

Ring Amplifiers for Switched Capacitor Circuits

Benjamin Hershberg, Skyler Weaver, Kazuki Sobue, Seiji Takeuchi, Koichi Hamashita, Un-Ku Moon

Abstract—In this paper the fundamental concept of ring amplification is introduced and explored. Ring amplifiers enable efficient amplification in scaled environments, and possess the benefits of efficient slew-based charging, rapid stabilization, compression-immunity (inherent rail-to-rail output swing), and performance that scales with process technology. A basic operational theory is established, and the core benefits of this technique are identified. Measured results from two separate ring amplifier based pipelined ADCs are presented. The first prototype IC, a simple 10.5-bit, 61.5dB SNDR pipelined ADC which uses only ring amplifiers, is used to demonstrate the core benefits. The second fabricated IC presented is a high-resolution pipelined ADC which employs the technique of Split-CLS to perform efficient, accurate amplification aided by ring amplifiers. The 15-bit ADC is implemented in a $0.18\ \mu\text{m}$ CMOS technology and achieves 76.8 dB SNDR and 95.4 dB SFDR at 20 Msps while consuming 5.1 mW, achieving a FoM of 45 fJ/conversion-step.

Index terms - ring amplification, ring amplifier, ring amp, ringamp, RAMP, Split-CLS, CLS, correlated level shifting, switched-capacitor, nanoscale CMOS, scaling, scalability, analog to digital conversion, A/D, analog to digital converter, ADC, low power, rail-to-rail, slew-based, stabilized ring oscillator, high resolution.

I. INTRODUCTION

In many ways, the task of performing amplification in switched capacitor CMOS circuits is at a crossroads. The effects of technology scaling have made it increasingly difficult to implement accurate, efficient amplifiers from topologies that were conceived of at a time when 2.5V supplies were considered low-voltage, and the intrinsic properties of transistors were quite different from that of a 14nm FinFET [1]. And yet, for each scaling challenge that has arisen, a multitude of techniques which seek to solve it have emerged, and there are indeed many highly effective digital correction, gain-enhancement, and output-swing enhancement techniques now available to designers. However, the ability of these techniques to deliver favorable amplifier scaling characteristics in actual fabricated designs has fallen somewhat short of expectations. For example, the ADC performance surveys conducted in [2] and [3] indicate that there is both a notable scarcity of high-resolution ADCs implemented in nanoscale CMOS as well as a progressive decline in power efficiency for those that are reported. These observations are also a reflection of the net effect of opamp scaling challenges on system-level performance, since amplification is used in almost all high-resolution ADC architectures. By contrast, certain low and medium resolution ADC architectures that do not use opamps (such as SAR ADCs) have scaled very well into modern technologies.




Fig. 1: Fundamental structure of a ring amplifier. A ring amplifier is created by splitting a ring oscillator into two signal paths and embedding a different offset in each path. When placed in switched-capacitor feedback, a set of internal mechanisms generate stability and allow the oscillator to be used as an amplifier.

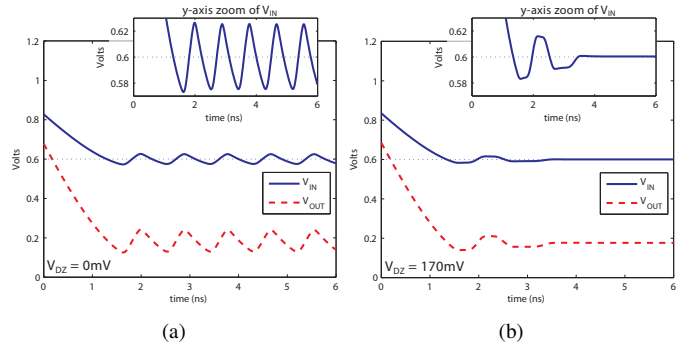


Fig. 2: Input and output charging waveforms of Fig. 1 when placed in the switched capacitor feedback structure of Fig. 3. In (a), when $V_{DZ} = 0\text{mV}$, the ringamp is functionally identical to a three-inverter ring oscillator. In (b), the dead-zone is set large enough to generate stability ($V_{DZ} = 170\text{mV}$) and the ringamp functions as an amplifier.

For the most part, this failure-to-scale seems to have happened because the underlying structure - an opamp - is fundamentally ill-suited to scaling. Applying additional techniques may enable an opamp to function in nanoscale environments, but it won't grant it the ability to scale at the same pace as digital performance improvements. A truly scalable amplifier must operate in a way that implicitly uses the characteristics of scaled CMOS to its advantage, transforming potential weaknesses into inherent strengths. Because technology scaling is deliberately designed to favor the time-domain world of high-speed digital, such scalable analog techniques will likely be found in this realm as well. In order to fully exploit the abilities of a transistor, the biasing and small-signal properties of the device must be viewed as highly coupled, time-dependent variables which can be applied as feedback to each other with respect to time. In this




Fig. 4: Example ring amplifier operation for an exaggerated design biased at the edge of stability, showing the three key phases of operation: 1) initial ramping, 2) stabilization, and 3) steady-state.

behind ring amplification is considerably more complex. The steady-state, small signal, and transient characteristics of a ring amplifier are highly co-dependent, and as such, its behavior cannot be fully explained by considering each operational domain (DC, AC, transient) separately, as is often the preferred approach in opamp design. Computer aided simulation is necessarily a major component of ringamp design, and for this reason a good conceptual intuition is perhaps of more practical value than a strict analytic model. It is in this conceptual manner that we will proceed, by breaking the theory of operation down to several simple sub-concepts. First, ringamp operation can be subdivided with respect to different phases of operation in time: slewing, stabilization, and steady-state. Second, the theory of operation within each phase can be reduced to a chain of cause-and-effect mechanisms.

To illustrate key concepts, we will use the exaggerated charging waveform of Fig. 4 (taken from the ring amplifier of Fig. 3) that has been designed with relatively low bandwidth, excessive drive current, and a dead-zone size that biases the ringamp right at the edge of stability. Although one would never wish to make a real design in this way, as a teaching example it is quite useful. V_{CMX} is set to 0.6V, and thus the ideal settled value of V_{IN} will also be 0.6V. For the sake of simplicity and generality V_{OUT} is not shown (because it is simply a scaled, shifted, signal-dependent replica of V_{IN}). Unless otherwise stated, any mention of the amplitude of the fed-back signal will refer to the amplitude seen at V_{IN} .

In Fig. 4 we can clearly see three main phases of operation. Initially, from 0ns to 2ns, the ringamp rapidly charges toward the dead-zone. Then, from 2ns to about 14ns it oscillates around the dead-zone region as it attempts to stabilize. By 15ns, with the output transistors M_{CP} and M_{CN} both completely cutoff, the ring amplifier reaches a steady-state solution within the dead-zone, and remains locked.




Fig. 5: Conceptual model of a ringamp during the initial slew-charging phase of operation. This model only applies to the initial charging phase and does not include the key ringamp stabilization mechanisms. $V_{OS(IN)}$ is the input-referred value of the dead-zone offset.

A. Initial Ramping

In the initial slew-charging phase of operation, the ring amplifier is functionally equivalent to the circuit of Fig. 5. The first two stages of the ring amplifier act like a pair of bi-directional continuous-time comparators that correctly select which output transistor (M_{CN} or M_{CP}) to use depending on the value of the input signal. The selected output transistor then operates as a maximally-biased current source and charges the output load with a ramp. In this initial charging phase the ringamp behaves similar to a zero-crossing based circuit [4][5].

The ramping phase ends when the input signal crosses the threshold of the comparator and the current source turns off. Due to the finite time delay of the comparator, there will be some amount of overshoot beyond the comparator threshold, which will be given by:

$$\Delta V_{overshoot} = \frac{t_d \cdot I_{RAMP}}{C_{OUT}} \quad (2)$$

where t_d is the time delay of the comparator decision, I_{RAMP} is the current supplied by the active current source, and C_{OUT} is the total loading capacitance seen at the output. This overshoot is with respect to the trip point, which will be on the boundary of the dead-zone. It will be more useful later on if we consider Eq. 1 as well, and express the input-referred overshoot with respect to the ideal settled value (the center of the dead-zone):

$$\Delta V_{init} = \frac{t_d \cdot I_{RAMP}}{\psi C_{OUT}} - \left| \frac{V_{DZ}}{2 \cdot \hat{A}_1} \right| \quad (3)$$

where ψ is the scaling factor that refers the output overshoot to the ringamp's input (and depends on feedback factor, parasitics, and feedback structure) and \hat{A}_1 is the effective gain of the first stage inverter at the end of the ramping operation (explained later).

