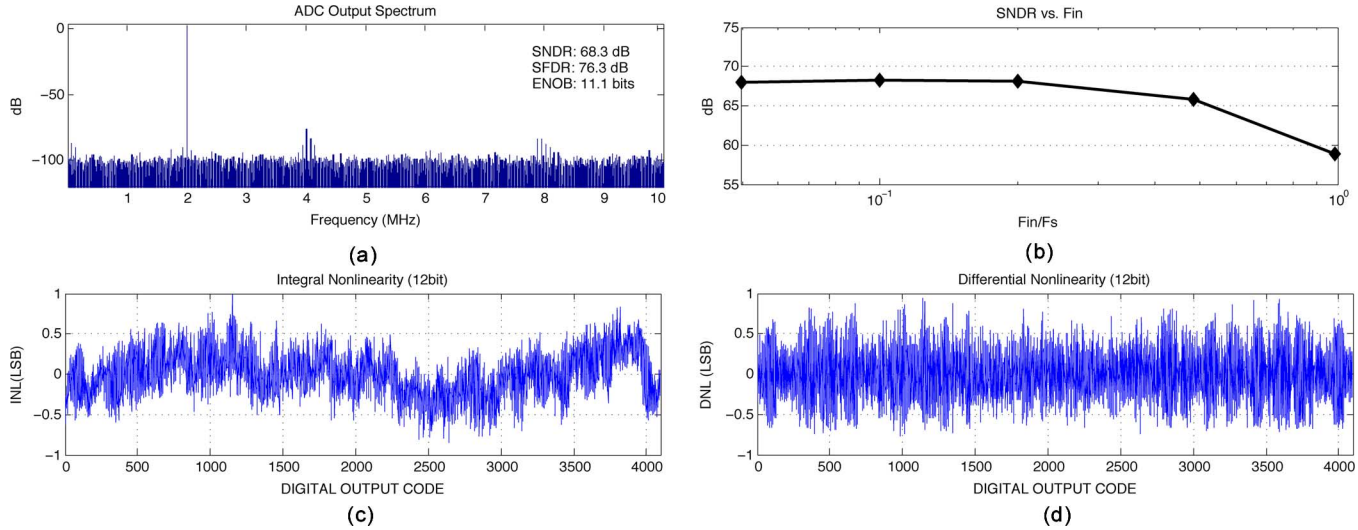

Fig. 15. Measured results. (a) ADC output spectrum for $f_s = 20$ MHz (b) SNDR versus f_{in} up to $2f_{Nyquist}$ (c) INL (d) DNL.

The FoM is 344 fJ/conversion-step at $f_s = 10$ MHz and 406 fJ/conversion-step at 20 MHz. Application of common design techniques such as stage scaling, multi-bit quantization, and opamp sharing, as well as the removal of the cancellation DACs in favor of a global ZCD input offset control could be applied to significantly extend power efficiency.

VII. Conclusion

In this paper we have introduced the concept of Split-CLS to enable high accuracy switched capacitor amplification in modern processes by creating extremely high effective gain. In the prototype implementation, power efficiency is accomplished by using a ZCBC for fast, coarse charging of the output load followed by fine settling with a low power, high gain double cascode telescopic opamp. From the perspective of ZCBC systems, the implemented Split-CLS structure allows us to relax design requirements and increase the robustness of ZCBCs by the application of linear feedback. The dynamic ZCD introduced in this paper provides further efficiency improvements to ZCBC systems.

Acknowledgment

The authors would like to thank Jazz Semiconductor for providing fabrication.

References

[1] P. R. Gray, P. J. Hurst, S. Lewis, and R. G. Meyer, *Analysis and Design of Analog Integrated Circuits*. New York: Wiley, 2001, pp. 425–452.

[2] K. Bult and G. Geelen, "A fast-settling CMOS op amp for SC circuits with 90-dB DC gain," *IEEE J. Solid-State Circuits*, vol. 25, no. 6, pp. 1379–1384, Dec. 1990.