B. Stabilization

After the initial charging ramp, the ring amplifier will begin to oscillate around the target settled value with amplitude ΔV_{init} . With no dead-zone, the structure is functionally identical to a three-inverter ring oscillator, and will continue to oscillate indefinitely (Fig. 2a). However, as the size of the dead-zone is increased, the ringamp will eventually reach an

operating condition where it is able to self-stabilize, such as in Fig. 4. If the dead-zone size is increased further still, the time required to stabilize decreases substantially, and for most practical designs, a ringamp will stabilize in only one or two periods of oscillation (i.e. Fig. 2b).

The most fundamental mechanism in the process of stabilization is the progressive reduction in the peak overdrive voltage applied to the output transistors M_{CN} and M_{CP} on each successive period of oscillation. This effect is illustrated in Fig. 4 by the progressive decrease in amplitude of the signals V_{BP} and V_{BN} . When the following relation is true, the trough (minimum value) of V_{BP} will be limited by the finite-gain of the first two stages, and begin to de-saturate from rail-to-rail operation:

$$\hat{A}_2[\hat{A}_1(\min(\tilde{V}_{IN}) - V_{CMX}) - V_{OS}] \geq V_{SS} - V_{CM} \quad (4)$$

(where \tilde{V}_{IN} is the peak-to-peak amplitude, and \hat{A}_1, \hat{A}_2 are the negative-valued effective instantaneous inverter gains). A similar relation can also be expressed for the lower signal path and V_{BN} :

$$\hat{A}_2[\hat{A}_1(\max(\tilde{V}_{IN}) - V_{CMX}) + V_{OS}] \leq V_{DD} - V_{CM} \quad (5)$$

The key point to notice in these expressions is that each signal path is being fed a different shifted replica of the oscillatory waveform generated at V_A . The upper path is given a replica where the *peaks* of the wave are *lowered* closer to the second stage inverter's threshold, and the lower path is given a replica where the *troughs* of the wave are *raised* closer to the threshold of the second stage inverter. For a sufficiently large shift in each path (V_{OS}), this creates the possibility that even for relatively large values of \tilde{V}_{IN} , finite gain effects will simultaneously limit the overdrive voltage that is applied to both M_{CP} and M_{CN} . This stands in stark contrast to the behavior of a three-inverter ring oscillator, where the decrease in V_{OV} of one output transistor necessarily means an increase in V_{OV} applied to the other.

When Eqs. 4 and 5 are true, the resulting reduction in V_{OV} applied to the output transistors M_{CN} and M_{CP} will reduce the magnitude of the output current I_{RAMP} . This decrease in output current will also cause a decrease in the amplitude of \tilde{V}_{IN} by a proportional amount, due to Eq. 3. The left sides of Eqs. 4 and 5 are therefore reduced further, and the V_{OV} 's of M_{CN} and M_{CP} will decrease even more for the next oscillation cycle. This effect will continue to feedback until the input signal amplitude becomes smaller than the input-referred value of the dead-zone, at which point the ring amplifier will stabilize and lock into the dead-zone.

If we combine Eqs. 4 and 5 and rearrange, we see that in order to trigger this progressive overdrive reduction effect, the input signal must satisfy the following relation:

$$\tilde{V}_{IN} \leq \frac{1}{\hat{A}_1} \left(\frac{V_{DD} - V_{SS}}{\hat{A}_2} - V_{DZ} \right). \quad (6)$$

Furthermore, at the beginning of the stabilization phase:

$$\tilde{V}_{IN} = 2 \cdot \Delta V_{init} \quad (7)$$

Finally, using Eqs. 2, 3, 6, and 7, we can express the stability criterion in terms of the dead-zone (i.e. settled accuracy) and the initial slew rate (i.e. speed):

$$\frac{t_d \cdot I_{RAMP}}{\psi C_{OUT}} \leq \frac{1}{2 \cdot \hat{A}_1} \left(\frac{V_{DD} - V_{SS}}{\hat{A}_2} - 2 \cdot V_{DZ} \right) \quad (8)$$

Recall once again that \hat{A}_1 and \hat{A}_2 are negative valued gains.

From this relation we see that there is a clear design tradeoff between accuracy, speed, and power. Let's assume for a moment that only t_d , I_{RAMP} , and V_{DZ} can be adjusted. To increase speed, one can either increase the initial ramp rate or decrease the time required to stabilize. Both options require sacrificing either accuracy (by increasing V_{DZ}) or power (by decreasing t_d). Likewise, to increase accuracy (by decreasing V_{DZ}), one must either decrease I_{RAMP} or decrease t_d accordingly. While these simple tradeoffs serve as a good starting point, as we will soon discover, every parameter in Eq. 8 is variable to some extent.

The discussion thus far is only a first-order model, and there are additional bandwidth, slewing, and device biasing dynamics which are not represented. Let's take a moment to evaluate this model in the form of a practical example. Consider a pseudo-differential ringamp where $A_1 = A_2 = -25 \frac{V}{V}$, $V_{DZ} = 100mV$, $V_{DD} = 1.2V$, and $V_{SS} = 0V$. By Eq. 1, the input-referred size of the dead-zone will be about 4mV, which for a 2V pk-pk input signal would ideally be accurate enough to achieve an input-referred SNDR of 54dB. By Eq. 6, the maximum allowable peak-to-peak amplitude of \tilde{V}_{IN} is approximately 6mV, and by Eq. 3, the maximum allowable input-referred overshoot at the end of the initial ramping phase must be less than 5mV.

This isn't a very encouraging result, since such a small overshoot will place a tight constraint on the parameters in Eq. 2. However, if one were to simulate this same scenario, it will turn out that the peak-to-peak amplitude of oscillation can be significantly larger than the predicted 6mV and still achieve stability. A closer look at Fig. 4 reveals an important contributor to this disparity between theory and practice. Although the AC small-signal gain of the first stage inverter, A_1 , may be $-25 \frac{V}{V}$, the effective instantaneous value

$$\hat{A}_1(t) = \frac{V_A(t)}{V_{IN}(t)} \quad (9)$$

in the actual transient waveform will be several times smaller at the beginning of stabilization. Thus, although the overall accuracy of the ringamp is determined by the final, settled, small-signal value of A_1 , the stability criterion is determined by the initial, transient, large signal effective value of A_1 . This reduction in A_1 occurs because the first stage inverter inherently operates around its trip point, where it will be slew limited. The maximum slewing current that the inverter can provide will be

$$I_{slew} = I_P - I_N \quad (10)$$

and for a square law MOSFET model, this will become:

$$I_{slew} = 2k' \frac{W}{L} \left(\frac{V_{DD} - V_{SS}}{2} \right) V_{IN} \quad (11)$$




Fig. 6: DC sweep of V_{IN} vs. I_{OUT} for a typical configuration of Fig. 1, illustrating the full input-referred characteristic near the dead-zone region. In addition to a true “dead-zone” where both output transistors are in cutoff, there is also a small boundary region “weak-zone” where the output pole location is low enough to create stability.

Notice here that the slew current is linearly related (not quadratically) to the input voltage. Thus, for the first stage inverter, slew rate limiting (and finite bandwidth) has an important impact on determining the effective value of \hat{A}_1 during stabilization (and to a lesser extent, the value of \hat{A}_2). This dynamic adjustment of the effective inverter gain is a very attractive characteristic, and improves the design tradeoff between speed, accuracy, and power by a significant factor. Similar effects also influence the operation of the second stage inverters in an additional way: although Eqs. 4 and 5 assume rail-to-rail swing for the second stage inverters when \hat{V}_{IN} is large, in reality the output swing of the second stage inverters may never completely reach rail-to-rail, regardless of the value of \hat{V}_{IN} due to slew rate limiting, finite bandwidth, and triode device operation.

The discussion of progressive overdrive reduction in this section can be conceptualized as a dynamic adjustment of the ringamp’s output pole corner frequency. The decrease in output current due to V_{OV} reduction increases the output impedance (R_o) of the ringamp, and pushes the output pole (formed by R_o and C_{LOAD}) to lower frequency. As the V_{OV} reduction effect gains momentum on each successive oscillation half-period, the output pole progressively pushes to lower and lower frequency. By the time the ringamp is locked into the dead-zone and the output transistors are in cutoff, R_o is infinite and the output pole is at DC.

C. Steady State

Thus far, we have defined the steady-state condition for a ring amplifier as the complete cutoff of both output transistors, with the input signal lying solidly within the dead-zone, such as is the case in Fig. 4. However, considering the discussion about pole adjustment in the previous paragraph, it’s clear that the ringamp can in fact be stable for a range of low frequency output pole locations down to DC. Such a situation will in practice occur often, even for a large dead-zone, since there is always a finite probability that the ring amplifier will

happen to stabilize right at the edge of the dead-zone. If that happens, one of the output transistors will still conduct a small amount of current to the output, and may never fully shut off before the amplification period ends. The existence of this stable, boundary-region “weak-zone” is illustrated in the V_{IN} vs. I_{OUT} plot of Fig. 6. The weak-zone isn’t an inherent problem for ring amplification operation, since any low-bandwidth settling will only serve to further improve accuracy. However, there are sometimes higher-level structural considerations that make it advantageous to ensure that both output transistors are completely non-conducting once settled. The design presented in Section VI is one such case, and it is there that we will explore this issue in more detail.

IV. KEY ADVANTAGES

Ring amplifiers are in many ways both structurally and functionally quite different from conventional opamps, and it is in these differences that the ringamp finds a unique advantage in the context of modern low-voltage CMOS process technologies. In this section, we will examine several of these important benefits.

A. Output Compression Immunity

In low-voltage scaled environments, kT/C noise, SNR, and power constraints will typically be dictated by the usable signal range available [2], and any practical amplification solution for scaled CMOS must therefore utilize as much of the available voltage range as possible. As it turns out, ring amplifiers are almost entirely immune to output compression, and this enables them to amplify with rail-to-rail output swing.

To understand the basis of this output compression immunity, we must consider two scenarios. First, imagine a ringamp whose dead-zone is large enough that when the ringamp is locked into the center of the dead-zone, both M_{CN} and M_{CP} will be in cutoff. In other words, when:

$$V_{DZ} \geq \left| \frac{V_{DD} - V_{SS} - 2V_T}{A_2} \right| \quad (12)$$

As a rule of thumb, this relation will usually hold for low and medium accuracy ringamps up to about 60dB. Under this scenario, M_{CN} and M_{CP} function as current sources whose linearity and small-signal gain has no appreciable effect on settled accuracy. The internal condition of the ringamp depends only on the signal at the input, and it will continue to steer toward the dead-zone until M_{CN} and M_{CP} are completely cut-off, regardless of whether they are in saturation or triode. Final settled accuracy will be governed by Eq. 1, independent of the characteristics at the output.

Now let’s consider the condition where Eq. 12 does not hold. This will occur when the dead-zone is very small, and accuracies in the 60dB to 90dB range are desired. Although other practical issues in the ringamp structure of Fig. 3 may hinder such design targets, a theoretical discussion is still quite useful in understanding the issues relevant to high accuracy ringamp topologies in general. In this scenario, the stability region of Fig. 6 is so small that the two weak-zones touch,




Fig. 7: Zoomed stabilization waveform of V_{IN} for three output swing cases: small (output near mid-rail), medium, and large (final output near the supply).

and M_{CN} and M_{CP} will still conduct a small amount once settled. The ringamp's steady state condition will essentially be that of a three stage opamp, and the open loop gain will be the product of the three stage gains. With no true dead-zone, the distortion term of Eq. 1 becomes zero, and finite loop gain will become the fundamental limitation on accuracy. At first glance, generating sufficient loop gain appears to be a problem, since the gain of M_{CN} and M_{CP} will depend on output swing (which must be as large as possible in nanoscale CMOS). Consider the case where all three stages have a gain of 25dB when operating in saturation. In the best case, the open loop gain will be 75dB, and in the worst case perhaps 50dB. Even in the best case, this seems to suggest that to build an 80dB accurate ringamp, an additional gain stage is required.

Luckily, there is another effect at play here. In the ideal square-law MOSFET model M_{CN} and M_{CP} will be in saturation when V_{OV} is less than V_{DS} . Furthermore, the small signal output impedance, r_o , is inversely proportional to the drain current, I_D . In the context of the progressive overdrive voltage reduction that occurs in ringamp stabilization, both V_{OV} and I_D will in fact trend towards zero. This implies that during steady-state, M_{CN} and M_{CP} will remain in saturation or weak-inversion even for very small values of V_{DS} , and moreover, that their gain will be enhanced by a dynamic boost in r_o . Thus, even for a nominal open loop gain of 75dB, with a wisely chosen topology it is possible to have an enhanced steady-state gain of at least 90dB, even when swinging close to the rails.

Although output swing has little effect on ringamp accuracy, it will indeed affect speed, both with respect to slewing and settling. In the initial ramping phase, the selected current source transistor will be biased with the maximum possible V_{OV} , and this guarantees that for much of the possible output range it will initially be operating in triode. As seen in Fig. 7, for settled output values near mid-rail, I_{RAMP} will be the highest and the initial ramping will be faster, but more time will be required to stabilize for the reasons discussed in Section III-B. Likewise, for values close to the rails, I_{RAMP}

will be smaller, so the initial ramping will be slower but the stabilization time will be shorter. For the most part, this works out quite nicely, since the total time required to reach steady state in each case turns out to be approximately the same. However, for extreme cases very close to the rails, the large RC time constant of the output transistor in triode operation will require a comparatively long time to reach its target value. Ultimately, it is this RC settling limitation that will usually dictate the maximum output swing possible for a given speed of operation.

B. Slew-Based Charging

Whereas a conventional opamp charges its output load with some form of RC-based settling, the output transistors M_{CN} and M_{CP} in the ring amplifier behave like digitally switched current sources, and charge the output with slew-based ramping. This is a much more efficient way to charge, since only one of the current sources in Fig. 5 will be active at a time, and the only power dissipated will be dynamic. Furthermore, during the initial ramping operation, M_{CN} or M_{CP} (whichever is selected) will be biased with the maximum V_{OV} possible for the given supply voltage. This is a major benefit, because it means that even for large capacitive loads, small transistor sizes can still produce high slew rates, and with small output transistors, the second stage inverters will be negligibly loaded by M_{CN}/M_{CP} . This effectively decouples the internal power requirements from that of the output load size, and for typical load capacitances in the femto and pico-farad range, the internal power requirements are more-or-less independent of output capacitance. This unique property stands in stark contrast to the power-loading relationship for a conventional opamp, where settling speed is typically proportional to g_m/C_{LOAD} . Even for large load capacitances, where the size of M_{CN}/M_{CP} does have an appreciable effect on the internal power requirements, the ratio of static-to-dynamic power will scale very favorably.

C. Performance Scaling with Process

In order for a technique to be truly scalable, it must meet two criteria. First, the given technique must operate efficiently in a scaled environment. This requirement has been our primary focus thus far. Second, the technique must inherently scale with advancing process technology, improving in performance simply by migrating into a newer technology. It is this second criteria that we will discuss now.

Intuitively, the ring amplifier seems like a prime candidate to benefit from process scaling, simply by its structural similarity to a ring oscillator. After all, the performance of ring oscillators track so closely with process technology that they are often used by foundries as a means of characterizing a given technology. And indeed, the stability criterion of Eq. 8 suggests this to be true. As we explored in Section IV-B, the internal power consumption of a ringamp is governed much more by inverter power-delay product and internal parasitics than the size of the output load (in stark contrast to conventional opamps). Since the power-delay product of

TABLE I: Design Requirements for Scaling Test

V_{DD}	process defined
L_{min}	process defined
SNDR	> 66dB (input-referred)
Output Range	$0.8 \cdot V_{DD}$ (per side)
Sampling Speed	$13.5/L_{min}$ MHz
Total Load	800fF (differential)
Power	minimize




Fig. 8: Simulated ring amplifier scaling trends, characterized using the simple structures of Figs. 9 and 10 and predictive technology models for nanoscale CMOS [7]. Device sizes are adjusted for each technology node as necessary in order to meet the fixed design requirements given by Table I.

an inverter decreases approximately linearly in accordance with decreasing feature size [6], the ringamp's inverter chain propagation delay, t_d , can be expected to scale according to digital process performance as well. With the relationships in Eq. 8, this reduction in t_d can be directly traded for an improvement in any of the three main design specifications: speed, accuracy, and power.