[3] C. Yun, Gray, P. R. Gray, and B. Nikolic, "A 14-b 12-MS/s CMOS pipeline ADC with over 100-dB SFDR," *IEEE J. Solid-State Circuits*, vol. 39, no. 12, pp. 2139–2151, Dec. 2004.
[4] B. Murmann and B. E. Boser, "A 12-bit 75-MS/s pipelined ADC using open-loop residue amplification," *IEEE J. Solid-State Circuits*, vol. 38, no. 12, pp. 2040–2050, Dec. 2003.
[5] A. N. Karanicolas, H. S. Lee, and K. L. Barcrania, "A 15-b 1-Msample/s digitally self-calibrated pipeline ADC," *IEEE J. Solid-State Circuits*, vol. 28, no. 12, pp. 1207–1215, Dec. 1993.
[6] C. Enz and G. C. Temes, "Circuit techniques for reducing the effects of op-amp imperfections: Autozeroing, correlated double sampling, and chopper stabilization," *Proc. IEEE*, vol. 84, no. 11, pp. 1584–1614, Nov. 1996.
[7] B. R. Gregoire and U. Moon, "An over-60 dB true rail-to-rail performance using correlated level shifting and an opamp with only 30 dB loop gain," *IEEE J. Solid-State Circuits*, vol. 43, no. 12, pp. 2620–2630, Dec. 2008.
[8] J. Hu, N. Dolev, and B. Murmann, "A 9.4-bit, 50-MS/s, 1.44-mW pipelined ADC using dynamic source follower residue amplification," *IEEE J. Solid-State Circuits*, vol. 44, no. 4, pp. 1057–1066, Apr. 2009.
[9] J. K. Fiorenza *et al.*, "Comparator-based switched-capacitor circuits for scaled CMOS technologies," *IEEE J. Solid-State Circuits*, vol. 41, no. 12, pp. 2658–2668, Dec. 2006.
[10] L. Brooks and H. S. Lee, "A zero-crossing-based 8-bit 200 MS/s pipelined ADC," *IEEE J. Solid-State Circuits*, vol. 42, no. 12, pp. 2677–2687, Dec. 2007.
[11] L. Brooks and H. S. Lee, "A 12b, 50 MS/s, fully differential zero-crossing based pipelined ADC," *IEEE J. Solid-State Circuits*, vol. 44, no. 12, pp. 3329–3343, Dec. 2009.
[12] C. Liu *et al.*, "A 10b 100 MS/s 1.13 mW SAR ADC with binary-scaled error compensation," in *IEEE ISSCC 2010 Dig. Tech. Papers*, Feb. 2010, pp. 386–387.
[13] M. Yoshioka *et al.*, "A 10b 50 MS/s 820 μW SAR ADC with on-chip digital calibration," in *IEEE ISSCC 2010 Dig. Tech. Papers*, Feb. 2010, pp. 384–385.
[14] W. Liu, P. Huang, and Y. Chiu, "A 12b 22.5/45 MS/s 3.0 mW 0.059 mm^2 CMOS SAR ADC achieving over 90 dB SFDR," in *IEEE ISSCC 2010 Dig. Tech. Papers*, Feb. 2010, pp. 380–381.
[15] C. T. Peach, U. Moon, and D. J. Allstot, "An 11.1 mW 42 MS/s 10 b ADC with two-step settling in 0.18 μm CMOS," *IEEE J. Solid-State Circuits*, vol. 45, no. 2, pp. 391–400, Feb. 2010.
[16] D. A. Johns and K. Martin, *Analog Integrated Circuit Design*. New York: Wiley, 1997, p. 291.
[17] M. Dessouky and A. Kaiser, "Very low-voltage digital-audio $\Delta\Sigma$ modulator with 88-dB dynamic range using local switch bootstrapping," *IEEE J. Solid-State Circuits*, vol. 36, no. 3, pp. 349–355, Mar. 2001.
[18] T. Musah and U. Moon, "Pseudo-differential zero-crossing-based circuit with differential error suppression," in *Proc. ISCAS 2010*, Jun. 2010, pp. 1731–1734.
[19] B. Hershberg, S. Weaver, and U. Moon, "A 1.4 V signal swing hybrid CLS-opamp/ZCBC pipelined ADC using a 300 mV output swing opamp," in *IEEE ISSCC 2010 Dig. Tech. Papers*, Feb. 2010, pp. 302–303.

Benjamin Hershberg received the H.B.S. degree in electrical engineering and the H.B.S. degree in computer engineering from Oregon State University, Corvallis, in 2006, and is currently pursuing the Ph.D. degree in electrical and computer engineering at Oregon State University.

His research interests include scalable and low power amplification techniques, analog to digital conversion, and mixed-signal systems.

Skyler Weaver received the H.B.S. degree in electrical engineering and the Ph.D. degree in electrical and computer engineering from Oregon State University, Corvallis, in 2006 and 2010, respectively.

He is currently with Intel Corporation, Hillsboro, OR. His research interests include highly scalable and synthesizable analog-to-digital converters.

Un-Ku Moon (S'92–M'94–SM'99–F'09) received the B.S. degree from the University of Washington, Seattle, in 1987, the M.Eng. degree from Cornell University, Ithaca, NY, in 1989, and the Ph.D. degree from the University of Illinois at Urbana-Champaign in 1994.

He has been with the School of Electrical Engineering and Computer Science, Oregon State University, Corvallis, since 1998. Before joining Oregon State University, he was with Bell Laboratories from 1988 to 1989, and from 1994 to 1998. His technical contributions have been in the area of analog and mixed-signal circuits including high linearity filters, timing recovery, PLLs, data converters, and low voltage circuits for CMOS.

Dr. Moon has served as an Associate Editor of the IEEE JOURNAL OF SOLID-STATE CIRCUITS (JSSC) and the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS II: ANALOG AND DIGITAL SIGNAL PROCESSING, as the Editor-in-Chief and Deputy Editor-in-Chief of the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS II: EXPRESS BRIEFS, and on the Technical Program Committee of the IEEE Custom Integrated Circuits Conference. He also served on the IEEE Solid-State Circuits Society (SSCS) Administrative Committee (AdCom) and the IEEE Circuits and Systems Society (CASS) Board of Governors (BoG) as the SSCS representative to CASS. He is currently a Distinguished Lecturer of the IEEE SSCS and serves as the Editor-in-Chief of the IEEE JSSC. He also serves on the Technical Program Committee of the IEEE International Solid-State Circuits Conference and the IEEE VLSI Circuits Symposium.