As a general rule of thumb, using small internal device sizes will tend to yield the highest power efficiencies, since doing so will not only help to minimize the power-delay product, but also minimize crowbar currents, which are a major source of static power (since the first stage inverter is always close to its trip point when settled). Considering that the time delay of a minimum sized inverter scales as a function of minimum feature size, it also follows that the conversion speed at which a ringamp most efficiently operates will be a function of the technology's minimum feature size.

To demonstrate this scalability property of ringamps, the results of a simple experiment designed to predict scaling trends are presented in Fig. 8. The setup for the test is simple: using the ring amplifier structure and pseudo-differential switched-capacitor MDAC of Figs. 9 and 10, the structure must be designed to meet the required specifications of Table I for several different CMOS technology nodes while attempting

to minimize total power consumption. Only transistor width and length resizing is allowed. The speed and accuracy targets are set at the upper-end of the given structure's practical limits, and the load capacitance is sized to be sufficient for an 11b pipelined ADC with 10b ENOB. Although the ringamp structure could be more efficient (for example, disabling itself during ϕ_s), we are only interested in the *relative* power efficiencies across process technologies, and the bare-bones structure used here illustrates the underlying trends clearly.

The predictive technology models provided by [7] were used to implement and simulate the design in the 130nm, 90nm, 65nm, 45nm, and 32nm nodes. The results in Fig. 8 indicate that the core ring amplifier structure does indeed scale according to process technology (note the logarithmic scale of the y-axis). In the upper curve, which depicts the total energy-per-cycle with respect to f_s (and thus L_{min}), the slope decreases after 45nm. This is due to the fact that the load capacitance is not scaled (for thermal noise reasons), causing dynamic energy to eventually dominate the total energy per cycle. The fundamental lower bound for dynamic power can be directly calculated for this design, since the output load is deliberately reset to mid-rail every cycle and configured such that its envelope is that of a sine wave. This dynamic power is thus calculated as

$$P_{CV^2f} = 800\text{fF} \cdot \left(\frac{0.8V_{DD}}{\sqrt{2}} \right)^2 \cdot f_s \quad (13)$$

The energy-per-cycle with respect to the total power minus the CV^2f power of Eq. 13 is shown in the lower trend line of Fig. 8. The ringamp's internal power continues to scale at pace into deep nanoscale nodes, and somewhat remarkably, the primary power contributor eventually becomes the ideal dynamic charging power itself. While this is a very encouraging result, there are many scaling effects that are not represented here that will cause the real trend to differ somewhat. In particular, the interconnect R, L, and C parasitics of the circuit are not modeled, which in upcoming nanoscale processes will become an increasingly dominant effect [8]. However, because digital circuits are also influenced by interconnect parasitics, whatever effect this has on ringamp performance will likely also affect digital circuits, and relative scaling trends will persist. The intrinsic device gains that determine the properties of A_1 , A_2 , \hat{A}_1 , \hat{A}_2 , M_{CN} , and M_{CP} also change with scaling. For the designs represented in Fig. 8, sufficient inverter gains were maintained with device dimension adjustments alone (although architectural changes are also an option). Although scaling of planar transistors to 22nm would begin to create challenges in maintaining sufficient gain in a three stage ringamp structure, newer process technologies appear to avoid this problem altogether. The good intrinsic FinFET device gains demonstrated in [1] indicate that the technologies being used in sub-32nm nodes such as FinFET and silicon-on-insulator (SOI) are able to control short-channel effects much better than previous technologies, and can produce more than sufficient intrinsic device gains for practical ringamp applications.

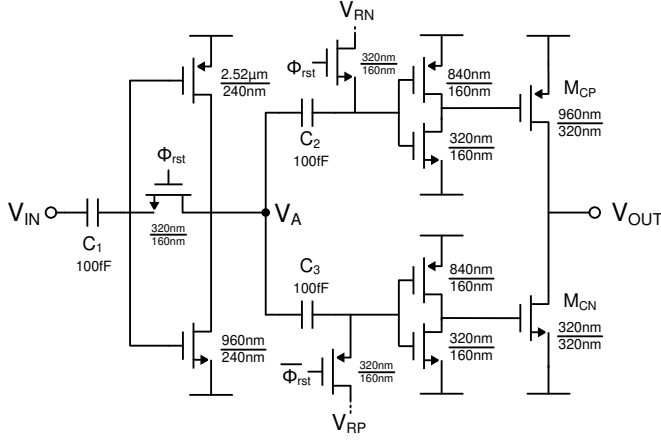


Fig. 9: Complete transistor-level ring amplifier structure of the 10.5b characterization ADC, with the actual component values used in the fabricated design.

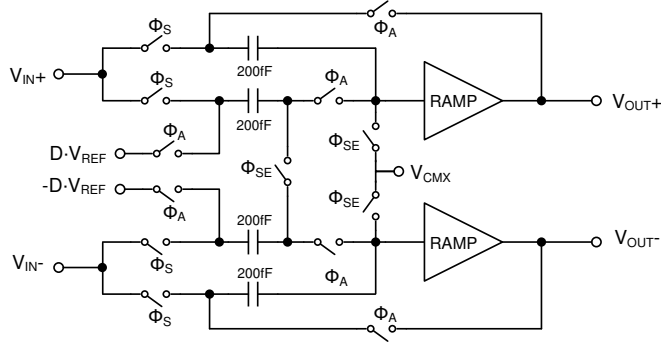


Fig. 10: Pseudo-differential float-sampled 1.5b flip-around MDAC. In this structure, the differential gain is 2 and the common-mode gain is 1, eliminating the need for additional common-mode feedback.

V. CHARACTERIZATION ADC

To demonstrate the properties that we have discussed in Sections III and IV, a 10.5-bit pipelined ADC was implemented and tested [9]. The ADC consists of nine identical 1.5-bit switched capacitor MDAC stages followed by a 1.5-bit backend flash. The transistor-level ringamp used is shown in Fig. 9. As we can see from the listed device sizes, the ringamp is quite small, and even the largest transistor in the design is only 2x the minimum W/L allowed by the process. In each stage, two of these single-ended ringamps are placed in the pseudo-differential configuration shown in Fig. 10. Due to the lack of common-mode feedback (CMFB) in such a configuration, the MDAC employs the 1.5b flip-around pseudo-differential float sampling scheme of [10]. This scheme sets the differential-mode gain to 2 and the common-mode gain to 1. Thus, any common-mode errors along the pipeline will simply add, rather than multiply and potentially saturate the common-mode level.

The pipelined ADC was fabricated in a 1P4M 0.18μm CMOS technology. A summary of performance is given in Table II. At 30MHz sampling rate, the ADC achieves 61.5dB SNDR, 61.9dB SNR, and 74.2dB SFDR. Total power




Fig. 11: Measured output spectrum of the characterization ADC for a 1MHz input tone sampled at 30Mps.

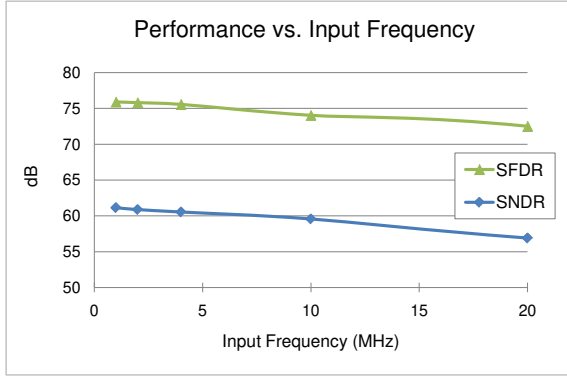
TABLE II: Characterization ADC Summary of Performance

Resolution	10.5 bits
Analog Supply	1.3 V
Sampling Rate	30 Msps
ERBW	15 MHz
Input Range	2.2 V pk-pk diff.
SNDR	61.5 dB
SNR	61.9 dB
SFDR	74.2 dB
ENOB	9.9 bits
Total Power	2.6 mW
FoM	90 fJ/c-step
Technology	0.18 μm 1P4M CMOS
Active Area	0.50 mm ² (2.00 mm x 0.25 mm)


consumption is 2.6mW, with approximately 90 μW consumed per stage by the psuedo-differential ring amplifier block. The measured ERBW is greater than 15MHz, and the Figure-of-Merit (FoM) is 90 fJ/conversion-step. Neither speed nor power were prioritized in this design (only accuracy), and with a more aggressive design and use of the power saving techniques discussed in Section VI, both can be improved significantly. The unit capacitor size for all MDAC stages is 200fF, yielding a total differential input capacitance of 800fF. The output spectrum is shown in Fig. 11, and the performance with respect to input frequency is given in Fig. 12a.

Despite such small device sizes, the actual noise contribution from the ringamp is quite small. This highlights another key benefit of ring amplifiers: because the output pole of the ringamp will be at very low frequency or even DC at the end of the amplification period, any internal noise sources will be heavily attenuated at the output.


To demonstrate the key scalability benefit of output-swing compression immunity, the analog supply was reduced to 1.2 V, the MDAC references were set to the supply voltages (0 V and 1.2 V), and the sampling frequency was reduced to 4MHz, with all other settings left unchanged. The results of an input amplitude sweep under this test setup is presented in Fig. 12b. Because the transfer function of a 1.5b MDAC spans the entire supply range, a rail-to-rail input signal will cause the ringamps to swing rail-to-rail at their outputs as well [11]. With this in mind, we see from Fig. 12b that the ringamp




(a)



(b)



(c)



(d)

Fig. 12: Measured performance and characterization data for the 10.5b ADC.




Fig. 13: Die micrograph of the 10.5b characterization ADC

maintains linearity even in true rail-to-rail operation, and only begins to degrade within $\pm 15\text{mV}$ of V_{DD} and V_{SS} due to insufficient RC settling time.

Fig. 12c shows the effect of dead-zone variation on performance. As can be seen, there is a wide, stable range of dead-zone values for which SNDR is largely unchanged. The roll-off on the right is due to Eq. 1, and the roll-off on the left is due to the ringamp becoming unstable. The SNDR plateau in the middle is a reflection of quantization and thermal noise limiting SNR in that region. The SFDR curve, by contrast, shows that distortion improves continuously to a peak value, not a plateau. The SFDR peak is somewhat compressed by the linearity of the frontend sampling switches, which are also the primary cause of the roll-off in Fig. 12a. The bootstrapped sampling switches in both this design and the design of Section VI did not operate as expected due to

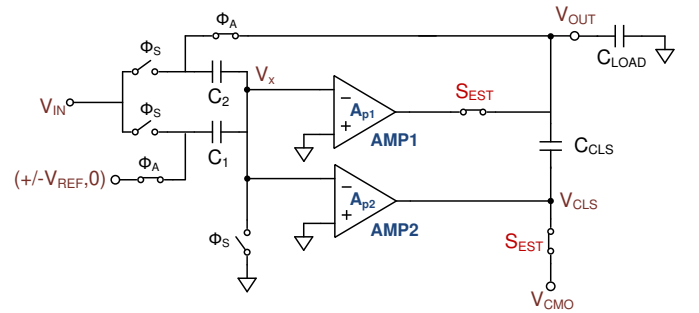


Fig. 14: A basic 1.5b MDAC implementation of the Split-CLS technique. The effective gain at the end of ϕ_A will be proportional to $A_{p1} \cdot A_{p2}$.

an unintentionally fabricated deep N-well conduction path, and caused additional distortion and power consumption. The plot of ringamp supply current consumption in the first pipeline stage with respect to applied dead-zone size in Fig. 12d indicates that the faster a ringamp stabilizes, the less power it will consume (as one would expect). On the left half of this plot, where the dead-zone is less than 0 mV, the structure becomes a ring oscillator.

VI. HIGH-RESOLUTION SPLIT-CLS ADC

In this section we will explore the design, features, and measured results of a 15-bit Split-CLS pipelined ADC. In a broader sense, this design describes a way to utilize ring

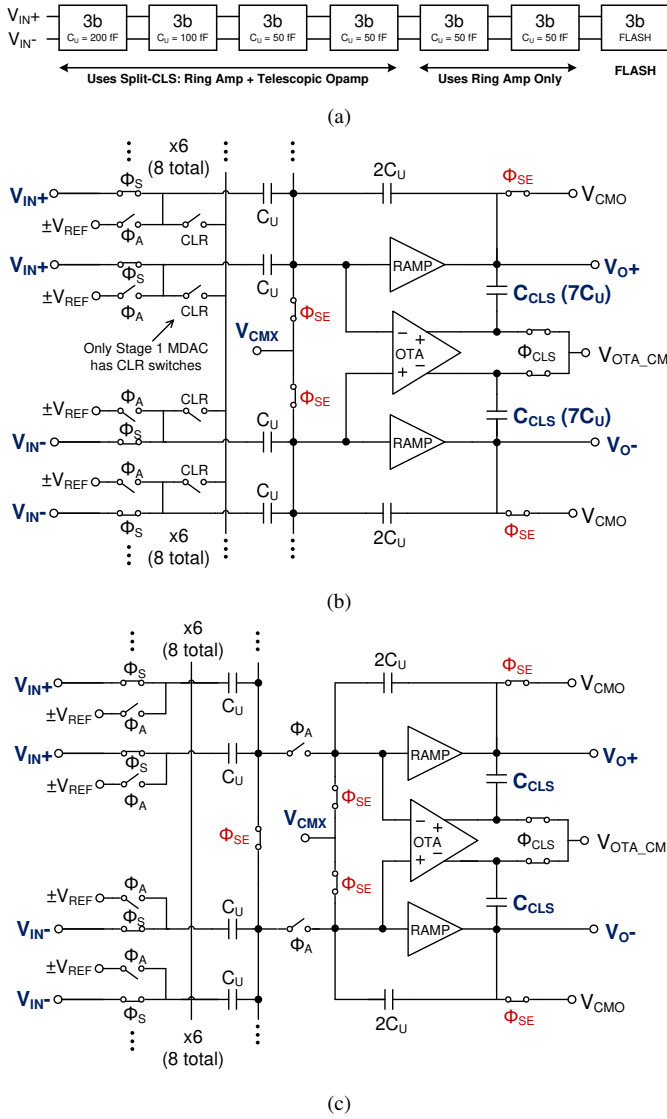


Fig. 15: Structural details of the 15b ADC. (a) Top level architecture, (b) the fully-differential Split-CLS 3b/9-level MDAC used for stage 1, and (c) the float-sampled, pseudo-differential Split-CLS MDAC used in stages 2-4. The stage 5-6 MDACs are the same as Fig. 15c except with the opamp and C_{CLS} removed.

amplifiers in the realm of high-accuracy switched-capacitor circuits.

A. Split-CLS

Beyond their usefulness as a stand-alone technique, ring amplifiers are also prime candidates to be integrated into other techniques which require an amplifier that is both fast and efficient at charging large capacitive loads. One such technique that is ideally suited for ring amplifiers is that of Split-CLS [12][13]. (In fact, it was the author's desire for a good Split-CLS coarse amplifier that inspired the idea of ring amplifiers in the first place). The basic concept of Split-CLS (split correlated level shifting) is illustrated in the single-ended switched capacitor MDAC of Fig. 14. The amplification phase, ϕ_A , is sub-divided into an estimation phase and a fine

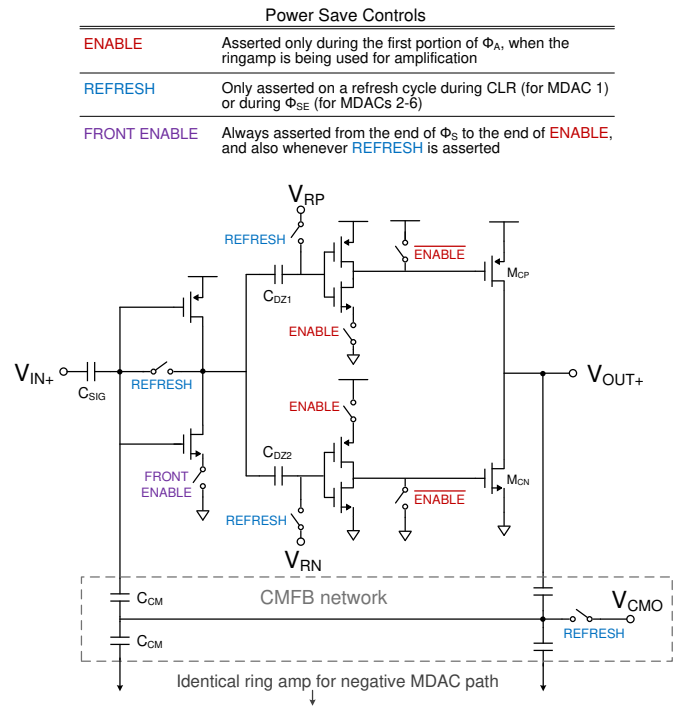


Fig. 16: Ring amplifier with power-save features and CMFB network used in the high-resolution ADC.




Fig. 17: Timing diagram for the first stage MDAC of Fig. 15b including the signals used by the ringamp of Fig. 16 and opamp of Fig. 18. The first clock period depicts a regular cycle, and the second depicts a refresh cycle. The timing for the latter MDAC stages is slightly different, with CLR removed (and ϕ_S becoming ϕ_{OFF}), and REFRESH active during all of ϕ_{SE} .

settling phase. In the first sub-phase, the switches S_{EST} are asserted, and AMP1 charges an estimate of the final settled value onto the output. Meanwhile, the capacitor C_{CLS} samples this estimation value and the output of AMP2 is shorted to mid-rail and held in standby. In the second sub-phase, the switches S_{EST} are de-asserted and AMP1 is disconnected from the output (and can be disabled). Meanwhile, AMP2 is coupled into the output via C_{CLS} and begins to fine-settle the output towards its ideal value. The key benefit derived from this operation is that by level-shifting the output of AMP2 back to mid-rail, AMP2 will process the error only, and the finite opamp gain error at its input will be substantially reduced. The final effective loop gain of the structure at the end of the amplification phase will thus be proportional to the product of

the two independent amplifier gains:

$$A_{eff} \approx \psi \cdot A_{p1} \cdot A_{p2} \quad (14)$$

where

$$\psi = \frac{C_{CLS}}{C_{CLS} + C_{LOAD} + \frac{C_1 \cdot C_2}{C_1 + C_2}} \cdot \left(\frac{C_2}{C_1 + C_2} \right)^2 \quad (15)$$

Thus, for a 55dB accurate AMP1, and a 65dB accurate AMP2, with $C_{CLS} = C_{LOAD} = 2C_1 = 2C_2$, the total effective loop gain at the end of ϕ_A will be more than 100dB. For a complete explanation of this analysis, see [14] and [15].

Consider for a moment the requirements of AMP1 and AMP2 during their respective phases of operation. AMP1 charges the output load directly, and must have a high slew rate and wide output swing. By contrast, AMP2 only processes the small error term left over from AMP1, and requires a much smaller output swing and slew rate. The key idea of Split-CLS is that by choosing amplifiers that are ideally suited for the differing requirements of AMP1 and AMP2, we can maximize the accuracy and efficiency of the overall structure. Ring amplifiers are clearly an attractive choice for AMP1, particularly in high-resolution designs where kT/C noise constraints require a relatively large loading capacitance. Due to the small output swing requirement of AMP2, a single stage telescopic opamp is a good candidate even in low supply voltages, and provides superior gain and bandwidth for a given power budget compared to other conventional opamp structures. For the cautious designer, this combination is particularly appealing - one can benefit from the advantages of ring amplifiers and still have the final settled accuracy be determined by a conventional opamp.

To demonstrate the effectiveness of this pairing, the high resolution pipelined ADC of Fig. 15a is presented. The pipeline resolves 15 bits with 6 MDAC stages and a 3b backend flash. The first four stages employ Split-CLS (and stages 5 and 6 use ring amplifiers only). The fully-differential 3-bit/9-level first stage MDAC, shown in Fig. 15b, uses two pseudo-differentially configured ringamps of Fig. 16 as the AMP1 coarse charging device and the opamp of Fig. 18 as the AMP2 fine settling device. A digitally programmable delay line controls the time allotted to the ringamp's coarse charge operation (ϕ_{CLS}), and the remainder of ϕ_A is used by the opamp. A timing diagram for the first stage MDAC is given in Fig. 17. Due to the fact that the opamp common-mode feedback (CMFB) is applied at the opamp output (and not the stage output), the pseudo-differential ring amplifiers must control the stage output's common-mode voltage. The simple capacitive CMFB network depicted in Fig. 16 provides an effective solution. The CMFB gain must be several times smaller than the gain of the primary feedback paths, and can be configured by selecting the appropriate capacitor ratio between C_{SIG} and C_{CM} . While this simple CMFB is adequate for most uses, it has inherently low common-mode accuracy, and under extreme circumstances may be insufficient. Therefore, to further relax the CMFB requirement, the stage 2-6 MDACs employ a 3-bit/9-level pseudo-differential MDAC similar to

the 1.5b flip-around MDAC of Fig. 10, but with a fixed feedback capacitor. Shown in Fig. 15c, this scheme sets the differential-mode gain to 4 and the common-mode gain to 1. Split-CLS is not needed in stages 5 and 6, where the accuracy requirement is low, and to save power the opamp and level-shifting capacitors (C_{CLS}) are removed from these stages.

B. Ringamp Power-Save Features

The ringamp shown in Fig. 16 is implemented almost identically in all stages, with only M_{CN} and M_{CP} sized differently in order to maintain the same slew rate across the differing output load sizes. For the vast majority of each period, the ringamp is unused, and can be disabled. This power-saving feature is implemented with the additional 'ENABLE' and 'FRONT ENABLE' switches depicted in Fig. 16. It is also important that the ringamp be completely disconnected from the output while the opamp is settling, and this operation guarantees that as well.

The input-offset and dead-zone voltages stored across capacitors C_{SIG} , C_{DZ1} , and C_{DZ2} must be refreshed periodically, and when doing so, the first stage inverter must be active (although the second and third stages can remain off). Therefore, during a refresh, the signals 'REFRESH' and 'FRONT ENABLE' will be asserted. These refresh operations must be done at a time when the ringamp is not in use. Although there are several options, a good choice for when to do the refresh in stage 1 is during the short window when 'CLR' is asserted and the input sampling capacitors are shorted together and cleared. At the end of this clearing pulse, the virtual node will be exactly equal to V_{CMX} , allowing an accurate value to be sampled on C_{SIG} . Performing the refresh operation while the input capacitors are sampling is also possible, but somewhat less accurate, because there will be a signal-dependent voltage drop across the virtual node sampling switch (ϕ_{SE}) that will also be sampled onto C_{SIG} . However, for the later stages this problem does not exist, since the sampled signal is completely settled at the end of ϕ_{SE} , and thus ϕ_{SE} is used as the refresh signal in these stages. The charge stored on C_{SIG} , C_{DZ1} , and C_{DZ2} will only be corrupted by small parasitic leakage currents, and it is sufficient to perform a refresh only once every N cycles (where N is controlled by an on-chip digital counter). Such an approach is used in this design, and reduces the contribution of additional static power during refresh periods down to negligible levels. For the majority of conversion cycles, the ringamp is only briefly on at the beginning of ϕ_A , and completely off during the rest of the period.

C. Charge-biased Switched Opamp

To save additional power in the opamps, the switched opamp scheme of Fig. 18 is used. The opamp is only enabled during ϕ_A (of Fig. 15), and uses the time at the beginning of ϕ_A when the ringamps are amplifying to power up. Shorting switches are placed on the source and drain nodes of the input transistors to ensure that any kickback onto the MDAC virtual node during power-on will be signal independent.

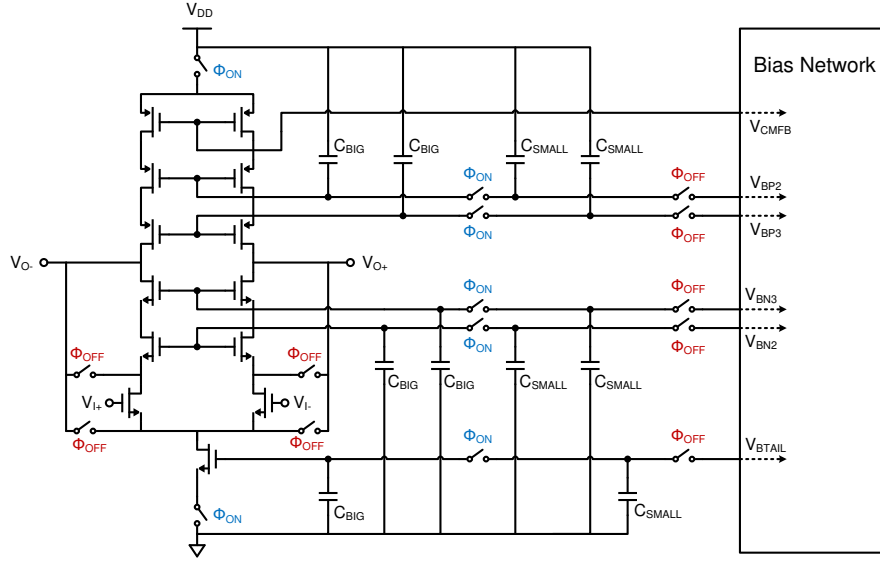


Fig. 18: The double-cascode fully differential opamp used in the first four MDACs. The charge-biased opamp switching scheme shields the bias network from power-up and power-down voltage kickback.

In the conventional switched opamp approach, the bias lines are directly connected to the opamp, and both the power-up and power-down operations will kick charge from the parasitic capacitances of the main opamp transistors onto the bias network. The bias network must then try to absorb this kickback and re-settle to the correct bias voltages, which will increase the required power-on time of the opamp [16]. Our proposed solution to this problem is to bias the opamp in charge-domain rather than voltage-domain, as shown in Fig. 18. Imagine for a moment that ϕ_{ON} is asserted, the opamp is amplifying, and the C_{BIG} capacitors hold the correct charge to properly bias the opamp. When the amplification period ends and ϕ_{OFF} asserts, the opamp will switch off and a large amount of voltage will kickback onto C_{BIG} . However, since the top-plate of C_{BIG} has no DC path to ground during ϕ_{OFF} , the charge stored on it is trapped. Meanwhile, a set of small capacitors (C_{SMALL}) sample the bias network. Then, when the opamp is powered on again, C_{BIG} will be shorted to C_{SMALL} and some amount of charge transfer will occur. Since this charge transfer occurs during ϕ_{ON} , there is no unwanted voltage kickback present, and C_{SMALL} will update C_{BIG} with an incremental piece of charge corresponding to the correct bias voltage. This bucket-brigade that C_{SMALL} provides between the bias network and C_{BIG} isolates the bias network from opamp kickback and enables rapid power-up.

D. Measurement Results

The 15-bit pipelined ADC was fabricated in a 1P4M 0.18 μm CMOS process. A summary of performance is given in Table III. At 20MHz sampling rate it achieves 76.8 dB SNDR (12.5 ENOB), 77.2 dB SNR and 95.4 dB SFDR, consuming 5.1mW. The ERBW is found to be above 10MHz, which results in a Figure-of-Merit of 45 fJ/conv-step. The MDAC references are set at 25 mV and 1275 mV, allowing the input signal to utilize 96% of the available supply range.

TABLE III: Split-CLS ADC Summary of Performance

Resolution	15 bits
Analog Supply	1.3 V
Sampling Rate	20 Msps
ERBW	10 MHz
Input Range	2.5 V pk-pk diff.
SNDR	76.8 dB
SNR	77.2 dB
SFDR	95.4 dB
ENOB	12.5 bits
Total Power	5.1 mW
FoM	45 fJ/c-step
Technology	0.18 μm 1P4M CMOS
Active Area	1.98 mm ² (3.05 mm x 0.65 mm)




Fig. 19: Measured output spectrum of the high-resolution Split-CLS ADC for a 1MHz input tone sampled at 20Msps.

Capacitor matching was good enough that no digital calibration was needed. The first stage MDAC unit capacitance is 200 fF, giving a total differential input capacitance of 3.2 pF (or 3.76 pF when the Flash sub-ADC capacitors are included).

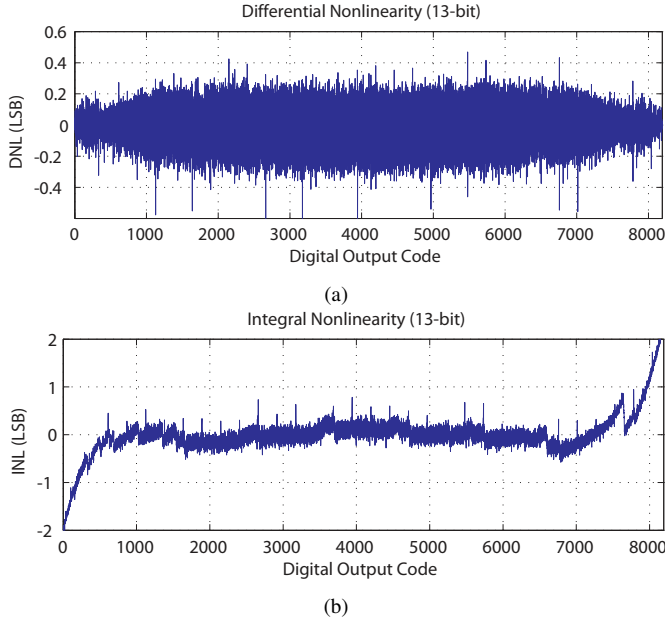


Fig. 20: Measured DNL (top) and INL (bottom) normalized to $\text{LSB}=2^{-13}$.

Total accuracy is fundamentally limited by noise, and the SNR is almost exactly equal to the kT/C noise limit predicted for this design. Due to jitter injected by the off-chip clock source and internal clock generator, noise also becomes a limitation with respect to input frequency, as seen in Fig 21a. A relatively conservative design target of 20 Msps was chosen in order to manage risk and ensure that the key features of accuracy and power efficiency would be demonstrated, and with a few modifications higher speeds are certainly possible.

Stage scaling is performed for the first three stages, with the unit capacitance for each stage listed in Fig. 15a. To minimize power, the 3b/9-level flash sub-ADCs of stages 2-7 use an architecture similar to [17], where the only load placed on the previous stage's output will be due to the gate capacitance of the 8 sub-ADC latches. In order to minimize aperture and skew related errors for fast moving inputs, the first stage's sub-ADC uses the opposite approach, in which the moving input signal is first sampled, then subtracted from the reference and quickly latched.

All analog portions of the circuit operate at 1.3 V, including the ringamps and opamps. Opamp switching is measured to reduce the total opamp power by 35%, and improve SNDR by 0.6 dB. An unfortunate issue with the physical layout caused a node in the bootstrapped input sampling switches to be shorted into a deep-nwell substrate, causing them to operate as NMOS-only statically biased switches. In order to compensate, the digital and switch supply had to be operated above 1.3 V. Based on follow-up simulations and measurements, this issue is estimated to have increased total power consumption by approximately 23%, and is included in the 5.1 mW reported. The impact on performance can be seen in the end-code non-linearity in Fig. 20b, and in the roll-off above -0.6 dBFS in Fig. 21b. For this reason, the output spectrum of Fig. 19 and

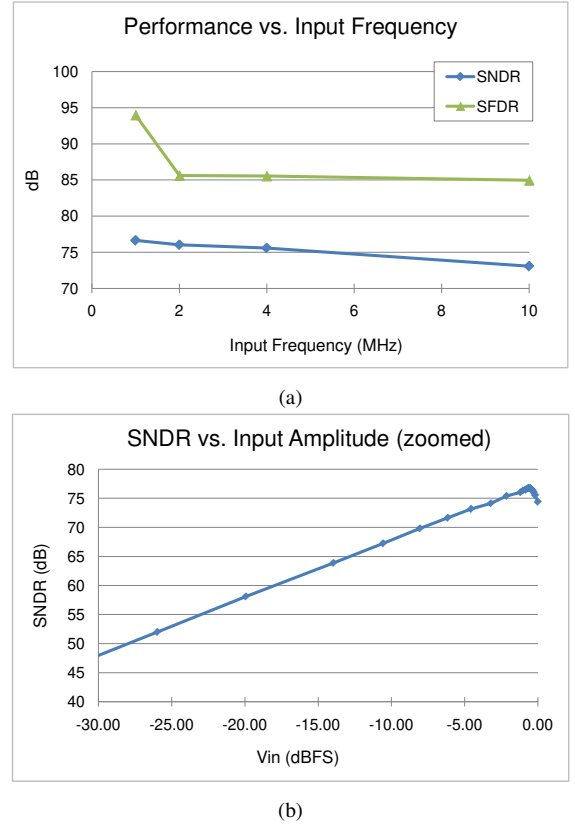


Fig. 21: Measured input signal performance data for the high-resolution Split-CLS ADC.

results in Table III are collected for a -0.6 dBFS input signal.

When the sampling frequency is increased to the point that the opamps never have a chance to amplify, the contribution of the ringamps to overall accuracy can be measured, and is found to be 55 dB SNDR at $f_s = 80$ MHz. This result is also confirmed by the result of Fig. 22d, where the digitally controlled timing generator that sets the ringamp amplification window is swept with the opamps disabled. What we find is that ringamp performance decreases logarithmically with respect to incomplete settling time, and that the ringamp settles to about 55 dB SNDR in roughly 6 ns.

The ADC exhibits a high tolerance to dead-zone variation, as shown in Figs. 22a and 22b. In Fig. 22a, the dead-zone applied to the ringamps in the 1st stage MDAC is swept with everything else held constant. For the most part, this plot tells the same story as Fig. 12c of Section V, with the roll-off above +50 mV due to the dead-zone becoming large and increasing distortion and the roll-off below -50 mV due to the ringamps going unstable. The curve's peak has a systematic offset of roughly -25 mV, and is caused by the asymmetry of the 'ENABLE' switches that control the second stage inverters of Fig. 16. The true effective dead-zone of the ringamp is the sum of V_{DZ} plus any inherent systematic offsets, including this asymmetry. The opamp, via Split-CLS, helps to absorb small errors and further flattens the SNDR curve of Fig. 22a into a wide stable plateau where kT/C noise is the fundamental




Fig. 22: Measured ringamp sensitivity and characterization data for the high-resolution Split-CLS ADC.

limitation of SNDR. The SFDR curve contains a plateau region as well, which occurs when distortion effects other than finite gain begin to dominate, such as input-switch linearity. When the ring amplifier supply voltage is swept with all other voltages and biases held constant (Fig. 22b), SNDR and SFDR are virtually unchanged. At 1275 mV the supply voltage equals the positive MDAC reference, and explains the roll-off seen below that point.

The benefit of the ringamp power-save feature is demonstrated in Fig. 22c. Total ringamp power consumption is reduced by about $4.5\times$ when the ringamps are only refreshed once every 64 cycles. For the ADC presented in Section V, switching the ringamps off during ϕ_S could have reduced ringamp power by almost 50%.

VII. CONCLUSION

Looking back at the individual benefits of a ring amplifier we find many parallels with other pre-existing techniques: zero-crossing based circuits also have the ability to charge efficiently and decouple internal power from external load requirements [4][5], dynamic current adjustment is a common attribute of class-AB amplifiers in general [18], using simple inverter-based amplifiers has long been a topic of interest [10][19][20][21], and CLS can already be used to achieve rail-to-rail output swing [15]. Ultimately, it is both logical and encouraging that we should find such similarities. Each




Fig. 23: Die micrograph of the active area of the high-resolution Split-CLS ADC.

one of these parallels represent a specific trait desired in a scalable amplifier topology, and the fact that a ring amplifier encompasses all of them indicates that it is a convergence point for many lines of research in scalable amplification. In other words, the unique aspect of a ring amplifier lies not in the traits that it possesses, but rather in the particular solution by which it is able to embody them all.

In this paper we have laid out the basic theory of ring amplification, discussed the key benefits that it brings to scaled CMOS, and studied two examples of its practical implementation. We have primarily considered the specific case of a medium-accuracy ring amplifier structure suited for use in switched capacitor circuits. As demonstrated in this

paper, this simple ringamp incarnation is highly effective, and can be integrated into more complex topologies, such as Split-CLS, to cover a range of speed and accuracy targets. And yet, the concepts of ring amplification are applicable to a broader range of structural approaches, and what we have examined here is simply the foundation for a deeper exploration of the topic.

ACKNOWLEDGMENT

The authors would like to thank Asahi Kasei Microdevices for supporting this research.

REFERENCES

- [1] M. Shrivastava, R. Mehta, S. Gupta, N. Agrawal, M. Baghini, D. Sharma, T. Schulz, K. Arnim, W. Molzer, H. Gossner, and V. Rao, "Toward system on chip (soc) development using finfet technology: Challenges, solutions, process co-development and optimization guidelines," *Electron Devices, IEEE Transactions on*, vol. 58, no. 6, pp. 1597–1607, june 2011.
- [2] B. Jonsson, "On cmos scaling and a/d-converter performance," in *NORCHIP, 2010*, nov. 2010, pp. 1–4.
- [3] B. Murmann. (2012) Adc performance survey 1997-2012. [Online]. Available: <http://www.stanford.edu/~murmman/adcsurvey.html>
- [4] J. Fiorenza, T. Sepke, P. Holloway, C. Sodini, and H.-S. Lee, "Comparator-based switched-capacitor circuits for scaled cmos technologies," *Solid-State Circuits, IEEE Journal of*, vol. 41, no. 12, pp. 2658–2668, dec. 2006.
- [5] L. Brooks and H.-S. Lee, "A 12b, 50 ms/s, fully differential zero-crossing based pipelined adc," *Solid-State Circuits, IEEE Journal of*, vol. 44, no. 12, pp. 3329–3343, dec. 2009.
- [6] H. Iwai, "Cmos scaling towards its limits," in *Solid-State and Integrated Circuit Technology, 1998. Proceedings. 1998 5th International Conference on*, 1998, pp. 31–34.
- [7] Y. Cao. (2012) Predictive technology model (ptm). [Online]. Available: <http://ptm.asu.edu/>
- [8] S. Thompson, R. Chau, T. Ghani, K. Mistry, S. Tyagi, and M. Bohr, "In search of "forever," continued transistor scaling one new material at a time," *Semiconductor Manufacturing, IEEE Transactions on*, vol. 18, no. 1, pp. 26–36, feb. 2005.
- [9] B. Hershberg, S. Weaver, K. Sobue, S. Takeuchi, K. Hamashita, and U. Moon, "A 61.5db sndr pipelined adc using simple highly-scalable ring amplifiers," *VLSI Circuits (VLSIC), 2012 Symposium on*, june 2012.
- [10] J. Li and U. Moon, "A 1.8-v 67-mw 10-bit 100-ms/s pipelined adc using time-shifted cds technique," *Solid-State Circuits, IEEE Journal of*, vol. 39, no. 9, pp. 1468–1476, sept. 2004.
- [11] I. Mehr and L. Singer, "A 55-mw, 10-bit, 40-msample/s nyquist-rate cmos adc," *Solid-State Circuits, IEEE Journal of*, vol. 35, no. 3, pp. 318–325, mar 2000.
- [12] B. Hershberg, S. Weaver, and U. Moon, "Design of a split-cls pipelined adc with full signal swing using an accurate but fractional signal swing opamp," *Solid-State Circuits, IEEE Journal of*, vol. 45, no. 12, pp. 2620–2630, dec. 2010.
- [13] B. Hershberg, S. Weaver, and U. Moon, "A 1.4v signal swing hybrid cls-opamp/zcbc pipelined adc using a 300mv output swing opamp," *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2010 IEEE International*, pp. 302–303, feb. 2010.
- [14] B. Hershberg, T. Musah, S. Weaver, and U. Moon, "The effect of correlated level shifting on noise performance in switched capacitor circuits," in *Circuits and Systems (ISCAS), 2012 IEEE International Symposium on*, may 2012.
- [15] B. 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," *Solid-State Circuits, IEEE Journal of*, vol. 43, no. 12, pp. 2620–2630, dec. 2008.
- [16] H.-C. Kim, D.-K. Jeong, and W. Kim, "A partially switched-opamp technique for high-speed low-power pipelined analog-to-digital converters," *Circuits and Systems I: Regular Papers, IEEE Transactions on*, vol. 53, no. 4, pp. 795–801, april 2006.
- [17] S. Lewis and P. Gray, "A pipelined 5-msample/s 9-bit analog-to-digital converter," *Solid-State Circuits, IEEE Journal of*, vol. 22, no. 6, pp. 954–961, dec 1987.
- [18] P. R. Gray, P. J. Hurst, S. H. Lewis, and R. G. Meyer, *Analysis and Design of Analog Integrated Circuits*. John Wiley And Sons, 2001.
- [19] F. Krummenacher, E. Vittoz, and M. Degrauwe, "Class ab cmos amplifier micropower sc filters," *Electronics Letters*, vol. 17, no. 13, pp. 433–435, 25 1981.
- [20] M. Figueiredo, R. Santos-Tavares, E. Santin, J. Ferreira, G. Evans, and J. Goes, "A two-stage fully differential inverter-based self-biased cmos amplifier with high efficiency," *Circuits and Systems I: Regular Papers, IEEE Transactions on*, vol. 58, no. 7, pp. 1591–1603, july 2011.
- [21] Y. Chae and G. Han, "Low voltage, low power, inverter-based switched-capacitor delta-sigma modulator," *Solid-State Circuits, IEEE Journal of*, vol. 44, no. 2, pp. 458–472, feb. 2009.